OXFORD ECONOMICS AUSTRALIA

DATA CENTRE ENERGY DEMAND

FINAL REPORT

JULY 2025

ABOUT OXFORD ECONOMICS AUSTRALIA

Oxford Economics Australia is a leading economic advisory firm. Following the acquisition of BIS Shrapnel in 2017, Oxford Economics Australia have unparalleled capabilities in helping clients to navigate local economic issues in the context of the global trends shaping the world economy.

Oxford Economics, the parent company, was founded in 1981 as a commercial venture with Oxford University's business college to provide economic forecasting and modelling to UK companies and financial institutions expanding abroad. Since then, we have become one of the world's foremost independent global advisory firms, providing reports, forecasts and analytical tools on more than 200 countries, 100 industries, and 8,000 cities and regions. Our best-in-class global economic and industry models and analytical tools give us an unparalleled ability to forecast external market trends and assess their economic, social and business impact.

Headquartered in Oxford, England, with regional centres in New York, London, Frankfurt, and Singapore, Oxford Economics has offices across the globe in Belfast, Boston, Cape Town, Chicago, Dubai, Dublin, Hong Kong, Los Angeles, Mexico City, Milan, Paris, Philadelphia, Stockholm, Sydney, Tokyo, and Toronto. We employ 450 staff, including more than 300 professional economists, industry experts, and business editors—one of the largest teams of macroeconomists and thought leadership specialists. Our global team is highly skilled in a full range of research techniques and thought leadership capabilities from econometric modelling, scenario framing, and economic impact analysis to market surveys, case studies, expert panels, and web analytics.

Oxford Economics is a key adviser to corporate, financial and government decision-makers and thought leaders. Our worldwide client base now comprises over 2,000 international organisations, including leading multinational companies and financial institutions; key government bodies and trade associations; and top universities, consultancies, and think tanks.

July 2025

The modelling and results presented here are based on information provided by third parties, upon which Oxford Economics Australia has relied in producing its report and forecasts in good faith. Any subsequent revision or update of those data will affect the assessments and projections shown.

To discuss the report further please contact:

Alex Hooper

Associate Director, Oxford Economics Australia

ahooper@oxfordeconomics.com

TABLE OF CONTENTS

1. Executive summary 3

2. Introduction 6

3. Modelling framework 10

3.1 Project approach 11

3.2 Economic approach 12

4. Forecast results 14

4.1 Headline results 14

4.2 Sydney deep dive 19

4.3 Melbourne deep dive 21

4.4 AI deep dive 23

4.5 Scenario results 24

5. Sensitivity analysis 26

5.1 Regional diversification sensitivity 26

5.2 Rapid AI uptake sensitivity 27

5.3 Global market share sensitivity 28

5.4 Additional pipeline sensitivity 29

5.5 Key risks and uncertainties 30

6. Technical appendix 32

We would like to thank the members of our expert panel, as well as all industry participants, who generously contributed their time and insights to support this research.

1. EXECUTIVE SUMMARY

This report presents Oxford Economics Australia's independent forecasts of electricity consumption by data centres in Australia, developed for the Australian Energy Market Operator (AEMO) to inform the 2025 Electricity Statement of Opportunities and the 2026 Integrated System Plan. It provides a forward-looking assessment of how data centre energy demand is likely to evolve under AEMO's scenarios.

The forecasts are based on a blended approach combining two independently constructed methodologies.

- The *Project Approach* uses project-level data from AEMO, Network Service Providers (NSPs), and industry to estimate electricity consumption of known existing and prospective data centre developments.
- The *Economic Approach* models the underlying demand for data centre services by sector, linking economic activity, technological adoption, and energy use.

This dual-framework captures both the physical pipeline of development as well as the macroeconomic demand-side drivers, including those associated with AI adoption and digital transformation. This approach allows for robust cross-validation and scenario testing under conditions of structural change and market uncertainty.

In FY25, Australian data centres are estimated to have consumed 3.9TWh of electricity. The NEM made up 98% of this and represented approximately 2% of NEM grid-supplied consumption. Under the *Step Change* scenario, data centre consumption in Australia is forecast to grow at an average annual rate of 25.1% to reach 12.0 TWh by FY30 and 34.5 TWh by FY50. This trajectory underscores the increasing importance of data centres as a structurally significant electricity load within Australia's energy system. Of the 34.5 TWh, 33.8 TWh of this is forecast to be on the NEM, representing 12% of NEM grid-supplied consumption by FY50.

Under the *Green Energy Industries* scenario, Australia-wide data centre demand reaches 13.1 TWh by FY30 (+9.1% relative to *Step Change*), while under the *Progressive Change* scenario, demand is materially lower at 8.5 TWh (−29.7%).

Under AEMO's *Step Change* scenario, Australia's existing stock of data centres is expected to consume 8.5 TWh of energy by FY30, up from 3.9 TWh today. These existing data centres are expected to grow at 7.8% on average to FY41 within their existing rated capacity due to increased computation. At FY41 all existing data centres are assumed to have reached their mature load.[1]

Prospective projects included in the *Step Change* scenario total 8 GW of rated capacity and are expected to consume 8.1 TWh of energy by FY30 and 33.3 TWh by FY50. Hyperscale data centres are

[1] Rated capacity refers to the maximum demand level specified in a connection agreement between the data centre and a network service provider.

driving capacity expansion. By FY40, co-locators will consume 50%, hyperscalers 47% with edge, telco and other data centre types accounting for the remainder.

Fig. 1. Data centre energy consumption in Australia (LHS) and NEM data centre energy consumption as a share of grid-supplied electricity (RHS) by scenario, Australia

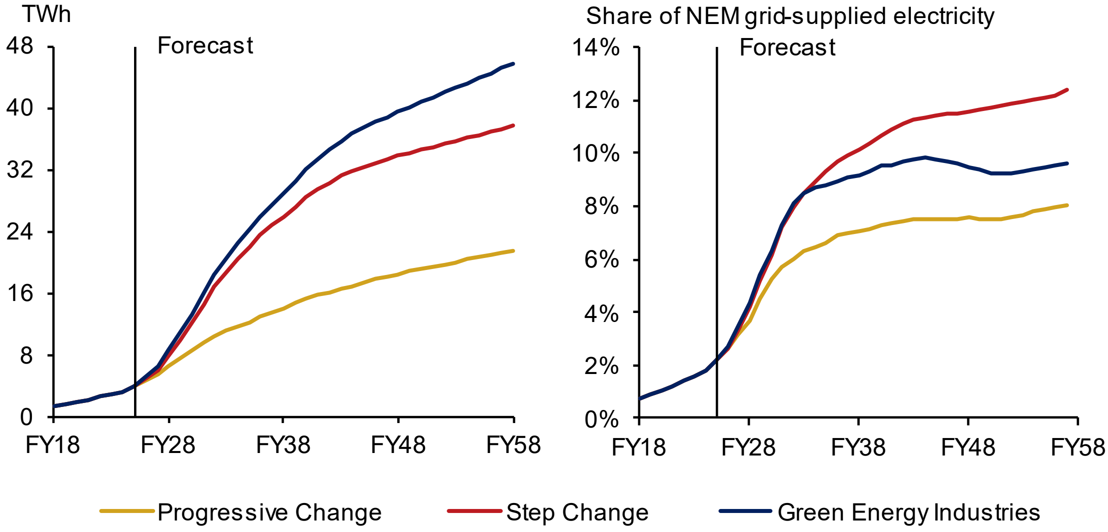

Source: Oxford Economics Australia based on AEMO data.
Note: The 'share of NEM grid-supplied electricity' is equal to OEA's estimates of 'NEM data centre energy consumption' as a share of 'NEM operational sent-out consumption' from the 2025 Electricity Statement of Opportunities.

Sydney remains the dominant hub for data centre activity, with a larger and more mature pipeline of existing and committed data centre projects compared to Melbourne, positioning it to meet near-term demand more reliably. However, Melbourne is positioned to absorb any upside in demand due to the significant scale of its prospective project pipeline.

Fig. 2. Existing, committed & anticipated project consumption, Sydney & Melbourne

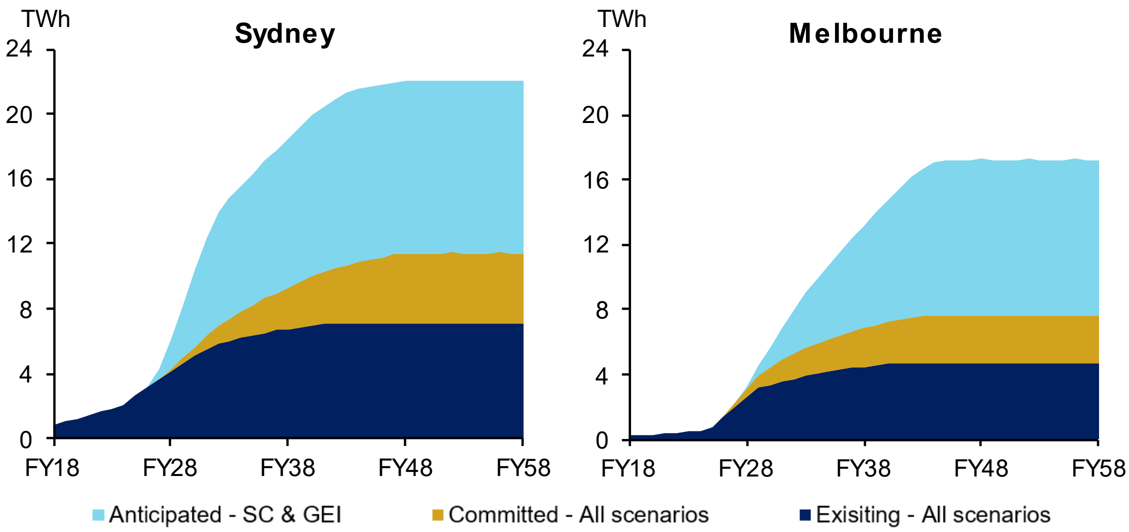

Source: Oxford Economics Australia based on AEMO data.

Melbourne is expected to significantly increase its share of data centre related energy consumption from 20% currently to 30% by FY30 and stabilise around 41% in the long run. Activity outside of Sydney and Melbourne is fairly minimal and concentrated around South Australia (1.2 TWh by FY50), Southern New South Wales and Canberra (0.8 TWh), and Western Australia (0.6 TWh).[2]

Currently only 1.5% of data centre electricity consumption is connected to the transmission network. This share is expected to grow rapidly to 32% by FY30 (2.0 TWh), driven by the progression of data centre developments in Melbourne, which account for the vast majority of prospective transmission-connected projects.

Fig. 3. Data energy consumption splits, *Step Change*

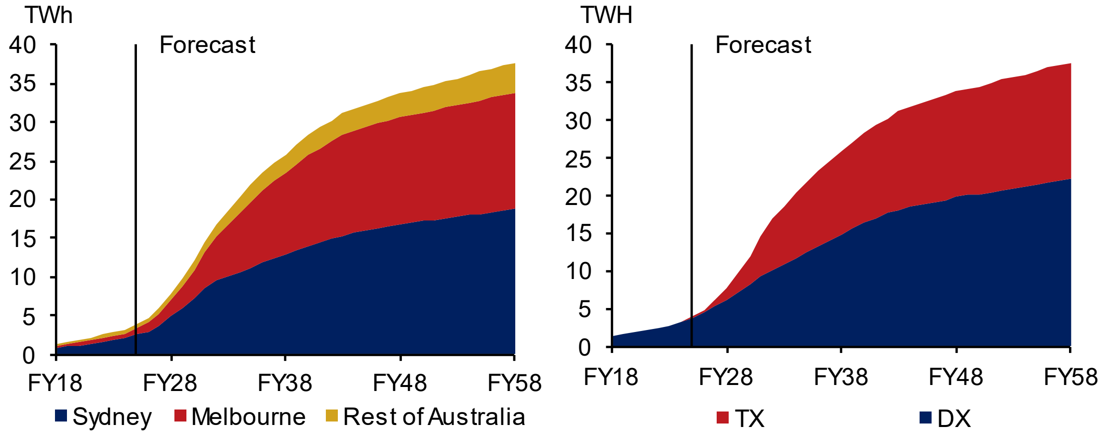

Source: Oxford Economics Australia based on AEMO data.

Forecasting electricity demand from data centres over a multi-decade horizon involves a high degree of uncertainty. On the upside, rapid AI uptake could lead to faster than expected demand realisation, compressing ramp-up periods and bringing forward peak loads. If a significant share of anticipated projects proceed, demand could significantly exceed current estimates. On the downside, project deferrals, efficiency gains, or slower digital adoption could lead to flatter demand, particularly beyond 2030.

[2] Note that anticipated projects are not included under the *Progressive Change* (PC) scenario. NSP survey respondents in Melbourne differentiated between anticipated project consumption under *Step Change* (SC) and *Green Energy Industries* (GEI) scenarios. In the Sydney survey, respondents did not.

2. INTRODUCTION

This report provides forecasts of the energy consumed by data centres at a sub-region level by scenario and was prepared for AEMO's forecasting and planning publications in 2025 and 2026.

Chapter three outlines the modelling framework used to generate the demand forecasts. It presents two approaches: a *Project Approach*, which captures industry expectations based on current and proposed data centre developments, and an *Economic Approach*, which estimates demand by analysing the broader economy's consumption of data centre services and the associated energy requirements.

Chapter four presents national electricity demand under the central *Step Change* scenario for data centres alongside more granular analysis of the of major data centre sub-regions. The chapter also examines the potential energy impact of the first wave of mass AI adoption. Finally, it considers the implications for national electricity demand under *Progressive Change* and *Green Energy Industries*.

Chapter five presents a series of sensitivity analyses to understand the implications of alternative assumptions for electricity demand. These include a *Regional Diversification* sensitivity, which reallocates data centre demand based on proxies such as professional workforce size, international fibre-optic connectivity, and population rather than observed market activity. A *Rapid AI Uptake* sensitivity explores the implications of accelerated demand growth driven by widespread AI adoption and shorter ramp-up periods. A *Global Demand* sensitivity considers Australia's potential share of Asia Pacific data centre expansion under high-growth global conditions. Finally, an *Additional Pipeline* sensitivity examines how outcomes under the *Step Change* scenario would evolve if a greater number of projects were assumed to proceed.

A data centre is a specialised facility that houses servers used to store, process, and deliver information. These centres support a wide range of digital services that households and businesses rely on every day – such as streaming, banking, cloud storage, and communications. They are essential infrastructure for the functioning of modern digital life and a modern economy.

Rather than each business or user hosting their own systems, data centres provide a centralised and secure platform that supports the delivery of digital services at scale. The key components of a data centre and their associated energy use are summarised well by the International Energy Agency.[3]

- **Servers** are computers that process and store data. They can be equipped with central processing units (CPUs) and specialised accelerators such as graphics processing units (GPUs). On average they account for around 60% of electricity demand in modern data centres, although this varies greatly between data centre types.
- **Storage systems** are devices used for centralised data storage and backup, and account for around 5% of electricity consumption.

[3] Energy and AI (2025) *International Energy Agency*. Available here

- **Networking equipment** include switches to connect devices within the data centre, routers to direct traffic and load balancers to optimise performance. Networking equipment accounts for up to 5% of electricity demand.
- **Cooling and environmental control** refers to equipment that regulates temperature and humidity to keep IT equipment operating at optimal conditions. The share of cooling systems in total data centre consumption varies from about 7% for efficient hyperscale data centres to over 30% for less-efficient enterprise data centres.
- **Uninterruptible power supply (UPS) batteries and backup power generators** are there to keep the data centre powered during outages. Both UPS and backup generators are rarely used, but necessary to ensure the extremely high levels of reliability that data centres must meet.
- **Other infrastructure**, such as lighting and office equipment for on-site staff, etc.

Fig. 4. Data centre components

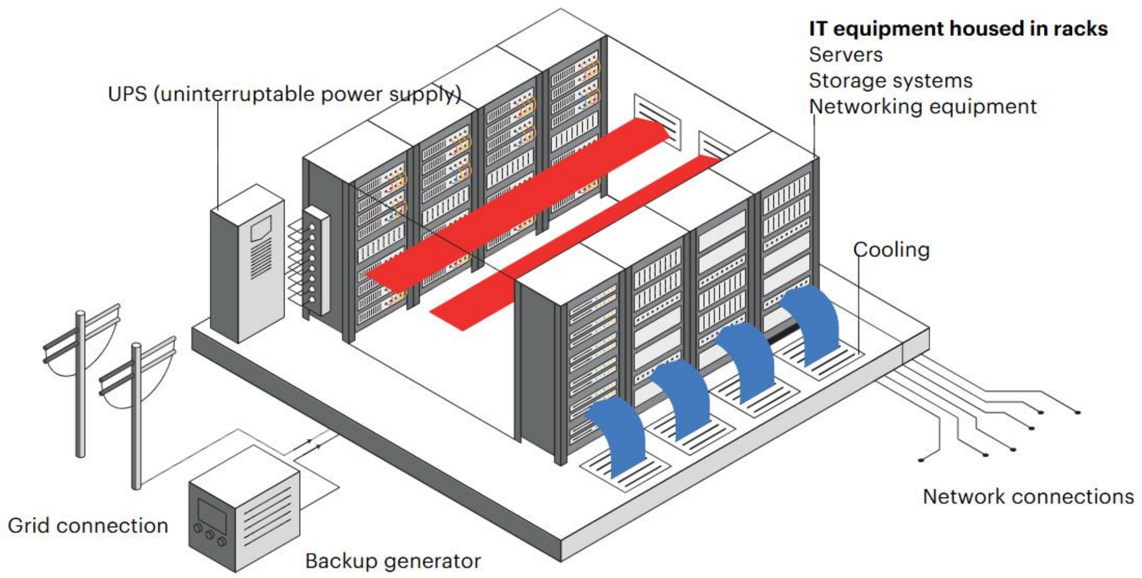

Source: Reproduced from International Energy Agency’s Energy & AI report (2025)

The share of these different components in data centre electricity consumption varies greatly by type of data centre, depending on the nature and efficiency of the equipment they have installed. For the purpose of this report, data centres are grouped into three distinct types:

1. **Hyperscale** data centres are large, high-capacity facilities designed to deliver computing, storage, and network services at scale. They are typically developed to support the operations of a single organisation, enabling highly centralised control over cloud platforms, artificial intelligence systems, and large-scale digital services. These centres are engineered for efficiency, automation, and rapid expansion to meet future demand for digital services. For example, Microsoft operates hyperscale data centres to deliver its global Azure cloud

platform, Google uses them to power services such as Search, YouTube, and Google Cloud, and OpenAI relies on hyperscale infrastructure to train and deploy AI models.

2. **Co-locator** data centres provide data centre services to multiple tenants. Tenants lease physical space, power, cooling, and network connectivity to house their IT equipment. Some co-locators also provide physical or virtual servers so the tenant does not need to supply and maintain their own hardware. Hyperscalers are also often tenants of co-locators.
3. **Other** data centres include edge, telecommunication & enterprise facilities. They are typically smaller and located closer to end-users to provide services that require fast response times, such as content delivery, mobile networks, and Internet of Things applications. They represent a small share of Australia's existing and prospective data centre development.

The rapid rise of digitalisation and cloud technologies over the last decade has led to an exponential growth in the demand for data centres' processing capacity. Simultaneously, the increasing demand for computational power was being offset by soaring increases in energy efficiency. In the decade to 2017, the number of computations that could be processed with a watt of energy grew thirteenfold.[4] However, despite computational energy efficiency continuing to grow at approximately 37% p.a., the energy consumption of data centres in Australia and globally[5] has outstripped these gains (Figure 5).

Fig. 5. Australian data centre energy consumption & server energy efficiency

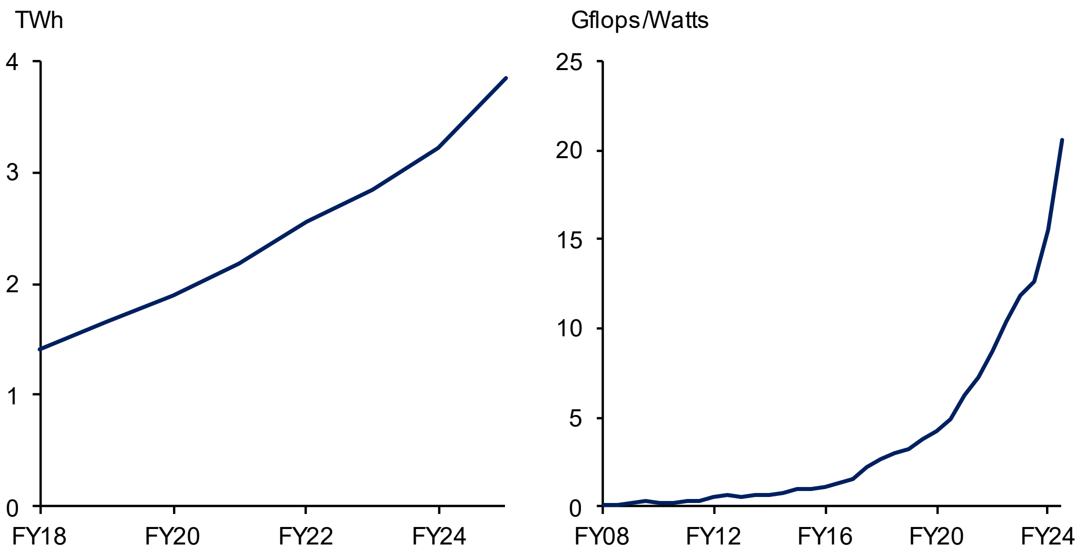

Source: Oxford Economics Australia based on AEMO data, Top 500.
Note: Gflop (giga floating point operations per second) is a measure of a computer's performance, representing one billion floating-point calculations completed per second. Gflops/Watts indicates the energy efficiency of a system, showing how much computational work it can perform for each watt of power consumed.

The majority of Australia's data centres are located in Sydney, with Melbourne acting as a secondary hub. Sydney has added 1,820 MW of rated capacity over the last 10 years, an average growth rate of 14% per year. Melbourne has also exhibited strong growth of 11% to add 930 MW of rated capacity. Rated capacity outside the two major cities has grown at a modest 4% p.a. and the share of

[4] The Green List (2024) *Top 500*. Available here
[5] Energy and AI (2025) *International Energy Agency*. Available here

consumption has subsequently decreased over the last few years as the pace of development has picked up in Sydney and Melbourne.

Fig. 6. Data centre energy consumption by region

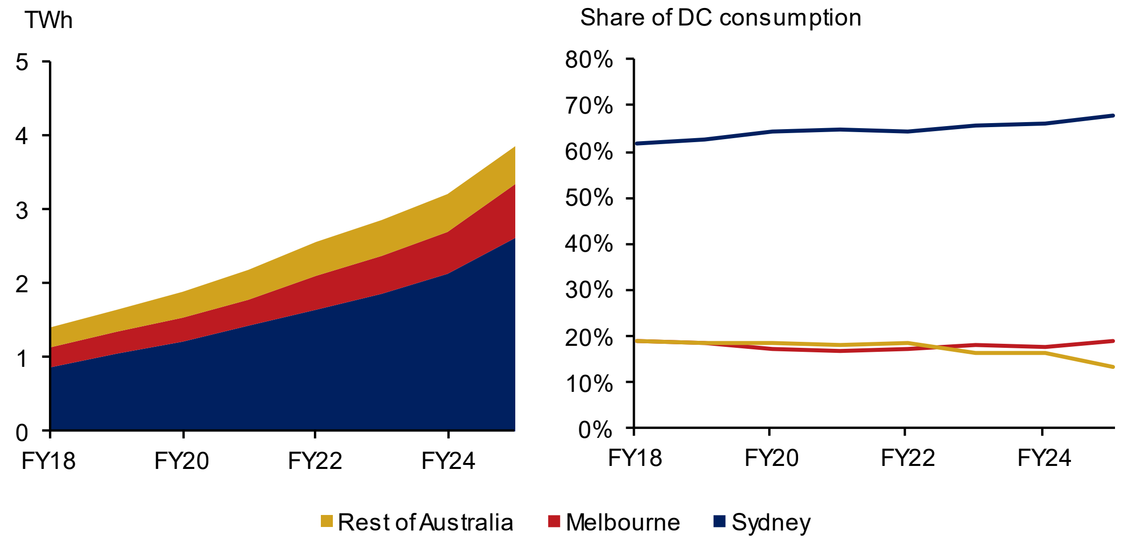

Source: Oxford Economics Australia based on AEMO data.

Strong growth in the number and rated capacity of data centre developments has led to 15% p.a. growth in the energy that data centres use from the grid. In FY18 data centres consumed 1.4 TWh rising to 3.9 TWh in FY25 over 98% of which is connected the distribution network. The swift increase has had implications for the grid with data centres consuming an increasing share of grid-supplied electricity, with NEM connected data centre consumption increasing from 0.9% of total NEM grid-supplied consumption in FY19 to 2.2% in FY25.

3. MODELLING FRAMEWORK

The Australian Energy Market Operator (AEMO) commissioned Oxford Economics Australia to forecast data centre demand to FY58 to inform the 2025 Electricity Statement of Opportunities and the 2026 Integrated System Plan.

Previous approaches to forecasting data centre demand have focused on shorter time horizons of 5 to 10 years. The Lawrence Berkeley National Laboratory[6] has produced several annual reports on the outlook for United States data centre energy usage over a 5-year time horizon. The International Energy Agency[7] recently built on this work, publishing a 10-year outlook for global data centre demand. Both estimates rely on a bottom-up methodology developed by the Lawrence Berkeley National Laboratory, which utilises IT equipment shipments as a key driver of data centre energy demand. Alternatively, the California Energy Commission[8] published data centre load forecasts based on a bottom-up methodology that estimated existing and proposed data centre capacity (MW) and applied average load profiles to determine annual consumption (TWh).

The modelling framework developed by Oxford Economics Australia utilises two independently derived forecast approaches capturing different market dynamics and data sources.

- The **Project Approach,** like the California Energy Commission, represents industry expectations of future demand for data centres as revealed by project development activity.
- An **Economic Approach** was also adopted to measure the relationship between economic demand for data centre services and the energy required to meet that demand.

The final forecasts represent a weighted average of the two approaches. This method was adopted to balance the risk that market activity as measured by the current data centre project pipeline and connection requests may be overestimating future demand for data centre services. In such cases, energy consumption may be lower than expected due to delayed or reduced projects, longer ramp-up periods, or lower load realisation factors.[9] The blended approach also accounts for the possibility that historical relationships between economic growth and data centre demand may underestimate future growth, due to structural shifts in this relationship.

[6] 2024 United States Data Center Energy Usage Report (2024) *The Lawrence Berkeley National Laboratory*. Available here

[7] Energy and AI (2025) *International Energy Agency*. Available here

[8] California Energy Commission (2024) *Data Centre Load Forecasts, 2024-2040*. Available here

[9] A detailed definition of ramp-up and load realisation can be found in Section 6.1 of the technical appendix.

Fig. 7. Modelling framework

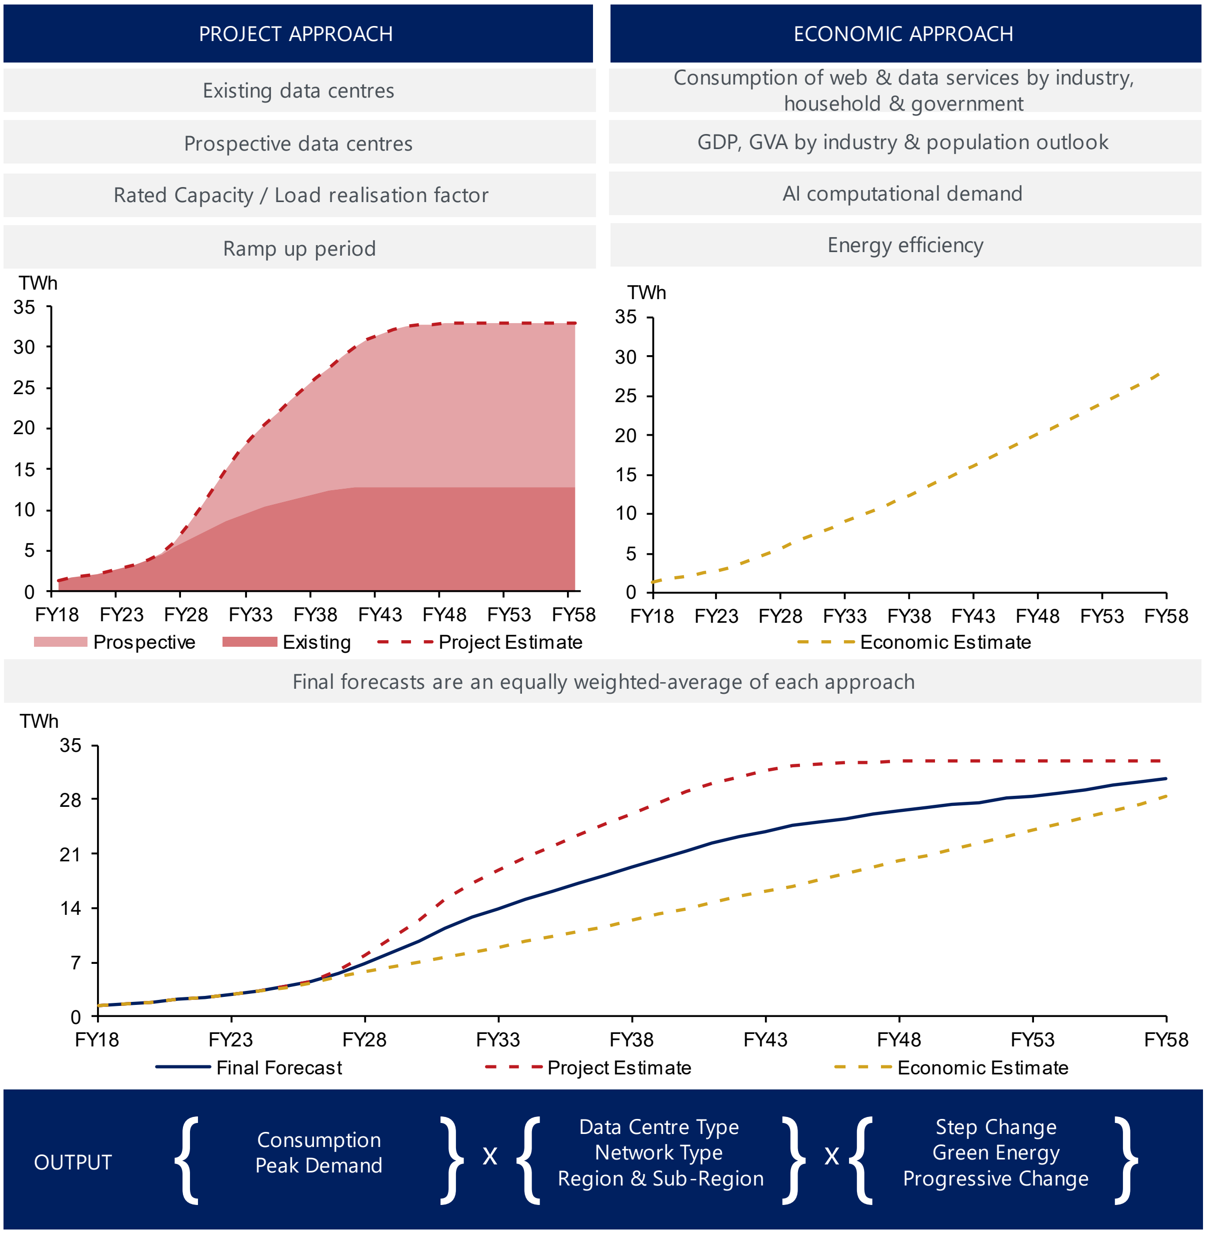

Source: Oxford Economics Australia

3.1 PROJECT APPROACH

The Project Approach leverages rich data sources from AEMO to measure energy consumption on a project-by-project basis[10]. Key data sources include:

- Existing data centre energy consumption by national metering identifier (NMI), AEMO
- Existing data centre energy consumption outlook, AEMO Large Industrial Load (LIL) surveys

[10] Data was deidentified to maintain confidentiality.

- Prospective project details including rated capacity, ramp-up period and load factor, AEMO prospective project list compiled based on NSP Standing Information Requests
- Load realisation factor for existing and prospective projects by data centre type, International Energy Agency[11]

Not all connection requests provided by NSPs in Standing Information Requests are included in the forecasts as not all connection requests are expected to be developed. Some key industry dynamics include:

- Natural project attrition as projects fail to pass final investment decision.
- Multiple developers submitting connection requests in a bidding process for a single project.
- Developers with prime land for a data centre seeking connection for future development without a current purchaser or anchor tenant.

See the Section 6 for a detailed explanation of which projects are included in the forecasts.

Data centres can take considerable time to reach a mature load. This is because developments are often built in stages, data centres take time to fully lease out their space and tenants often lease more space than their current needs, not fully utilising their contracted capacity. Both existing and prospective data centres are assumed to have a ramp-up period during which they operate below the level of energy demand. They are assumed to trend towards their mature load realisation factor over time.

The location, type of data centre and network connection type is based on the characteristics of each existing and prospective data centre. These breakdowns represent current market preferences for locations and network connection as well as expected demand for different types of data centres.

3.2 ECONOMIC APPROACH

The economic approach aims to forecast the energy required to meet Australia's economic demand for data centres over the next 30 years by estimating the economy's adoption of technology and consumption of data centre services by industry, government, and household sectors. Two key conceptual components drive the economic approach:

1. Data centre services required by a modern economy to provide goods and services.
2. The first wave of AI adoption and implementation.

Key data sources include:

- Existing data centre energy consumption by NMI, AEMO
- Prospective project details including rated capacity, ramp-up period and load factor, AEMO prospective project list compiled based on NSP Standing Information Requests
- AI intensities of data centres, Oxford Economics Australia, Data Centre Ownership and Location, NEXTDC, Microsoft and other AI focussed firms, Type and location of AI

[11] The mature load realisation factor is equal to the ratio of energy consumption to the multiplication of IEA's load factor and the data centre's rated capacity. A detailed explanation can be found in Section 6.1 of the technical appendix.

focussed data centres, Microsoft & Tech Council of Australia, Adoption rates of Generative AI. [12] [13] [14] [15]
- Technology adoption, Department of Industry Science and Resources, Adoption rates of technologies by industry[16]
- Technology processing requirements for inferencing and training, Luccioni & Jernite; EPOCH AI, IEA, Average technology compute processing requirements by industry and household sectors, segmented by intensity[17] [18] [19]
- Historic PUE trends, DCCEEW and NABERS[20] [21] [22]
- Energy Efficiency, Evolution of computing energy efficiency: Koomey's law revisited.[23] Moore's Law might be slowing down, but not energy efficiency[24]
- Server Utilisation by type of data centre, 2024 United States Data Center Energy Usage Report[25]

The first wave of AI adoption was modelled separately to account for the potential for data centre electricity demand to surge in the near-term driven by the rapid adoption of AI which would not be captured by historical relationships between industry consumption of web & data services and energy intensity.

Each industry, as well as the household and government sectors, are at different stages of digital maturity and as a result, we model their adoption of AI technology and consumption of web & data services individually. All else being equal, a larger economy will require more data centre services, resulting in data centres consuming more energy. The approach also accounts for differences in the composition of the economy by scenario.

The consumption of data centre services is then translated into energy, so that an additional unit of consumption on data centre services produces additional demand for energy.

The location of data centre demand is determined by which projects are most likely to meet incremental demand based on the maturity of prospective projects and market preference as revealed by AEMO's prospective project list.

[12] Oxford Economics Australian Building Masterplan (January 2025). Available here
[13] NEXTDC, 'S6 Sydney Data Centre', Available here
[14] Microsoft (2023) 'Microsoft announces A$5 billion investment in computing capacity', Available here
[15] Microsoft & Tech Council of Australia (2023) Australia's Generative AI Opportunity, Available here
[16] Department of Industry, Science and Resources (2024) *Exploring AI adoption in Australian businesses*. Available here
[17] Luccioni and Jernite (2024) *Power Hungry Processing*. Available here
[18] EPOCH AI (2025) *How much energy does ChatGPT use*? Available here
[19] International Energy Agency 2024, Energy and AI. Available here
[20] DCCEW (2021) *International Review of Energy Efficiency in Data Centres*. Available here
[21] NABERS (2025) *Data Centres*. Available here
[22] NABERS (2025) *Find a current rating*. Available here
[23] Prieto et al. (2024) *Evolution of computing energy efficiency: Koomey's law revisited*. Available here
[24] Koomey, J., Naffziger, S., (2015) *Moore's Law might be slowing down, but not energy efficiency*.
[25] Berkeley Lab (2024) *2024 United States Data Center Energy Usage Report*. Available here

4. FORECAST RESULTS

This chapter outlines Oxford Economics Australia's independent forecasts of electricity consumption by data centres in Australia under the three 2025 Inputs, Assumptions and Scenarios Report (IASR) scenarios: *Step Change*, *Progressive Change* and *Green Energy Industries*. The final forecasts are an equally weighted blend of:

- **Project Approach** – estimates the energy consumption of each existing and prospective data centre development; and
- **Economic Approach** – measures the economy's demand for data centres and the commensurate energy required to meet demand.

For more detailed information on these approaches see Section 3 and Section 6.

4.1 HEADLINE RESULTS

In FY25, data centres are estimated to have consumed 3.9 TWh of energy across Australia.[26] For the NEM, where the vast majority of data centres are located, this currently represents 2% of grid-supplied electricity.

Under the *Step Change* scenario, data centre energy consumption is forecast to grow 25.1% p.a. to reach 12.0 TWh by FY30, with NEM connected data centres representing 6% of the NEM's grid-supplied electricity. Over the subsequent two decades, growth is expected to slow to 5.4% p.a. to reach 34.5 TWh, with NEM connected data centres representing 12% of the grid by FY50. The slowing growth reflects the projected adoption curve of data services and the integration of AI demand in the business sector, with a greater portion of longer-term growth driven by economic fundamentals including economic growth in key data consuming sectors, population growth and business investment in digital uplift.

[26] Unless otherwise stated, data centre energy consumption figures include data centres connected to the NEM, SWIS and NTIS.

Fig. 8. Data centre energy consumption in Australia (LHS) and NEM data centre energy consumption as a share of grid-supplied electricity (RHS), *Step Change*

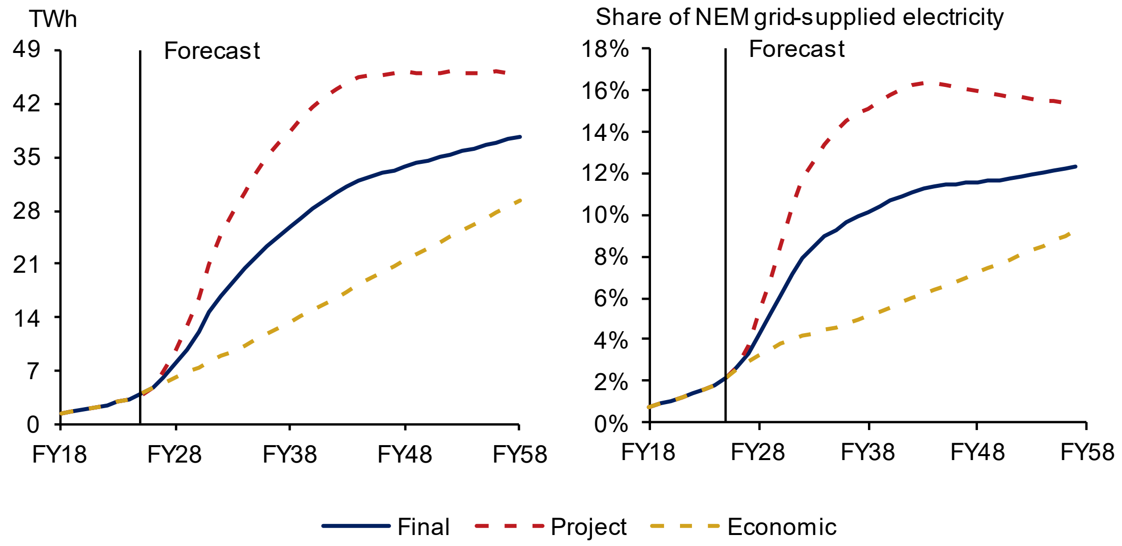

Source: Oxford Economics Australia based on AEMO data.
Note: The 'share of NEM grid-supplied electricity' is equal to OEA's estimates of 'NEM data centre energy consumption' as a share of 'NEM operational sent-out consumption' from the 2025 Electricity Statement of Opportunities.

Energy consumption in the short-term is driven by strong growth in the energy demand of existing data centres and a significant number of prospective data centre projects coming online (as estimated by the Project Approach).

Fig. 9. Existing & prospective project detail, *Step Change*

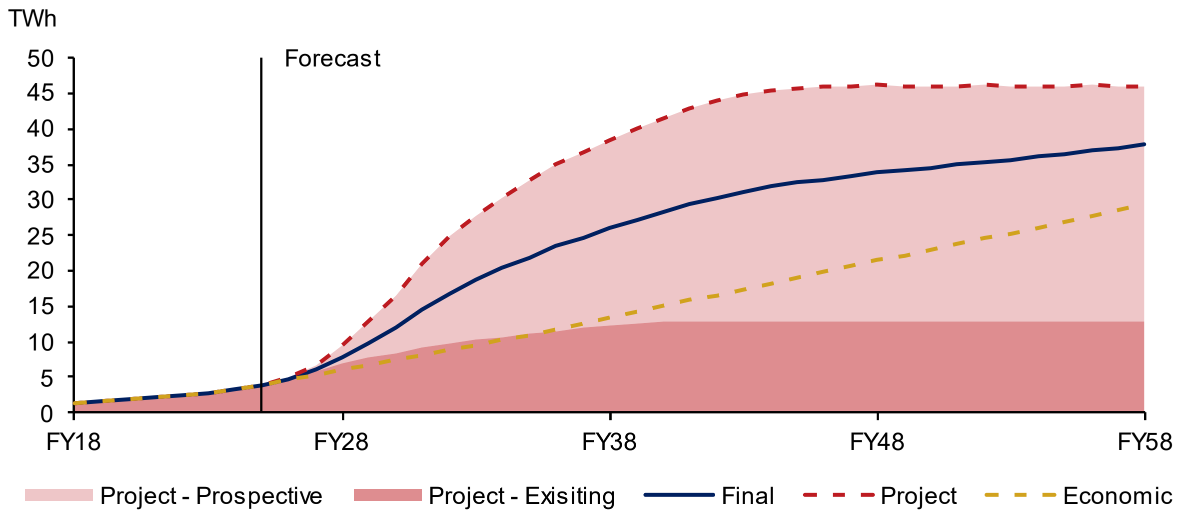

Source: Oxford Economics Australia based on AEMO data.

The Project Approach estimates consistently overperform relative to the Economic Approach estimates, though both approaches converge in the long run. Existing data centre energy consumption closely follows our Economic Approach estimates over the next 10 years. This is not by design, the two forecasts are completely independent. This finding suggests that:

- Development over the last 5 years is sufficient to satisfy the economic demand and technological progress captured by our Economic Approach in the short term.
- The significant capacity of prospective projects, of which we only include a subset, is much more uncertain and is likely to come online progressively to meet incremental additional demand.
- There is significant potential capacity in various stages of development to respond quickly to higher demand due to faster-than-expected AI adoption or demand from other technologies.

These dynamics were confirmed in our consultations with data centre providers and the expert panel, who noted that while planning for a large potential upside in AI demand, it was too early and too uncertain to commit investment at this point in time.

Australia's existing stock of data centres is expected to consume 8.5 TWh of energy by FY30, up from 3.9 TWh today. They are expected to grow at 7.8% on average until FY41, when all existing data centres have reached their mature load.

Prospective projects included in the *Step Change* scenario total 8 GW of rated capacity and are expected to consume 8.1 TWh of energy by FY30 and 33.3 TWh by FY50. Prospective projects are sourced from NSPs through AEMO's 2025 Standing Information Request and validated via interviews with data centre providers, discussion with NSPs, cross-referencing with Oxford Economics Australia Non-Residential Masterplan and desktop research. Not all potential projects or connection requests are included in the forecasts.[27] For example, in the case where multiple data centre providers or developers made several connection requests for what appeared to be an identical development as part of a proposal, project scoping or bidding process, only one of these was included.[28]

[27] See the Technical Appendix for a detailed explanation of which prospective projects are included from AEMO's Standing Information Request surveys in the *Step Change* scenario.

[28] For example, if a large tenant requested proposals to provide a 200 MW capacity site in a specific region by a specific date and multiple developers all submitted applications as part of their proposal process, one of these prospective projects was maintained in the list while the others were excluded.

Fig. 10. Data energy consumption, sub-region, *Step Change*

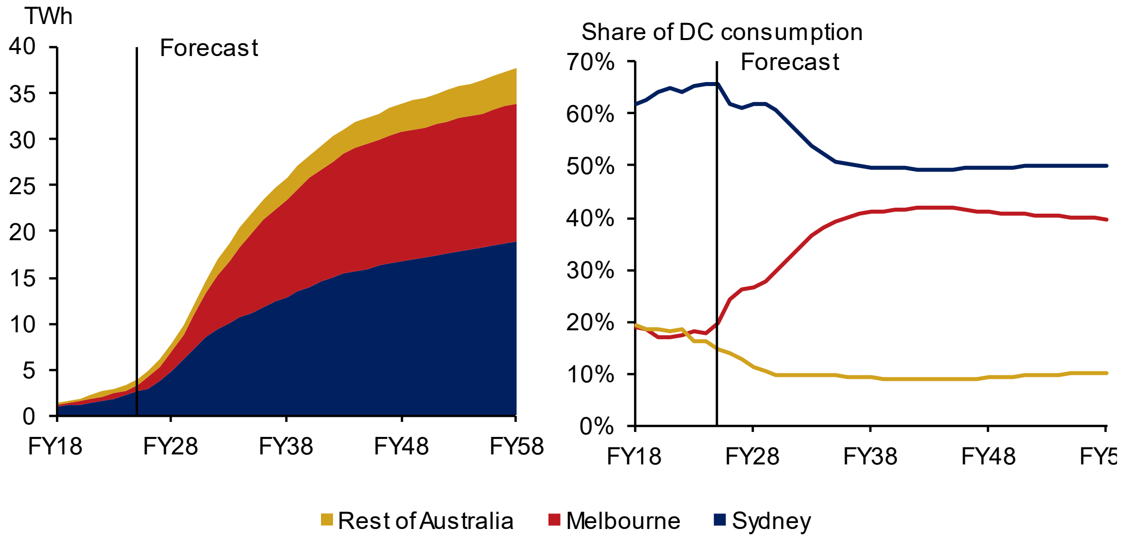

Source: Oxford Economics Australia based on AEMO data.

Sydney is expected to maintain its position as the data centre hub of Australia. Existing data centres in Sydney are still ramping towards their mature capacity, increasing consumption until the early 2040s.[29] Sydney also has a pipeline of committed and anticipated prospective projects that is larger and more mature than other sub-regions, which will absorb incremental data centre demand prior to earlier-stage developments in other sub-regions.

Consultation with market participants indicated a strong industry preference for Sydney over other locations. Sydney was preferred over other regions due to:

- A desire to be close to their client base. For hyperscalers, this means the end-users of cloud technologies. For co-locators, this means tenants as well as hyperscalers who are often anchor tenants.
- Agglomeration benefits of locating many data centres nearby, including the ability to shift workloads between data centres and workers being able to move between facilities.
- Proximity to skilled labour to build and maintain centres, in particular for co-locators whose tenants' ICT teams maintain their servers.

Despite a clear preference for Sydney, as the region's relatively more mature pipeline of existing and committed projects is absorbed, additional demand will be increasingly met by Melbourne's large stock of prospective projects. Melbourne is expected to increase its share of energy consumption from 20% currently to 30% by FY30 and stabilising around 41% in the long run.

Activity outside of Sydney and Melbourne is fairly minimal and concentrated in Central and Northern South Australia (1.2 TWh by FY50), Southern New South Wales and Canberra (0.8 TWh), and Western

[29] Mature capacity refers to the maximum level of consumption a data centre reaches in the long-run once fully ramped up. This can be though of as the stable, long-run consumption level achieved once the data centre is mature.

Australia (0.6 TWh). A more diverse regional allocation of data centres is explored as a model sensitivity in chapter five of this report.

The outlook for network connections is driven by the future development that occurs in Sydney relative to Melbourne. Grid connection requests have suggested a continued preference for the distribution networks in Sydney. In contrast, there is a significant pipeline of large developments in Melbourne expected to connect to the TX network representing a greater proportion of new demand.

Trends in network connection types historically and into the future are predominately driven by two key factors, availability and ease of connection and proximity to customers. Currently, less than 2% of data centre consumption is connected to the transmission network. This is expected to rapidly grow to 32% of total consumption by FY30 (3.8 TWh) as data centre development in Melbourne escalates. Beyond FY30, transmission's share of total consumption stabilises around 42% in the long run, reflective of the connection requests made to NSPs for prospective projects.

Fig. 11. Data centre energy consumption by network connection, *Step Change*

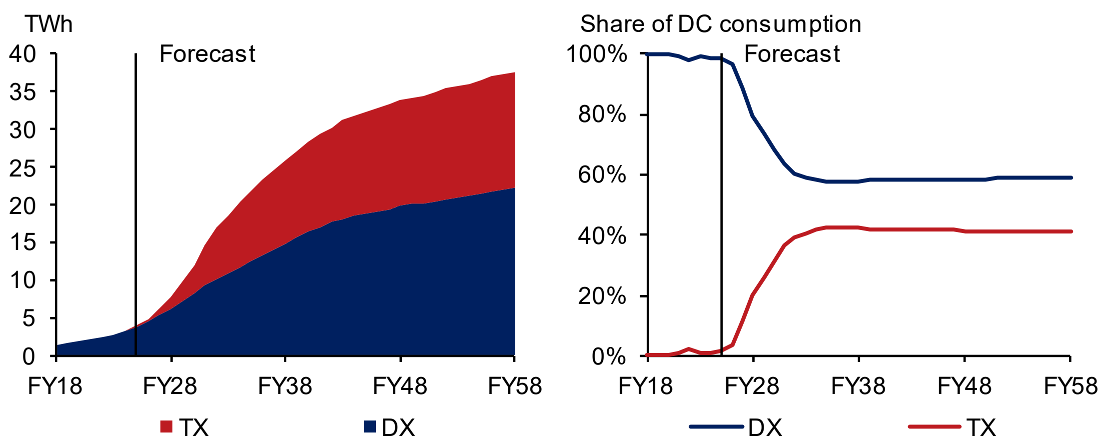

Source: Oxford Economics Australia based on AEMO data

The outlook for data centre types is driven by the market activity of different project developers. At present, market activity is heavily driven by demand from hyperscalers.

Hyperscalers accounted for just 14% of consumption in FY18 which has escalated rapidly to 55% in FY25. As a significant number of co-locator projects come online over the next few years, hyperscalers will start to lose some market share to co-locators such that by FY30 the co-locator share of consumption is expected to be 49% of total consumption.

Other data centres are set to see their share of consumption continue to decline from 13% today to just 2% by FY50. At present, there is no evidence suggesting rapid growth in edge data centres energy consumption. Although some nascent technological applications for edge centres have been proposed and tested in the financial sector, for localised gaming and VR/AR services, potential smart infrastructure support and local network AI, these applications are not currently leading to significant investment in the prospective project pipeline. Many telecommunication and enterprise data centres have been active for many years, have reached their mature load factors and therefore are not projected to significantly contribute to future consumption growth.

When using the data centre type results, it is worth considering that:

- Hyperscaler and co-locator consumption is somewhat substitutable as hyperscalers are often significant tenants of co-locators.
- The data centre type splits are based on the primary business model of each project's developer; however, we may reasonably expect business models to change over the forecast period as the industry matures.

Fig. 12. Data centre energy consumption by data centre type, *Step Change*

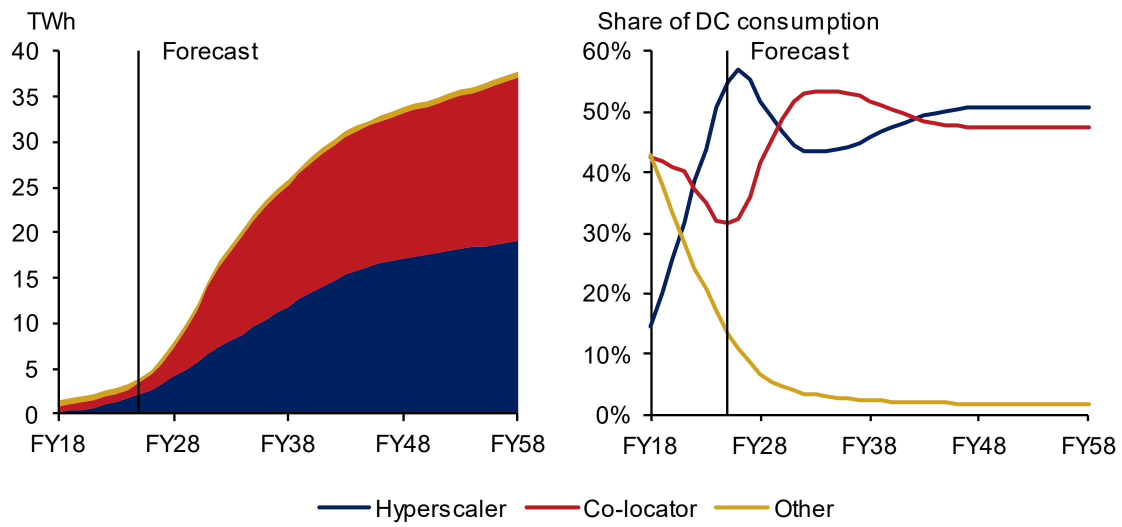

Source: Oxford Economics Australia based on AEMO data.

4.2 SYDNEY DEEP DIVE

Sydney[30] is Australia's most mature data centre market, and the outlook is characterised by a continuation of existing market trends.

Data centres in Sydney currently consume 2.6 TWh of electricity per annum, 4% of NSW's grid-supplied electricity. By FY30, this is expected to nearly triple to 11% and reach 18% by FY50, when data centres in Sydney are expected to consume 17.2 TWh of energy. Due to the relatively mature pipeline in Sydney, accounting for only existing and committed projects, data centres would consume 11.4 TWh in FY50.

[30] Sydney refers to the AEMO Integrated System Plan (ISP) sub-region Sydney, Newcastle and Wollongong (SNW).

Fig. 13. Sydney data centre energy consumption, *Step Change*

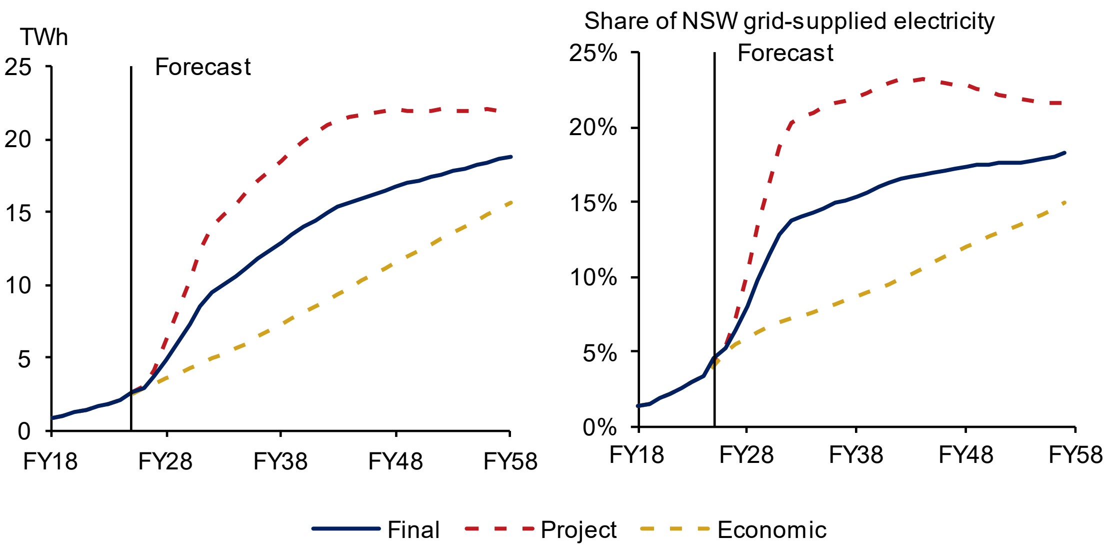

Source: Oxford Economics Australia based on AEMO data.
Note: The 'share of NSW grid-supplied electricity' is equal to OEA's estimates of 'SNW data centre energy consumption' as a share of 'NSW operational sent-out consumption' from the 2025 Electricity Statement of Opportunities.

Co-locators are expected to be the key driver of demand growth in Sydney over the rest of the decade. There is a strong pipeline of prospective (committed and anticipated) co-locators which ramp load quicker than hyperscalers, contributing to rapid growth in the short-term. From FY30 to FY45, as hyperscalers continue to ramp and prospective co-locators are nearing mature load realisation, growth for both hyperscalers and co-locators will moderate.

Fig. 14. Project Approach - Existing, anticipated and committed projects, Sydney, *Step Change*

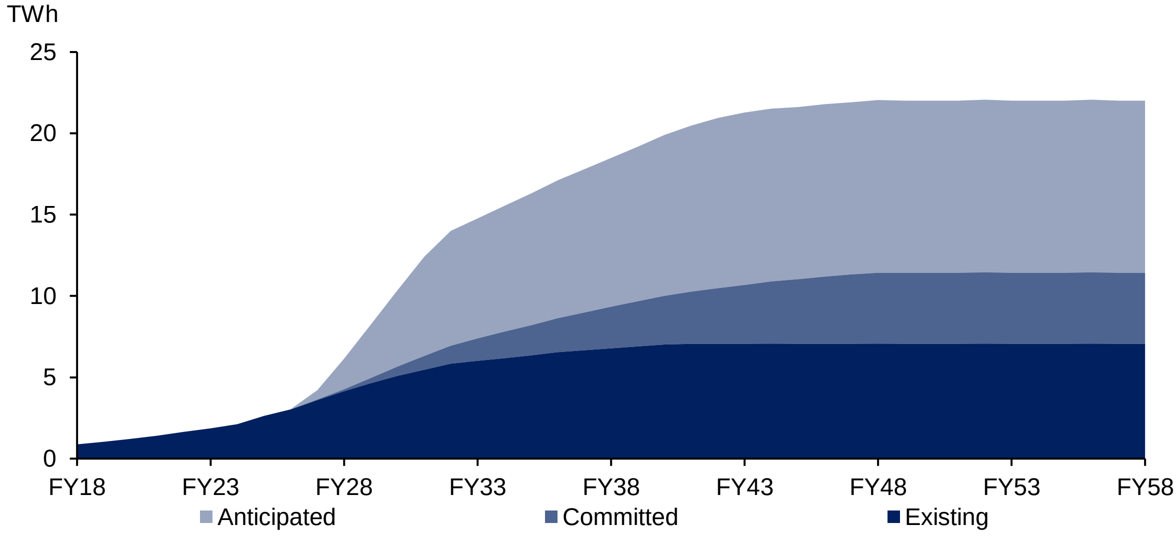

Source: Oxford Economics Australia based on AEMO data.

Grid connection requests have suggested a continued preference for the distribution networks in Sydney. By FY45, only 22% of forecast consumption in Sydney is expected to be connected to the transmission network, with around 3.9 TWh in total consumption through the transmission network expected to come online from prospective projects.

Fig. 15. Sydney data centre energy consumption splits, *Step Change*

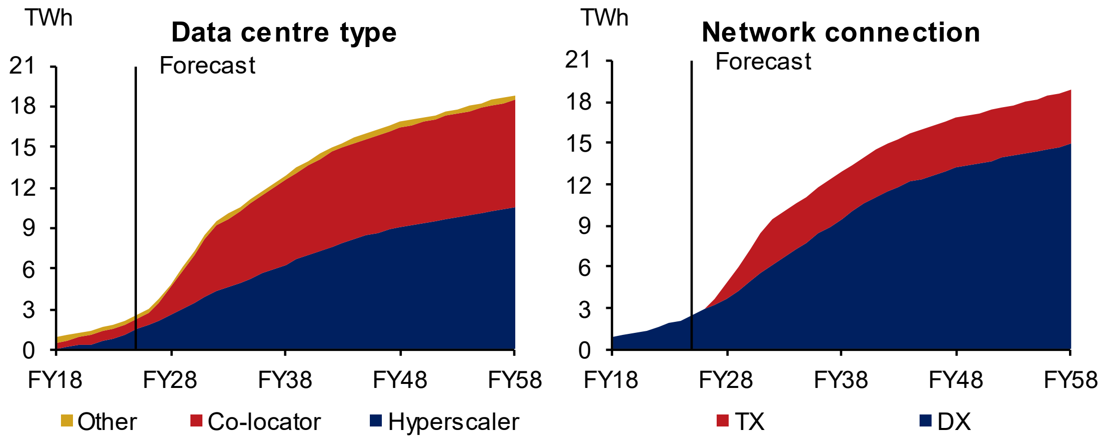

Source: Oxford Economics Australia based on AEMO data.

4.3 MELBOURNE DEEP DIVE

Growth in Melbourne[31] is expected to be driven by a rapid entry of large hyperscalers connecting to the transmission network which was little-used by data centres.

Current data centre energy consumption in Melbourne is 0.8 TWh, representing 2% of Victoria's grid-supplied energy. Under the *Step Change* scenario, Melbourne grows to 8% by FY30 and 19% by FY50 (14.1 TWh).

A significantly higher proportion of growth in Melbourne will be driven by prospective projects rather than existing developments ramping up. This drives a greater divergence between the Project Approach and Economic Approach estimates than Sydney, demonstrating that there is significantly more uncertainty over Melbourne's outlook.

Melbourne is likely to be an outsized beneficiary under any upside scenario for national data centre demand. Industry consultations indicated that Melbourne is seen as a more cost-effective and quicker location for construction. However, given high and persistent customer demand for Sydney as the industry incumbent and a key population hub, Melbourne is viewed as a secondary hub and unlikely to reach the scale of Sydney's data centre network unless there is sufficient demand to absorb the significant prospective pipeline.

[31] Melbourne refers to the AEMO sub-region Greater Melbourne and Geelong (MEL).

Fig. 16. Melbourne data centre energy consumption, *Step Change*

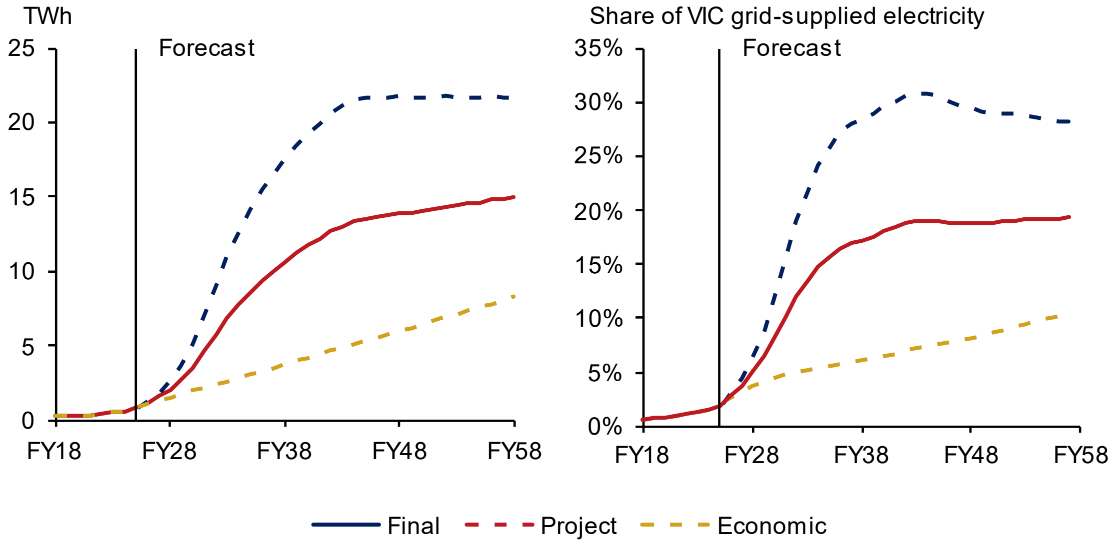

Source: Oxford Economics Australia based on AEMO data.
Note: The 'share of VIC grid-supplied electricity' is equal to OEA's estimates of 'MEL data centre energy consumption' as a share of 'VIC operational sent-out consumption' from the 2025 Electricity Statement of Opportunities.

Prospective projects are significantly larger than the existing Melbourne market and are heavily dominated by large hyperscale projects, which are expected to be connected to the transmission network. The strong pipeline means that prospective data centres could represent 73% of consumption by FY50, with consumption evenly split between hyperscalers and co-locators.

Fig. 17. Project Approach - Existing, committed & anticipated projects, Melbourne, *Step Change*

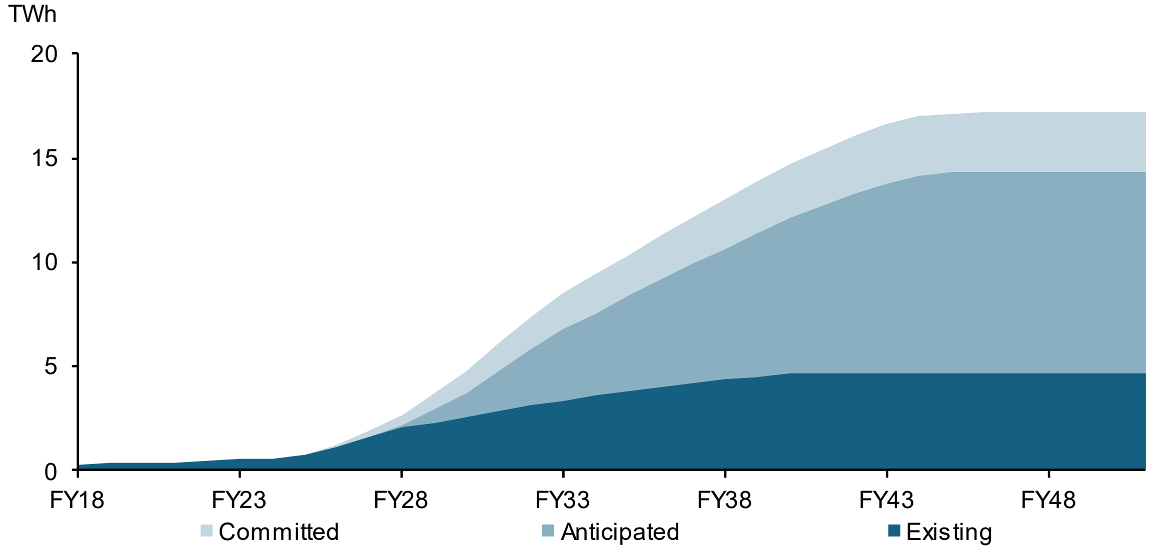

Source: Oxford Economics Australia based on AEMO data

Melbourne's prospective pipeline is much more likely to request connection to the transmission network – 83% of expected consumption in FY50. It is this pipeline that drives the share of transmission connected data centres at the national level.

Fig. 18. Melbourne data centre energy consumption splits, *Step Change*

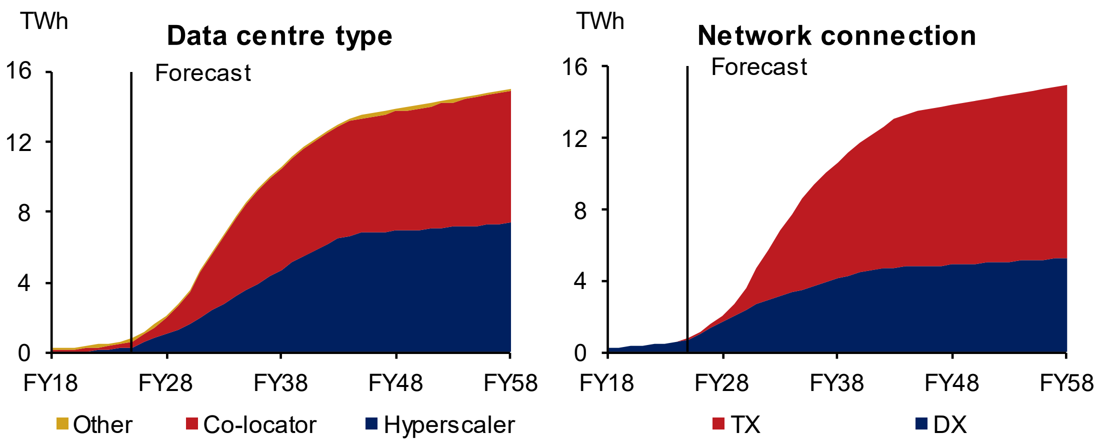

Source: Oxford Economics Australia based on AEMO data.

4.4 AI DEEP DIVE

AI demand could significantly increase energy consumption in data centres. The first wave of AI adoption is expected to be largely offset by ongoing improvements in software, hardware, and computing efficiency. In the long run however improvements in energy efficiency will induce greater demand as costs lower.

The forecasts indicate that without energy efficiency gains, demand from the first wave of AI through currently known AI applications will increase to 9 TWh by the end of the forecast period. However, energy efficiency gains offset most of this with adjusted AI energy demand peaking at around 0.9 TWh between 2027-2028, contributing only marginally to economic demand in the *Step Change* consumption forecasts, with the primary driver remaining demand for data centre services from industry uses in the economy. In practice, energy efficiency measures reduce the cost of computations, inducing additional demand. This relationship has been shown to be positive over time, that is, demand outpaces energy efficiency. This relationship is accounted for in the Economic Approach estimates. Please see the Technical Appendix for more information on the Economic Approach methodology.

The majority of AI energy demand is concentrated in hyperscaler data centres because they host the most compute-intensive AI workloads. Training large-scale models requires significant processing power, memory bandwidth, and infrastructure that only hyperscalers can provide at scale. Moreover, most enterprise and consumer access to AI services occurs via APIs hosted in the cloud, meaning that even inference activity is often centralised. Hyperscalers benefit from economies of scale, custom AI-optimised chips, and global reach, positioning them as the backbone of AI development and deployment worldwide[32].

This chart highlights a critical dynamic in the future of AI infrastructure. While AI adoption could significantly increase electricity demand, particularly from hyperscalers, much of this impact may be

[32] McKinsey 2024, 'AI power: Expanding data center capacity to meet growing demand'

mitigated through advances in hardware efficiency, model compression, and smarter workload scheduling. Key results across the data centre types include:

- Hyperscaler data centres (e.g. those operated by AWS, Google, and Microsoft) show the most dramatic growth, with gross AI-related energy demand rising steeply to over 7 TWh by 2058. However, once energy efficiency improvements are applied, this demand remains nearly flat, peaking just after 2030 before declining sharply due to expected gains in processing efficiency and model optimisation.
- Co-locators also see substantial growth, with unadjusted demand exceeding 1.5 TWh by 2058. The efficiency-adjusted trajectory is far more moderate, plateauing below 0.5 TWh.
- Edge/Telecom data centres experience slower and steadier growth, with both adjusted and unadjusted demand remaining relatively low (<0.5 TWh), reflecting more modest AI workloads and greater distribution across smaller devices.

Fig. 19. AI demand outlook for data centres, *Step Change*

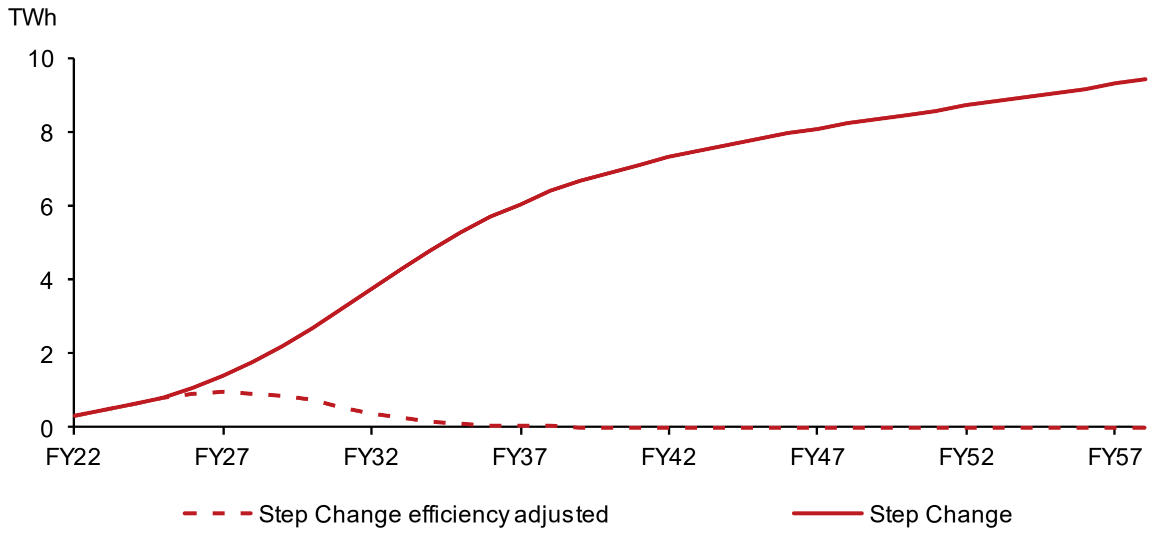

Source: Oxford Economics Australia based on AEMO data

4.5 SCENARIO RESULTS

Differences across AEMO's scenarios are driven by scenario assumptions underlying both forecast approaches. Under the Project Approach AEMO's scenario allocation of anticipated and prospective projects varies under each scenario. These assumptions are outlined in Section 6.1. Under the Economic Approach, economic growth and industrial composition, as well as rates of technological adoption are flexed as outlined in Section 6.2.

Under *Green Energy Industries*, by FY30, energy consumption is projected to be 13.1 TWh, 9.1% greater than under the *Step Change* scenario. Under *Progressive Change*, the forecast by FY30 is projected to be 8.5 TWh, 29.7% smaller than under the *Step Change* scenario.

Grid-supplied energy also changes under the scenarios, NEM connected data centres share of grid-supplied electricity under *Progressive Change* is 5% compared to 6% in *Step Change* and *Green Energy*

Industries by FY30, though *Progressive Change* only increases to 9% by FY50 relative to 12% under *Step Change*. Under *Green Energy Industries* this share tracks closely to *Step Change* through to FY33 where the additional grid supply under the *Green Energy Industries* scenario results in data centre demand flattening as a share of the grid, remaining near 10%.

Fig. 20. Data centre energy consumption in Australia (LHS) and NEM data centre energy consumption as a share of grid-supplied electricity (RHS) by scenario, Australia

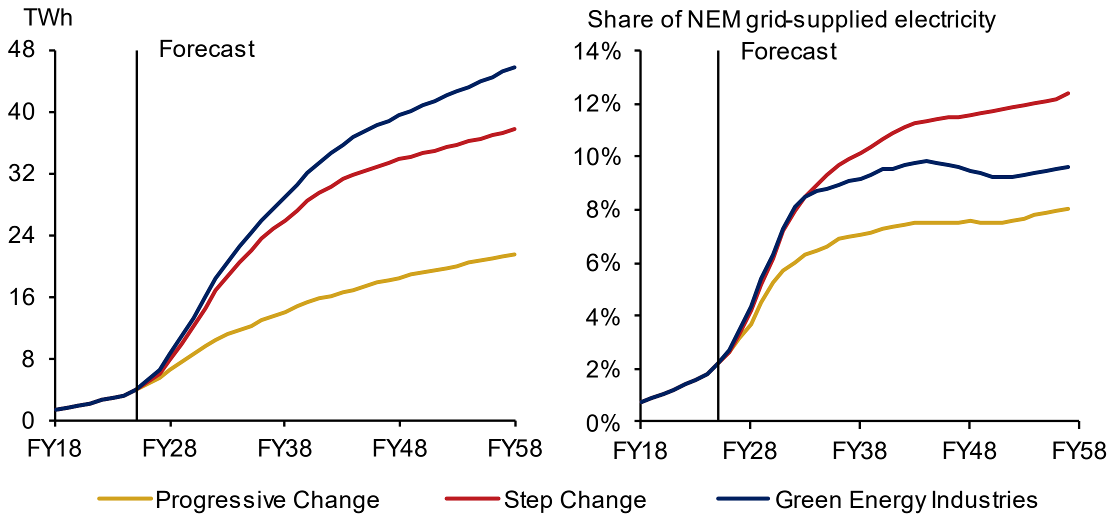

Source: Oxford Economics Australia based on AEMO data.
Note: The 'share of NEM grid-supplied electricity' is equal to OEA's estimates of 'NEM data centre energy consumption' as a share of 'NEM operational sent-out consumption' from the 2025 Electricity Statement of Opportunities.

There is a greater divergence in scenario forecasts for Melbourne relative to Sydney, reflecting the greater proportion of prospective projects that can be deployed to meet greater demand under *Green Energy Industries* or delayed/abandoned in a lower demand environment. Sydney has relatively more committed and existing projects, which are not as responsive to different demand outlooks.

Fig. 21. Major sub-region data centre energy consumption by scenario

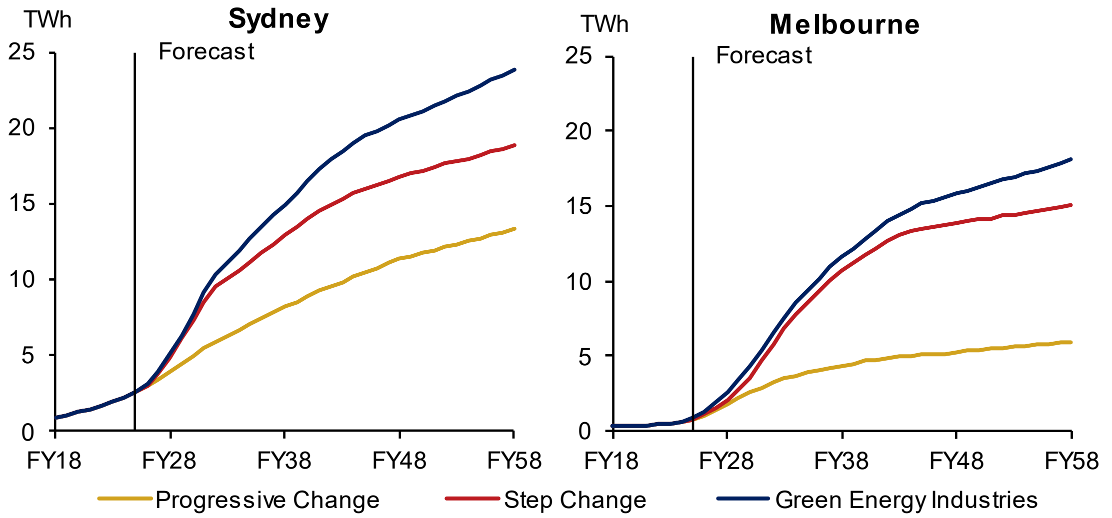

Source: Oxford Economics Australia based on AEMO data.

5. SENSITIVITY ANALYSIS

This section presents the results of four sensitivities designed to test the robustness of the central forecast and explore how data centre electricity demand could evolve under alternative assumptions. These sensitivities examine how outcomes vary when regional development patterns shift, AI adoption accelerates, Australia captures a greater share of Asia Pacific data centre demand, or a larger share of the project pipeline proceeds. Together, they help quantify the range of plausible demand trajectories and highlight key sources of uncertainty in the outlook. The chapter also includes a qualitative assessment of key risks and uncertainties affecting the forecasts.

5.1 REGIONAL DIVERSIFICATION SENSITIVITY

This sensitivity tests the question, 'How would data centre consumption be regionally allocated if it reflected professional workforce, proximity to international fibre-optic cables and relative population size rather than revealed market preference?'

Under this sensitivity we maintain aggregate final demand however, we allocate demand to each state based on their relative population size and outlook. In states with multiple sub-regions we weight sub-regions that the market has prioritised historically due to proximity to international fibre optic cables and large professional workforces that both demand data centres services and service data centres.

The outcome of this sensitivity is a much lower outlook for Sydney and Melbourne and therefore NSW and Victoria – states which currently dominate the outlook. The reasons for the current market preference for Sydney and Melbourne are discussed in chapter four.

NSW's demand adjusts downwards in the first year and remains structurally lower throughout the forecast. In Victoria, the results show that strong growth under the central *Step Change* forecast represents a rebalancing of demand commensurate with Victoria's population, however after the first few years the central forecast and sensitivity diverge as Melbourne's data centre consumption far outpaces population growth.

Fig. 22. Data centre consumption, regional diversification sensitivity, *Step Change*, NSW & VIC

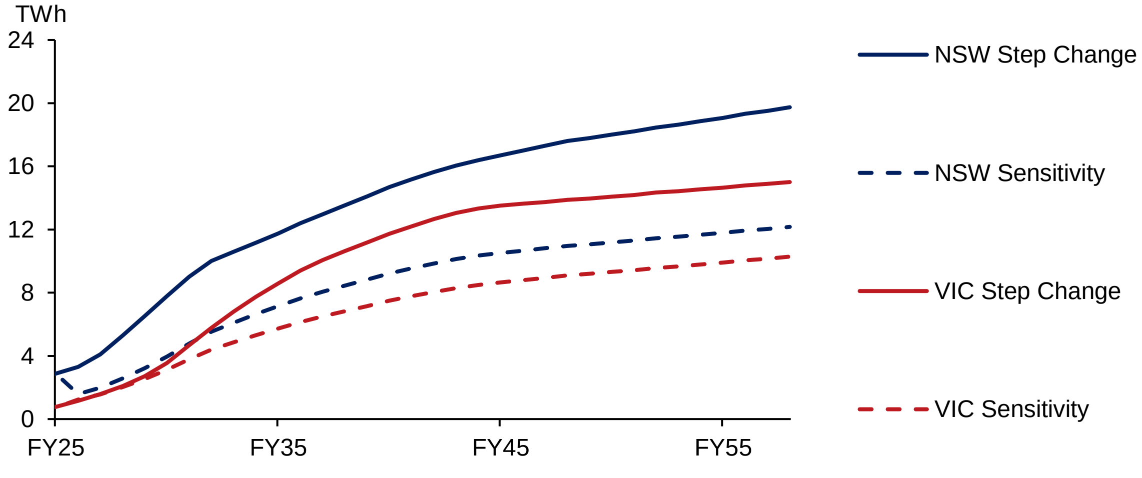

Source: Oxford Economics Australia based on AEMO data.

Queensland, Western Australia and South Australia have much higher consumption under the sensitivity, demonstrating that the current pipeline of data centre developments in these regions is much smaller than their population would suggest. Queensland has the greatest growth under the sensitivity, reaching equivalent consumption in the long run as is expected in NSW in the next 5 years.

See chapter four for discussion on the reasons for the current market preference for Sydney and Melbourne.

Fig. 23. Data centre consumption, regional diversification sensitivity, *Step Change*, QLD, WA & SA

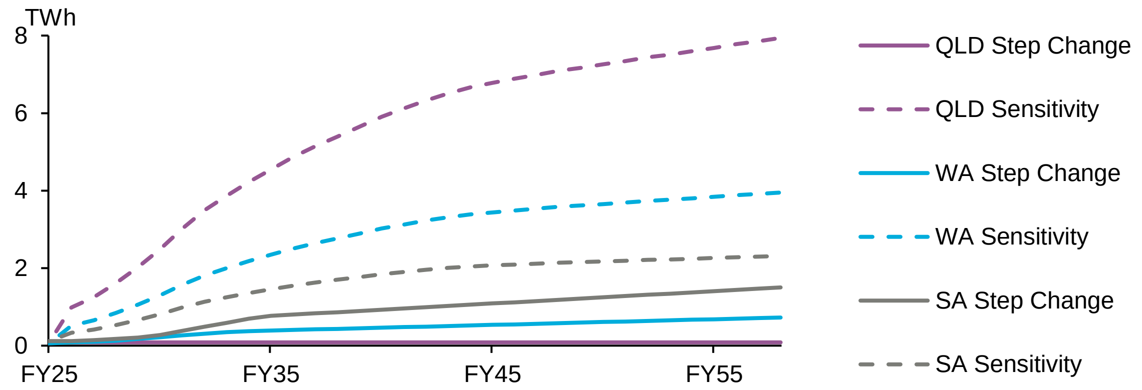

Source: Oxford Economics Australia based on AEMO data.

5.2 RAPID AI UPTAKE SENSITIVITY

This sensitivity tests the question, ‘How quickly could data centre consumption grow if AI demand escalates rapidly and ‘soaks up’ all available data centre capacity?’

Under this sensitivity we assume that recently-completed and committed prospective projects ramp up more rapidly to a higher load realisation factor in response to stronger near-term demand. Specifically, existing data centres completed from 2022 onwards and committed prospective projects ramp-up in 5 years to their full load realisation.

In a world where demand is sufficiently high, the existing stock of data centres could consume 8.5 TWh of energy in FY30. Together with prospective projects this could almost double to 16.2 TWh by FY30. Additional consumption due to these assumptions peaks in FY33 at 12.6 TWh before declining as ramp rates in the central forecast catch up to the sensitivity. The long run level of additional consumption reflects the impact of data centres reaching their full load rather than the central load realisation assumptions.

Fig. 24. Data centre consumption under rapid AI uptake sensitivity

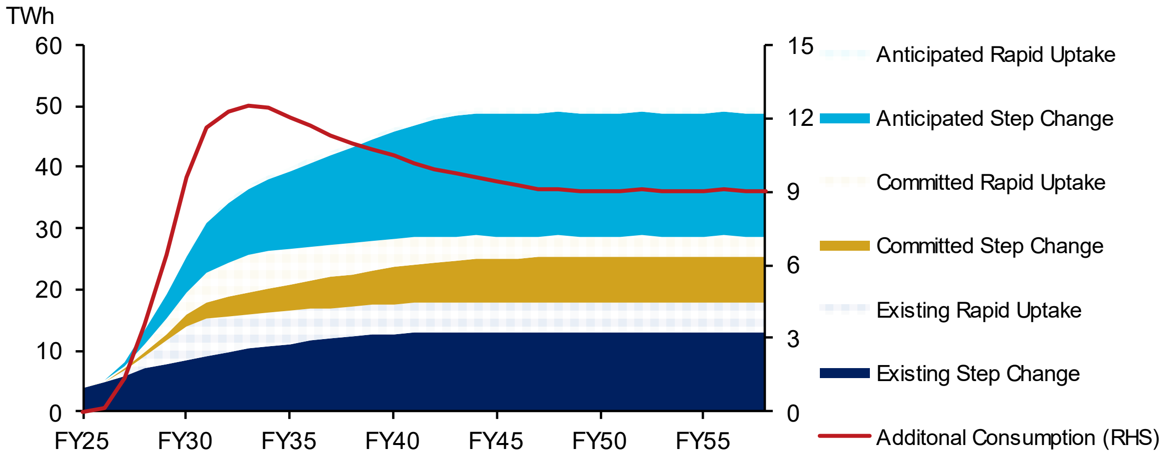

Source: Oxford Economics Australia based on AEMO data.

5.3 GLOBAL MARKET SHARE SENSITIVITY

This sensitivity tests, 'What would a stable or growing market share of Asia Pacific data centre growth imply for data centre energy demand in Australia?'. Rapidly increasing domestic demand or stronger comparative advantage of data centre services in Australia could drive increasing market share for Australian data centres. Consultation suggested that while strong demand growth to date is reflective of growing domestic demand, Australia is well placed to be relatively competitive in the Asia region in the future due to our stable regulatory environment and strong reputation for data security.

We test two assumptions under this sensitivity. The first keeps Australia's current share of Asia Pacific (excl. China) consumption constant at 6.7%. The second allows this share to grow in line with linear trend growth since FY20. We then apply this share to forecasts of Asia Pacific data centre energy consumption from the International Energy Agency's (IEA) recent estimates under two scenarios, 'Baseline' and 'Lift-Off'.[33]

Australia's market share of the Asia Pacific's data centre energy consumption has been steadily increasing from 6.1% in FY20 to 6.7% in FY24.

Until FY26 all the central scenario forecasts are within the range that Australia's constant or trended market share would imply. In FY27 *Step Change* and *Green Energy Industries* begin to grow at a rate which implies a much faster growth in market share than Australia has demonstrated historically. In general the sensitivities testing Australia's market share give results in line with *Progressive Change* and the raw economic approach estimates though it's important to note that there are many differing assumptions and methodologies underpinning the IEA, AEMO and OEA analyses.

[33] International Energy Agency 2024, Energy and AI. Available here

Fig. 25. Data centre consumption under global market share sensitivity

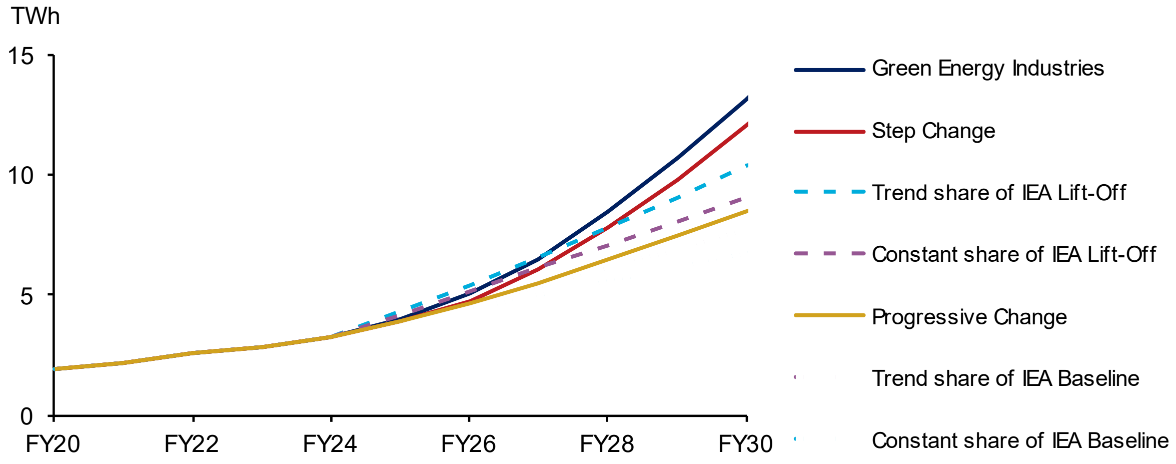

Source: Oxford Economics Australia based on AEMO data.

Fig. 26. Implied market share of Asia Pacific (excl. China) data centre consumption

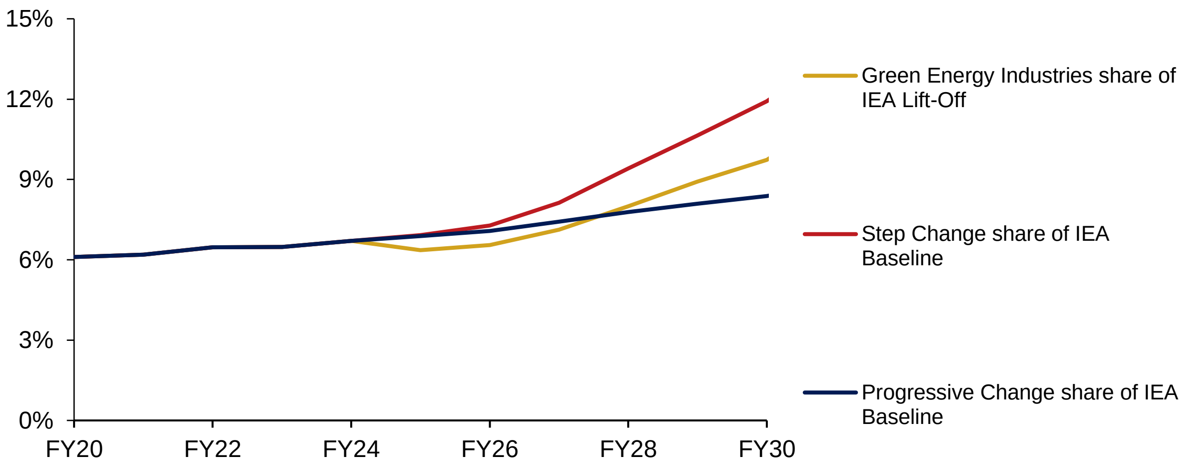

Source: Oxford Economics Australia based on AEMO data.

5.4 ADDITIONAL PIPELINE SENSITIVITY

This sensitivity tests the impact of including 8 additional data centres from AEMO's list of prospective projects. These projects were included as an upside sensitivity based on discussions with NSPs and a review of prospective projects from prospective developers. They were excluded from the central analysis either because they are not sufficiently developed or were judged to represent multiple instances of a prospective project as part of a bidding process.[34] This sensitivity tests the impact on the *Step Change* scenario if a selection of these projects were to proceed.

The impact of including the additional projects to the Project Approach is 10 TWh in the long run which equates to 5 TWh of additional consumption to the final forecast (which equally weights the Project Approach and Economic Approach).

[34] Further discussion of the treatment of prospective projects in the *Step Change, Progressive Change* and *Green Energy Industries* scenarios can be found in Section 6.1

Fig. 27. Data centre consumption under additional pipeline sensitivity, Project Approach

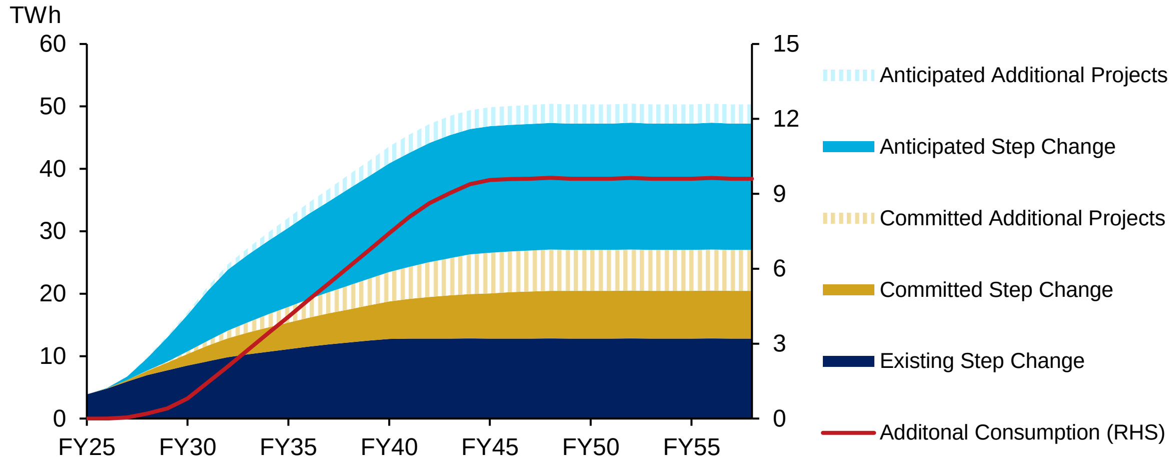

Source: Oxford Economics Australia based on AEMO data.

5.5 KEY RISKS AND UNCERTAINTIES

Forecasting electricity demand from data centres over a multi-decade horizon involves a high degree of uncertainty. This chapter sets out the key risks to the outlook, grouped according to their implications for the Project Approach, Economic Approach, and the Final Forecast, which blends the two.

The Project Approach is based on existing and prospective project-level information, including rated capacity, connection requests, and ramp-up assumptions. Risks to this approach primarily arise from factors that could delay, downsize or accelerate the physical buildout of data centre infrastructure. In particular:

- The pace of development is highly dependent on the availability of transmission and distribution infrastructure. AEMO and international experience show that grid connection delays, for example due to long lead times for substation equipment, slow the pace of consumption.
- We did not consider changes to government policy or planning approvals which could deter investment or if implemented at local or state levels could distort the modelled regional distribution.
- Global supply chains remain under pressure, with hyperscaler demand and geopolitical trade risks (e.g. chip export controls) adding to delivery uncertainty and Australia's relative competitiveness.
- The majority of capacity in the project pipeline is driven by a small number of global cloud providers. If these firms revise their investment strategies or shift workloads to other countries, demand may fall short of the projected capacity ramp up.
- Conversely, a rapid AI-driven demand spike could compress ramp up timelines and push existing projects to full load faster than modelled.

The Economic Approach estimates demand based on macroeconomic growth, digital adoption, and energy intensity trends. Risks to this approach largely reflect the uncertainty in how technology, AI, and economic activity shape demand for data centre services. In particular:

- The single largest source of demand-side uncertainty is the trajectory of AI adoption. If generative AI and other compute-intensive applications are adopted faster than anticipated, electricity demand could significantly exceed baseline forecasts. By contrast, delays in adoption due to slower consumer adoption, reduced willingness to pay, security concerns, regulatory constraints, limited business readiness, or fewer business applications than expected would reduce the expected growth.
- The model assumes continued growth in the electricity intensity of web and data services, based on recent trends in energy consumption outpacing energy efficiency gains. However, improvements in software, hardware, and AI model compression could overtake demand and reverse this trend.
- If Australian businesses increasingly rely on overseas data centres for AI workloads, domestic energy use may fall short of what economic demand implies. Similarly, international developments in sovereign data policies, cloud regulation, or hyperscaler strategy could influence the extent to which demand is serviced.

On balance, risks to the Project Approach tend towards the downside while risks to the Economic Approach are biased upwards. The methodological decision to equally weight each forecast in the final forecasts is an attempt to capture this asymmetric risk.

6. TECHNICAL APPENDIX

6.1 PROJECT APPROACH

Historical data

Historical data was sourced from AEMO metering data collected by site across the NEM. For Western Australia and Northern Territory historical metering data was incomplete and dual-use NMIs,(sites with both a data centre use and other large industrial loads), were identified. Historical consumption data was estimated for Western Australia and Northern Territory based on the OEA Building Construction Masterplan January 2025[35], back casting consumption using the same methodology applied to prospective projects. This means that in this analysis, historical consumption for the NEM regions can be considered actual while historical data for Western Australia and the Northern Territory should be considered as estimates.

Prospective projects

The prospective project list was compiled and categorised by AEMO based on NSP Standing Information Requests, interviews with data centre providers, discussion with NSPs, cross referencing with Oxford Economics Australia Non-Residential Masterplan and desktop research. In the case where multiple data centre providers or developers made several application requests for what appeared to be an identical project as part of a proposal, project scoping or bidding process, only one of these was included in the prospective project list. For example, if a large tenant requested proposals to provide a 200 MW capacity site in a specific region by a specific date and three providers or developers all submitted applications as part of their proposal process, one of these prospective projects was maintained in the list while the other two were excluded. Data collected from these sources and assigned to each project included:

- 2025 IASR Sub-region
- Rated capacity (MW)[36]
- Ramp-up period (years)[37]
- Ramp start date
- Data centre type[38]
- Project status (Constructed, Committed, Anticipated, Proposed)[39] [40]
- Grid connection status (Completion, Application, Enquiry, Prefeasibility)

Prospective projects were then allocated to scenarios based on the following:

[35] Oxford Economics (2025) *Australian Building Masterplan – January 2025.* Available here

[36] Rated capacity refers to the maximum demand level specified in a connection agreement between the data centre and a network service provider.

[37] Ramp-up period refers to the duration in years for the new data centre connection to mature load

[38] Please see Introduction for definition of data centre types

[39] AEMO (2024), *Forecasting Approach – Electricity Demand Forecasting Methodology*. Available here

[40] Project status as defined in AEMO's Electricity Demand Forecasting Methodology.

Fig. 28. Scenario allocation criteria

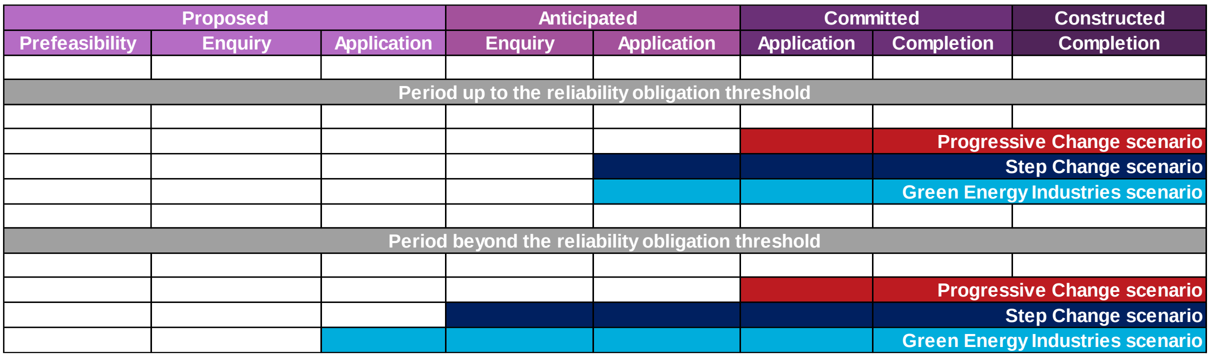

*The forecast reliability gap period in this case this date is 1 July 2029. For more information please see National Electricity Law, Part 2A (Retailer Reliability Obligation). Available here
Source: AEMO

Ramp-up period

Assumed ramp-up periods are applied on a project-by-project basis from data provided by NSPs when compiling the prospective project list. Five years was assumed to be the minimum ramp up period and in the few cases where ramp-up periods were not provided, the ramp-up period of similar prospective developments was used. Descriptive statistics for these assumptions are below:

Fig. 29. Ramp-period assumptions (years)

	Minimum	Maximum	Mean	Median
Hyperscaler	5	20	14	15
Co-locator	5	20	11	15
Edge/Telco/Enterprise/Other	5	15	10	10
Total	**5**	**20**	**12**	**15**

Source: Oxford Economics Australia based on AEMO data

Load Factor

Load factor is equal to the ratio of average consumption to peak demand for a data centre.

Mature load realisation factor

The mature load realisation factor is equal to the ratio of energy consumption to the multiplication of load factor (defined above) and rated capacity. We have based our assumptions by data centre type on the Base 2030 assumption from the International Energy Agency's (IEA) recent report, Energy and AI[41].

In their report IEA use "Load Factor" to define the ratio of energy consumption and rated capacity. As "Load factor" has traditionally been used to define the relationship between peak demand and average consumption, we have converted their assumptions of "Load Factor" (IEA definition) to the

[41] International Energy Agency (2025) *Energy and AI*. Available here

AEMO definition of mature load realisation factor based on our assumption of load factor (AEMO definition) as per the below:

$$Mature\ load\ realisation\ factor = \frac{Load\ factor\ (IEA\ definition)}{Load\ factor\ (AEMO\ definition)}$$

Fig. 30. Mature load realisation assumptions

	Load factor (IEA definition)	Load factor (AEMO definition)	Mature load realisation factor assumption
Hyperscaler	51%	0.75	**68%**
Co-locator	47%	0.85	**55%**
Edge/Telco/Enterprise/Other	45%	0.80	**56%**

Source: Oxford Economics Australia based on AEMO data

Fig. 31. Project Approach key assumption sources summary table:

Input assumption	Source
Historical data	Key source: AEMO metering data by site Supplemented by: OEA Building Construction Masterplan January 2025 for Western Australia and Northern Territory where Metering data was incomplete and dual-use NMIs including both data centre use and other large industrial loads were identified[42]
Prospective projects	Key source: AEMO prospective project list compiled from NSP 2025 Standing Information Requests, LIL surveys and Data centre developers Validated by: Data Centre developments from OEA Building Construction Masterplan January 2025.[43]
Ramp-up period	Key source: AEMO prospective project list compiled from NSP 2025 Standing Information Requests. Validated by: Data centre provider surveys, Trends in data centre vacancy rates[44], consultation with the Expert Panel, Australian benchmarks and International benchmarks.[45] [46]

[42] Oxford Economics (2025) *Australian Building Masterplan – January 2025.* Available here
[43] Oxford Economics (2025) *Australian Building Masterplan – January 2025.* Available here
[44] Cushman & Wakefield (2024) *APAC Data Centre Update*. Available here
[45] Transgrid (2025) *Meeting demand growth in the Western Sydney Aerotropolis 'Priority Growth Area'*. Available here
[46] EirGrid (2023) *Ireland Capacity Outlook 2022-2031.* Available here

Load factor	Key source: AEMO peak demand data in comparison with AEMO consumption data and desktop research Validated by: Consultation with the Expert Panel and international benchmarks.[47]
Mature load realisation factor	Key source: IEA (2025) Energy & AI report[48] Validated by: Data centre provider surveys, historical analysis of AEMO consumption data, consultation with the Expert Panel, CSIRO benchmarks. International benchmarks.[49] [50]
Allocation to data centre type	Assumption based on developer in OEA Building Construction Masterplan January 2025 and additional desktop validation for larger developers that commonly build both co-locators and hyperscalers. Validated by: Consultation with the Expert Panel and manual allocation where appropriate.

Existing data centres

The demand outlook for existing data centres was modelled separately from prospective projects to account for differing historical trends in specific market segments. Demand from existing data centres was built from the bottom up on a project-by-project basis based on:

- Trend analysis of historical consumption
- Large industry load surveys
- Mature load realisation factor assumptions
- Ramp-period assumptions

Modelling of existing data centre demand used six treatments based on the amount and type of data available from the above sources that had the following criteria:

[47] Independent Commodity Intelligence Services, *Data centres: Hungry for power*. Available here
[48] IEA (2025) *Energy & AI*, Available here
[49] CSIRO (2021) *Data Centres and the Australian Energy Sector*. Available here
[50] EirGrid (2023) *Ireland Capacity Outlook 2022-2031*. Available here

Fig. 32. Project Approach treatment of existing data centre consumption outlook

	LIL Survey response available	Survey response aligns with historical consumption trends	Historical consumption stable over extended period	Historical consumption aligned with mature market assumptions
Survey response adopted	Yes	Yes	-	-
Consumption approaches surveyed mature load realisation factor over the corresponding ramp period assumption	Yes	Mature load realisation factor consistent with mature load assumptions but ramp rate inconsistent with historical consumption	-	-
Consumption approaches the corresponding mature load realisation factor assumption over the surveyed ramp period.	Yes	Mature load realisation factor inconsistent with mature load assumptions but ramp rate aligned with historical consumption	-	-
Consumption approaches the lower end of mature load factor identified from data centre provider surveys* over the corresponding ramp period assumption	No	-	No	Historical consumption not aligned with IEA mature load realisation factor and NSP ramp period assumptions
Consumption approaches the corresponding mature load realisation factor assumption over the corresponding ramp period assumption	No	-	No	Historical consumption relatively aligned with IEA mature load realisation factor and NSP ramp period assumptions
Historic consumption maintained	No	-	Yes	-

*Surveys of data centre providers suggested mature load factors ranging from 20-60%. Existing data centres fitting the criteria of treatment 4 were assumed to reach a mature load realisation factor of 20-30%.
Source: Oxford Economics Australia based on AEMO data

These treatments produce an individual demand outlook for each existing data centre which was then aggregated by sub-region, type of data centre and connection type to produce the short-run demand outlook.

Prospective projects

Individual prospective projects were allocated to scenarios based on AEMO's criteria outlined in **Fig 26**. The short-run demand outlook for each project was then calculated based on the corresponding ramp period, load factor and mature load realisation factor assumptions. That is, for each data centre i of type *T* in year *t*, demand was calculated as:

$$
\begin{aligned}
&Consumption\ (GWh/year)_t^i \\
&\qquad = Rated\ capacity\ (MW)^i \times \frac{\#\ hours\ in\ year\ t}{1000} \times \frac{\#\ years\ since\ ramp\ start_t}{Ramp\ period_T} \\
&\qquad \times Load\ factor_T \times Mature\ load\ realisation_T
\end{aligned}
$$

The market outlook for consumption across sub-regions, data centre types and network types was then summed across those characteristics to produce the aggregated short-run demand outlook for prospective projects.

6.2 ECONOMIC APPROACH

Economic forecasts

Oxford Economics Australia utilised AEMO's FY25 to FY58 economic forecasts provided by Deloitte Access Economics (DAE) to estimate the size of the economy under the three scenarios to FY58. We assume that, despite recent volatility, the outlook for longer-term economic growth has not materially changed.

Fig. 33. Australian GDP by scenario, FY24 to FY58

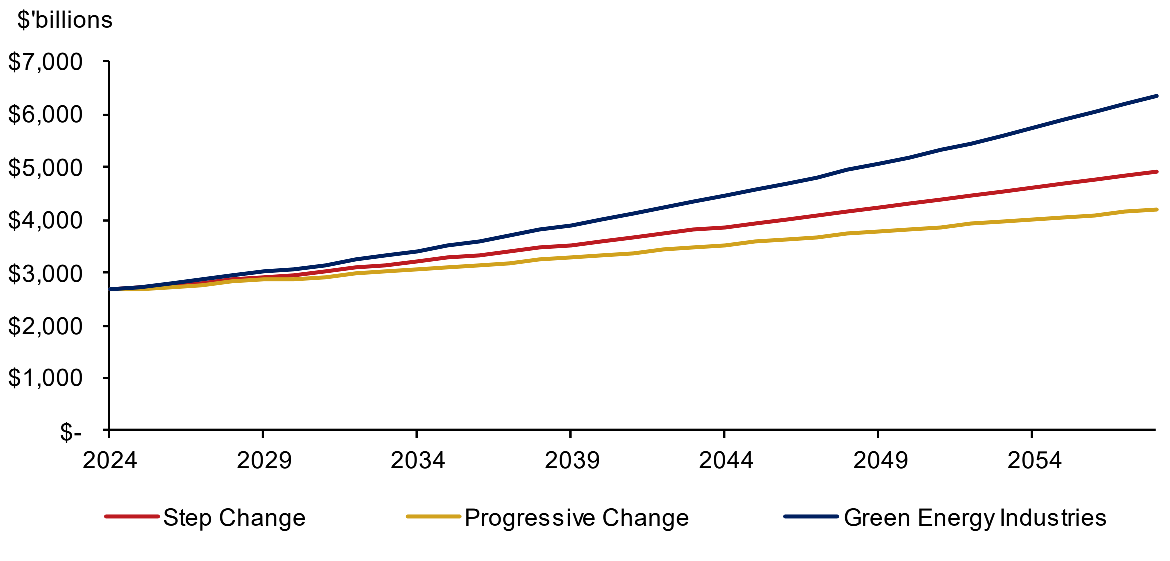

Source: AEMO

Differences in the size of the economy are forecast to impact the level of web and data services consumed, which in turn will affect electricity demand. Under a *Progressive Change* and *Green Energy Industries* Scenario, the economy will be 14.9% smaller and 29.0% larger, respectively, than under the *Step Change* Scenario. Furthermore, the composition of the economy will differ under each scenario; AEMO's estimates of the industrial structure are based on forecasted growth rates of individual industries by scenario. Underlying demand for data centres is likewise estimated at a 1-digit ANZSIC classification, creating greater deviations in underlying demand than the deviation in the size of the economy. Household consumption is estimated based on DAE population forecasts by state and scenario, accounting for relative changes in the population outlook.

Fig. 34. Industry GVA CAGR by scenario, FY23 to FY58

	Agriculture	Mining	Manufacturing	Utilities	Construction	Services
Progressive Change	0.5%	0.7%	-0.9%	-0.2%	0.8%	1.7%
Step Change	0.8%	1.0%	-0.5%	0.1%	1.1%	2.3%
Green Energy Industries	1.4%	1.6%	0.0%	0.8%	1.8%	3.0%

Source: AEMO

Intermediate consumption

Data centre services are classified under ANZSIC industry subdivision 59 – Internet Service Providers, Web Search Portals, and Data Processing Services (web and data services). Australian National Accounts Use Tables provide detailed estimates of where web and data services are consumed at the household, government, and detailed industry levels.

Fig. 35. Ratio of use of web and data services to GDP, FY95 to FY23

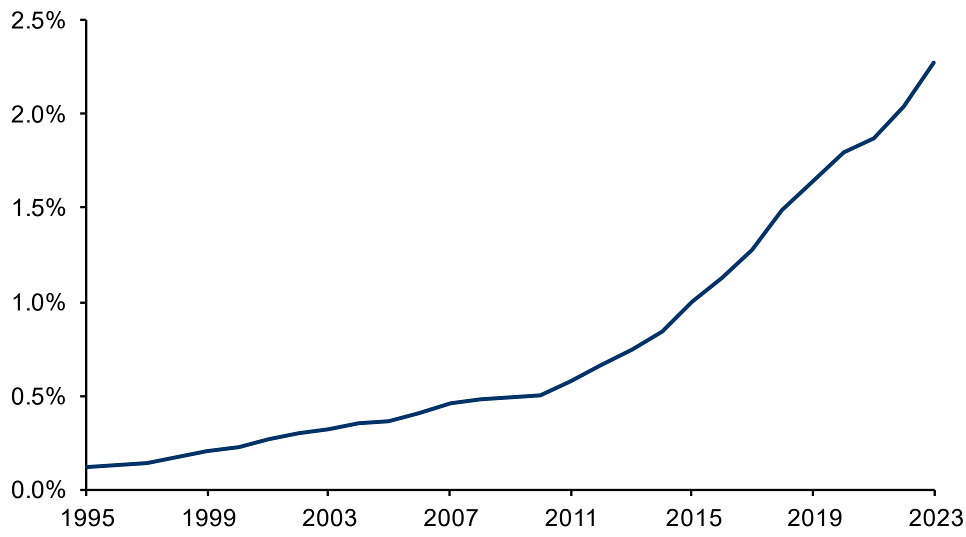

Source: ABS, Oxford Economics Australia

Aggregate consumption of web and data services as a ratio of GDP has risen steadily from 0.1% in FY95 to 2.3% in FY23. At no point over the past 28 years has this ratio declined; however, it has undergone several growth cycles fuelled by innovations that directed additional consumption to web and data services. Throughout the late 1990s and early 2000s, internet adoption drove peaks in the growth of the share, followed by a rise in cloud computing throughout the mid-2010s. Growth has begun to accelerate again since FY22, likely due to the continued integration of cloud services, the adoption and integration of AI and onshore requirements for data centre services use in certain sectors such as health and defence.

Fig. 36. Annual growth in consumption of web and data services

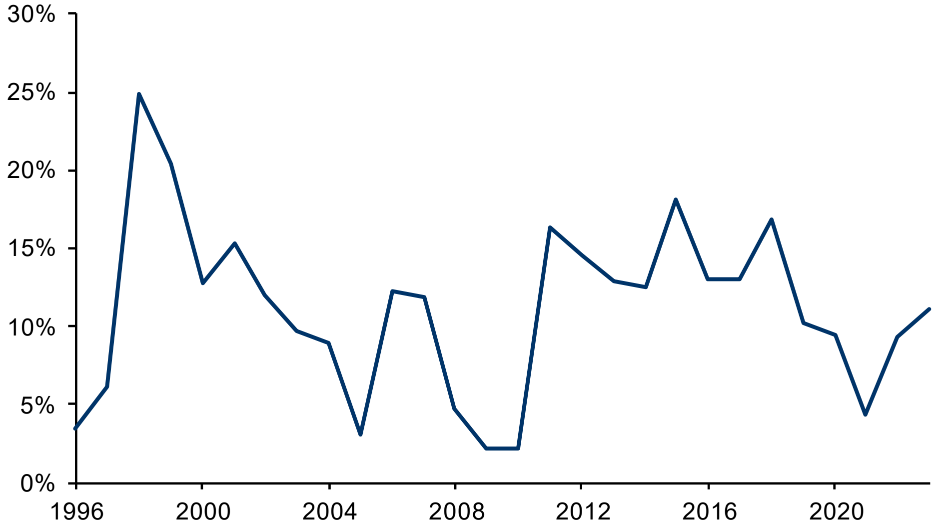

Source: ABS, Oxford Economics Australia

With the increasing digitisation of the economy, the ratio of web and data services is expected to continue growing, as novel technologies fuel additional consumption of these services, while remaining tied to the overall level of economic activity by scenario.

$$ratio\ of\ web\ and\ data\ services\ (\%) = \frac{Use\ of\ web\ and\ data\ services}{GDP}$$

However, the ratio of web and data service use to GDP is unlikely to continue increasing with the exponential shape that has characterised it over the past three decades, as these inputs of web and data services would begin to consume too great a share of business inputs to be economically feasible. It is assumed that, in the long run, the share will stabilise as the integration of digital services into the economy reaches a more mature stage; accordingly, a logistic curve is used to model the long-run relationship of the ratio.

Under *Step Change* and *Progressive Change* 2027 is adopted as the midpoint (x_0) of the ratio, as it theoretically precedes the peak adoption of AI, reflects the historical growth of earlier technologies in this domain, and sits at the midpoint of the historical dataset and the end of the forecast period.[51] The logistic curve is fitted by solving for *L* and *l* that minimises the mean-square errors in the historical data for 1-digit ANZSIC industries and the household and government sectors.

$$f(x) = \frac{L}{1 + e - k(x - x_0)}$$

[51] Under the GEE scenario, the midpoint is extended to 2029 to account for the faster pace of investment and economic growth over the forecast period and the likely acceleration in technological development that would follow.

GVA and the ratio of web and data service forecasts are adopted to estimate the real spend on web and data services over time (*t*) by scenario (*i*).

$$Web\ and\ data\ services_t^i = GVA_t^i \times ratio\ of\ web\ and\ data\ services_t^i$$

Intensity of electricity demand

The intensity of electricity demand is calculated historically as AEMO's estimate of total electricity demand from data centres from FY18 to FY24, over the web and data services level estimated by the ABS National Accounts Use Table.

$$Intensity\ of\ electricity\ demand = \frac{Data\ centre\ (GWh)}{Web\ and\ data\ services\ (\$)}$$

Over the past seven years, the GWh consumed per dollar of use has risen linearly. We adopt this linear trend over the forecast period to account for the increasing electricity consumption generated by one dollar of historical use of web and data services. The short time series of data for this variable indicates a medium degree of uncertainty surrounding it. If this relationship were to level off, we would see a reduction in overall electricity demand, equivalent to the demand for web and data services. Conversely, an acceleration would result in higher overall electricity demand.

Fig. 37. GWh per dollar of web and data services, Real FY23

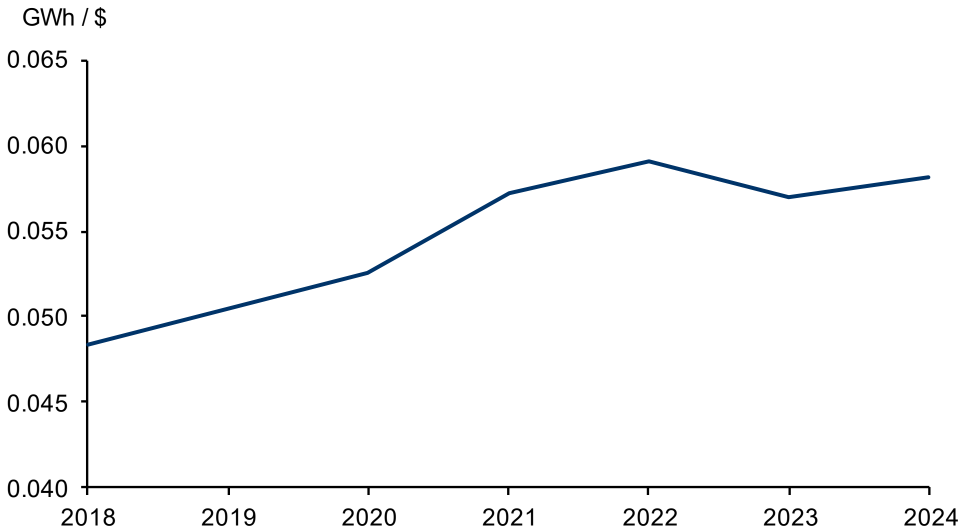

Source: Oxford Economics Australia based on AEMO data

The underlying demand for data centres is expressed by multiplying the real consumption of web and data services by the intensity of electricity demand:

$$Consumption\ (TWh) = Web\ and\ data\ services\ (\$) \times Intensity\ of\ electricity\ demand$$

Fig. 38. Economic approach key assumption sources summary table:

Input assumption	Source
Historical data centre electricity consumption	Key source: AEMO provided estimates of historical data centre electricity consumption.
GVA by industry by scenario	Key source: Deloitte Access Economics, Economic forecasts 2023/24.[52]
Population by state by scenario	Key source: Deloitte Access Economics, Economic forecasts 2023/24.[53]
Historic consumption of web and data services	Key source: Australian Bureau of Statistics, Australian National Accounts: Supply Use Tables 2022/23

AI demand

AI demand was assessed through three main components:

1. Short-Run Validation – Pipeline of AI-Driven Data Centre Projects

We identified and analysed the current and planned data centre developments in Australia that are specifically designed to support AI, machine learning, and other emerging technologies. This involved:

- Compiling a project-level pipeline of existing and upcoming data centres with a focus on AI workloads.
- Reviewing public disclosures, planning approvals, and industry announcements from hyperscalers, colocation providers, and enterprise operators.
- Estimating near-term additional load profiles based on facility size, design, and use-case (e.g. AI model training, inference, or hybrid workloads).

2. Long-Run Forecast – AI and Emerging Tech Adoption by Industry and Consumers

[52] Deloitte Access Economics (2024) *Economics forecasts 2023/24 Prepared for the Australian Energy Market Operator Limited*. Available here

[53] Ibid

To capture long-run impacts, a forecast model was developed linking technology adoption rates to expected compute intensity and resulting energy demand across both business and household use cases. This included:

- Sector-specific adoption curves for AI inferencing and training, by household and industry.
- Intensity multipliers based on known energy profiles of AI workloads (e.g. model training vs inference) and GPU-based processing compared to traditional CPU-based applications.[54] ,[55] [56]
- Differentiation by industry, based on anticipated uptake in sectors such as finance, health, retail, government, and media. [57] [58]
- A scenario of significant AI model training in Australia, for the *Green Energy Industries* Scenario

3. Top-Down Validation – Benchmarking Against Global and Domestic Estimates

We validated our bottom-up projections by comparing them to top-down estimates of total electricity demand attributable to AI and emerging technologies. This involved:

- Reviewing benchmarks from international agencies, industry reports, and Australian-specific sources.[59] [60] [61] [62]
- Testing the alignment of our model outputs against published projections of AI-related electricity demand growth as a share of total grid load. This testing showed that our bottom up inferencing estimates were within a reasonable range (+/- 20%) of the top down forecast.
- Performing sensitivity analysis to understand the range of possible outcomes under different technology adoption and efficiency scenarios. Key assumptions such as the pace of AI rollout, adoption rates and inferencing task energy consumption were stress tested with the expert panel and FRG.

[54] Luccioni and Jernite (2024) *Power Hungry Processing*. Available here
[55] EPOCH AI (2025) *How much energy does ChatGPT use*? Available here
[56] International Energy Agency 2024, Energy and AI. Available here
[57] Department of Industry, Science and Resources (2024) *Exploring AI adoption in Australian businesses.* Available here
[58] AiGroup (2024) *Technology adoption in Australian Industry.* Available here
[59] International Energy Agency 2024, Energy and AI. Available here
[60] Australian Financial Review (2024), '*How much will hungry data centres take out of the power grid?*' Available here
[61] Goldman Sachs (2024), '*AI is poised to drive 160% increase in data centre power demand*' Available here
[62] Avelar et al. (2023), 'The AI Disruption: Challenges and Guidance for Data Center Design', White Paper 110, Energy Management and Research Centre. Available here

Fig. 39. AI demand model framework

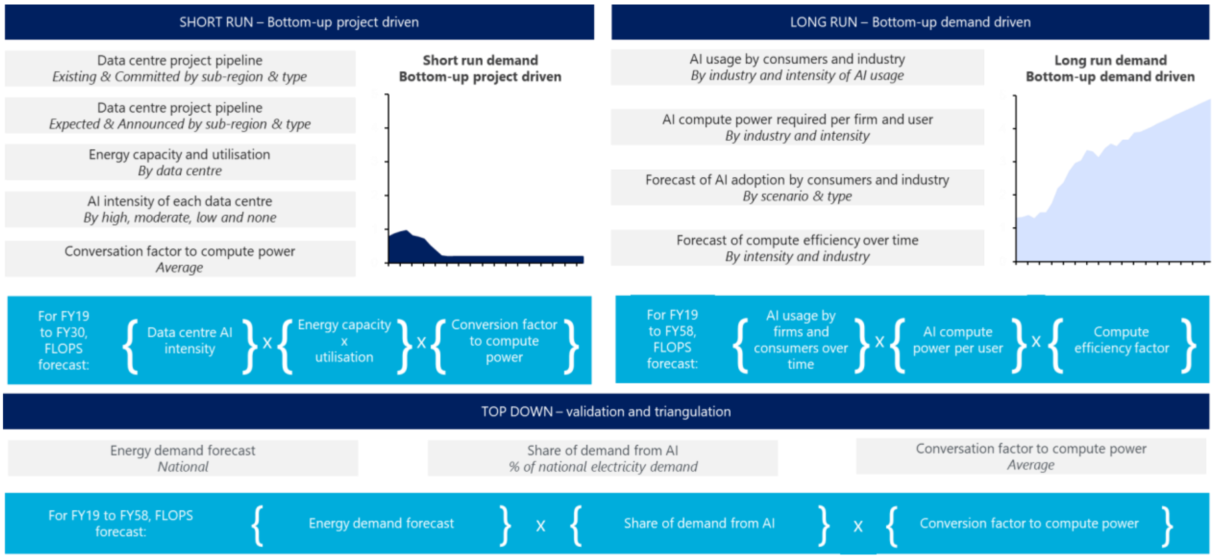

Fig. 40. AI demand key assumption sources summary table:

Input assumption	Source
Project pipeline of data centres in various stages of development	Key source: AEMO prospective project list compiled from NSP 2025 Standing Information Requests, LIL surveys and Data centre developers Validated by: Data Centre developments from OEA Building Construction Masterplan January 2025.[63]
AI intensities of data centres	Key source: Oxford Economics Australia, Data Centre Ownership and Location.[64] NEXTDC, Microsoft and other AI focussed firms, Type and location of AI focussed data centres.[65,66] Microsoft & Tech Council of Australia, Adoption rates of Generative AI.[67] Validated by: Top down estimates

[63] Oxford Economics (2025) *Australian Building Masterplan – January 2025.* Available here
[64] Oxford Economics Australian Building Masterplan (January 2025). Available here
[65] NEXTDC, 'S6 Sydney Data Centre', Available here
[66] Microsoft (2023) 'Microsoft announces A$5 billion investment in computing capacity', Available here
[67] Microsoft & Tech Council of Australia (2023) Australia's Generative AI Opportunity, Available here

Technology adoption	Key source: Department of Industry, Science and Resources, Adoption rates of technologies by industry[68]. AiGroup, Adoption rates of technologies by business size.[69] Count of businesses by size[70]. Validated by: Comparison to top-down estimates
Technology processing requirements for inferencing and training	Key source: Luccioni & Jernite; EPOCH AI, IEA, Average technology compute processing requirements by industry and household sectors, segmented by intensity.[71, 72, 73] Validated by: Short run estimates and top down estimation

Efficiency gains in data centre demand for energy

This section details the impact of adoption rates and compute intensity on energy demand based on the current known applications of AI and the energy intensity associated with them. This component reflects improvements in the energy efficiency of data centre infrastructure and equipment, as well as their resulting impact on energy consumption over time and the implications of these efficiency gains on the required consumption of electricity for the first wave of AI applications in the long-run.

Infrastructure efficiencies encompass improvements in data centre structural efficiencies and their subsequent impact on energy consumption over time. This is primarily based on Power Usage Effectiveness (PUE). Equipment efficiencies are based on server utilisation rates and computational energy efficiency (commonly referred to as Koomey's Law).

PUE

PUE is calculated as the ratio of the total facility energy consumed by a data centre to the energy used by IT equipment. Global industry averages of PUE for both traditional enterprises and hyperscalers have been steadily approaching the optimal range of 1.0 for the past 15 years. Holding computing demand constant, a better PUE can lead to lower overall energy consumption per data centre. However, hyperscale leaders such as Google and Meta historically have much lower PUEs than traditional enterprise data centres.

$$PUE = \frac{Total\ facility\ energy\ consumption}{IT\ equipment\ energy\ consumption}$$

[68] Department of Industry, Science and Resources (2024) *Exploring AI adoption in Australian businesses*. Available here
[69] AiGroup (2024) *Technology adoption in Australian Industry.* Available here
[70] ABS (2024), 'Counts of Australian Businesses, including Entries and Exits'. Available here.
[71] Luccioni and Jernite (2024) *Power Hungry Processing*. Available here
[72] EPOCH AI (2025) *How much energy does ChatGPT use*? Available here
[73] International Energy Agency 2024, Energy and AI. Available here

Server utilisation

Server utilisation measures the actual usage of servers (AE) over maximum capacity (MC). Global industry averages of server utilisation rates for both traditional enterprises and hyperscalers have typically been sitting in the 10-20% and 50-70% range, respectively. This alludes to the potential relationship between data centre type and utilisation (or efficiency). Consequently, server utilisation by type of data centre is analysed. Holding computing demand constant, a higher utilisation rate will reduce aggregate IT equipment requirements and subsequently total energy demand.

$$SU = \left(\frac{AE}{MC}\right)$$

Koomey's Law

Koomey's Law states that, historically, the energy efficiency of computing has doubled approximately every 1.5 years.[74] More recent research suggests that, like Moore's law, this relationship has slowed over recent years.[75] Energy efficiency (EE) of computing power can be measured as the number of computations per kilowatt-hour (kWh). Under *Step Change* and *Progressive Change*, growth in computations per kWh is set at the historic rate observed from 2008 to 2024 using the methodology specified in Prieto et al (2024), with computational efficiency doubling over 2.3 years. Under *Green Energy Industries*, where investment and economic growth outpace the central scenario, growth in computations per kWh is set at the Department of Energy's target of doubling every 2 years.[76]

$$EE = \frac{\text{Computations}}{\text{KWh}}$$

A fall in the PUE ratio, increases in server utilisation, and the continuation of Koomey's law would all have an offsetting effect on energy demand that will come about from increasing computational demand.

Implications of increases in computational energy efficiency

The energy demands of currently known AI applications are likely to diminish in the long run through increases in computational energy efficiency.

Mass adoption of known AI tools would substantially increase data centre energy demand if energy efficiency is not accounted for. Efficiency improvements in software, hardware, and infrastructure have the potential to offset a significant share of this uplift, mitigating strain on the energy system.

Green Energy Industries projects the highest energy demand, exceeding 20 TWh by FY58. This scenario uniquely includes large-scale AI model training activity occurring domestically in Australia, which significantly lifts energy consumption in the long term. However, when efficiency improvements are factored in, the adjusted line remains significantly lower than its unadjusted counterpart and stabilises over time.

---

[74] Koomey, J., Naffziger, S., (2015) *Moore's Law might be slowing down, but not energy efficiency.*

[75] Prierto et al. (2024) *Evolution of computing energy efficiency: Koomey's law revisited*. Available here

[76] Department of Energy (2024) *Energy Efficiency Scaling for 2 Decades.* Available here

The *Step Change* and *Progressive Change* scenarios forecast more moderate increases in data centre energy demand, peaking below 10 TWh by FY58. These scenarios assume widespread digitisation and electrification, but not the same scale of AI training workloads. Efficiency adjustments in these scenarios also lead to a considerable flattening of the demand curve, with *Step Change* and *Progressive Change* both stabilising.

Fig. 41. Energy efficiency adjusted data centre demand, current known mass AI adoption applications, by scenario, FY22 to FY58

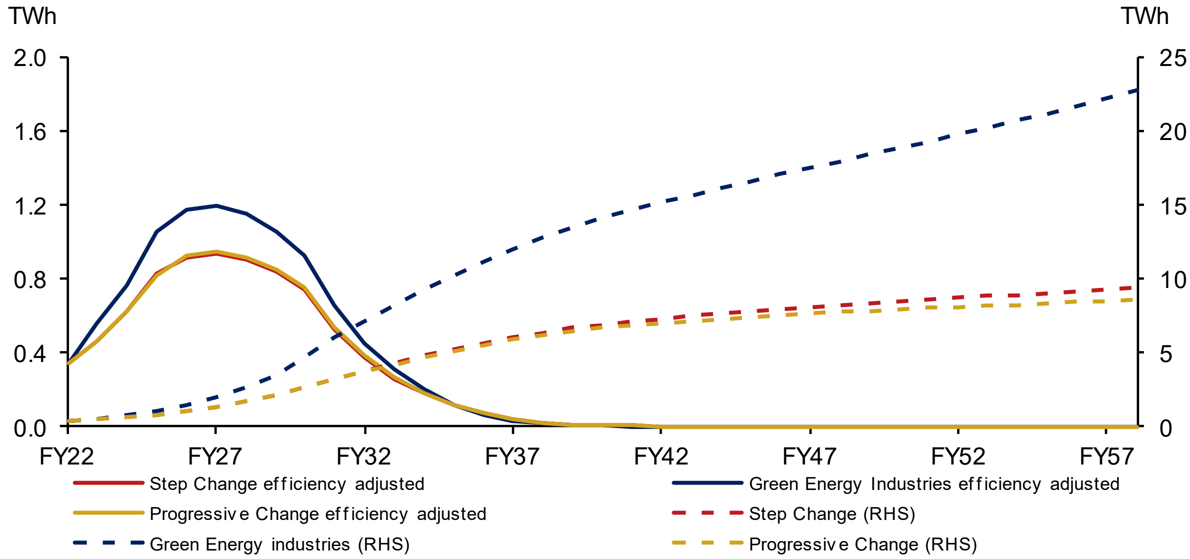

Source: Oxford Economics Australia

Fig. 42. Energy efficiency adjusted AI demand key assumption sources summary table:

Input assumption	Source
Historic PUE trends	Key source: DCCEEW[77], NABERS[78 79] Validated by: International benchmarks from best practice data centres[80], Statista[81], UpTime Institute[82] and Azura[83], and industry consultation[84].

[77] DCCEEW (2021) *International Review of Energy Efficiency in Data Centres.* Available here
[78] NABERS (2025) *Data Centres.* Available here
[79] NABERS (2025) *Find a current rating.* Available here
[80] Google (2025) *Growing the internet wile reducing energy consumption*. Available here
[81] Statista (2024) *Data center average annual Power Usage Effectiveness (PUE) worldwide from 2007 to 2024*. Available here
[82] Uptime Institute (2024) *Uptime Institute Global Data Center Survey 2024.* Available here
[83] Azura (2024) *Power Usage Effectiveness (PUE) Trends in Data Centers.* Available here
[84] Industry consultation highlighted that while globally, examples exist of hyperscaler data centres with PUEs of below 1.1, Australia's climate makes it less viable to achieve these results.

	Assumption: The average PUE data centre for non-hyperscalers falls from the current average of 1.53 to 1.34 (5-star NABERS rating). Hyperscaler average falls from 1.4 to 1.2 by FY58.
Increasing computations per kWh (Koomey's Law)	Key source: Evolution of computing energy efficiency: Koomey's law revisited.[85] Moore's Law might be slowing down, but not energy efficiency.[86] Validated by: Discussion with the expert panel to understand their view on the link between previous computational energy efficiency gains and potential future gains. Assumption: Growth in computational energy efficiency under *Step Change* and *Progressive Change* follows Prieto et al (2024) finding that it has increased annually at 35.3% p.a. since 2008. Under *Green Energy Industries*, increased global investment boosts computational efficiency; we align to the US Department of Energy semiconductor efficiency roadmap which aims for an increase in computational energy efficiency of 41.3% p.a. over the next two decades.[87]
Server Utilisation by type of data centre	Key source: 2024 United States Data Center Energy Usage Report.[88] Validated by: Consultation with industry to determine potential changes to server utilisation rates and any relative differences in server utilisation by type in Australia relative to the United States, alongside additional desktop research. Assumption: Server utilisation assumptions are adopted from the 2024 United States Data Center Energy Usage Report.

6.3 FINAL FORECASTS

Final consumption forecasts

The final forecasts represent an even weighting between Project and Economic Approach estimates.

$$Final\ consumption\ (TWh) = 0.5\ \times Project\ estimate\ (TWh) + 0.5\ \times Economic\ estimate\ (TWh)$$

[85] Prieto et al. (2024) *Evolution of computing energy efficiency: Koomey's law revisited*. Available here

[86] Koomey, J., Naffziger, S., (2015) *Moore's Law might be slowing down, but not energy efficiency.*

[87] United States Department of Energy (2024). *Energy Efficiency Scaling for Two Decades: Research and Development Roadmap.* Available here

[88] Berkeley Lab (2024) *2024 United States Data Center Energy Usage Report.* Available here

Peak demand

The monthly peak demand for each subregion is calculated by data centre type. AEMO provide estimates of load factor by subregion and average load factor by data centre type. We assume that in the long run, the load factor by type does not differ by region.

Data centre load factors by type are applied to consumption (MWh) by subregion (*r*) and data centre type (*i*), and then aggregated to forecast peak demand.

$$\text{Peak demand}_r^i\ (\text{MW}) = \Sigma\,(consumption\,_r^i(MWh) \div (no.days\ in\ year \times no.hours\ in\ day) \div load\ factor_r^i\ (\%))$$

6.4 DEMAND FLEXIBILITY

Demand flexibility refers to the ability of electricity consumers to adjust their consumption in response to supply conditions, such as price signals or network constraints. Internationally, data centres have been able to increase demand flexibility by shifting loads temporally, spatially and via accessing onsite generation and uninterruptible power supply systems. The opportunities and challenges for demand flexibility in data centres are discussed below.

Shifting demand temporally

Currently, data centre electricity loads are relatively stable throughout the day, with load factors averaging between 70% and 85% across all data centre types from FY22 to FY25. [89]

Fig. 43. Load factor by data centre type, FY22 to FY25

Hyperscaler	Co-locator	Edge	Telecom
75%	85%	70%	85%

Source: AEMO

Some data centres can shift some daily workloads to absorb excess renewable generation or support grid stability, with this being more feasible on an intraday basis. Shifting energy demand across days is viewed as less feasible.[90] Opportunities exist for some data centres to engage in *job scheduling,* whereby non-urgent computing tasks are shifted temporally to promote grid stability during times of stress and to increase their exposure to renewable energy demand by aligning to daily peaks in renewable energy generation however this is dependant on the type of tasks being conducted.[91]

Internationally Google has adopted demand scheduling initially to reduce the carbon impact of its compute load by scheduling non-urgent compute tasks to occur at times of peak renewable energy generation. However, despite initially being developed to target periods of high renewable supply, the

[89] Mandala (2024), *Empowering Australia's Digital Future*. Available here

[90] Bloomberg NEF (2021), *Data Centers and Decarbonization*. Available here

[91] Crozier & Liska (2025), *The Potential of Data Center Energy Demand to Provide Grid Flexibility.* Available here

ability to shift demand temporally has also been utilised to promote grid stability during periods of peak stress on several occasions, including in Europe, Taiwan, and the US.[92]

There are limitations on the capacity for data centres to shift demand temporally. Consultation with data centre providers indicated uncertainty around the economics of demand shifting as a flat load profile allows greater server use optimisation rather than increasing hardware requirement to respond to potential peaks generated from demand shifting. Furthermore, shifting demand is likely to be more viable for hyperscalers than co-locators. Co-location data centres host independent customers and have limited control over how or when computing tasks are executed, presenting an additional obstacle to increased flexibility in data centre demand.

Shifting demand spatially

Urgent workloads, which can be categorised as more *latency-dependent*, can also be geographically shifted across a fleet of data centres, enabling operators to optimise compute placement and minimise total energy use by operating the fleet at maximum collective efficiency.[93] Industry consultation highlighted that data centre operators are already actively considering these agglomeration benefits both in Australia and internationally, with Sydney and, more recently, Melbourne acting as hubs for data centres in Australia.

Spatial benefits for latency-dependent computing can be better leveraged by large, integrated data centre fleets with multiple centres to shift between. Opportunities to dynamically shift compute and optimise energy use are constrained for smaller standalone facilities.[94]

Backup onsite generation & UPS

A UPS (Uninterruptible Power Supply) in a data centre provides instant, short-term backup power to keep systems running during outages or power fluctuations, bridging the gap until generators start. Backup generation in data centres refers to on-site power systems—most commonly diesel generators—that automatically activate during grid outages to maintain uninterrupted power to critical IT infrastructure.

UPS systems can potentially support demand flexibility by reducing a data centre's net electricity demand on the grid for *latency-sensitive computing* that cannot be shifted temporally or spatially, or by discharging stored energy to provide a flexible resource to the grid. Like traditional BESS systems, they could make this viable through arbitrage – charging when surplus energy is available and discharging during peak periods however this requires specialised design and hardware.[95]

Microsoft data centres in Ireland have been approved to be connected to the grid, allowing their UPS systems to interact directly with the grid. The ability of Microsoft to deploy its UPS systems is argued to be mutually beneficial, as it allows Microsoft to unlock an additional source of revenue from its data centres while also providing additional grid stability as more variable renewable power generation

[92] Google (2023), *Supporting power grids with demand response at Google data centers.* Available here
[93] Crozier & Liska(2025), *The Potential of Data Center Energy Demand to Provide Grid Flexibility.* Available here
[94] Lindberg, Lesieutre & Roald (2022), *Using geographic load shifting to reduce carbon emissions*. Available here
[95] Crozier & Liska(2025), *The Potential of Data Center Energy Demand to Provide Grid Flexibility.* Available here

comes online. Furthermore, they note that these UPS systems will be built regardless of whether they can engage with the grid; as such, they are likely to be a cheaper source of additional BESS capacity than the equivalent marginal cost of newly installed BESS capacity.[96]

On-site back-up generation (typically diesel generators) needs to be 'turned over' to maintain warranty. As a result, there is an opportunity to 'switch on' during peak demand events.

Industry consultation conducted by Oxford Economics Australia highlighted that, currently in Australia, many large data centre providers are limited by their customer service agreements from using their UPS or backup generation for anything other than providing backup power for their data centres.

On-site backup generation presents an additional challenge in terms of compliance with emission limits and environmental standards, such as the EPA or NABERS, which are often required under planning regulations and government contracts. Additionally, uptime requirements and reliability standards can constrain opportunities to shift power use. Strict service availability requirements mean that even short interruptions or minimal performance degradations could compromise critical services and violate service-level agreements.

[96] Microsoft (2022), *Microsoft data center batteries to support growth of renewables on the power grid.* Available here

OXFORD ECONOMICS
AUSTRALIA

Sydney
Level 6
95 Pitt Street
Sydney
NSW 2000
Tel: +61 (0)2 8458 4200

Melbourne
Level 22
120 Spencer Street
Melbourne
VIC 3000

Global Headquarters
Oxford Economics Ltd
Abbey House
121 St Aldates
Oxford, OX1 1HB
Tel: +44 1865 268 900

London
Tel: +44 (0)20 7803 1400

Belfast
Tel: +44 (0)2982 635400

Frankfurt
Tel: +49 69 95 925 280

Paris
Tel: +033 (0)1 78 91 50 52

Milan
Tel: +39 02 9406 1054

Paarl
Tel: +27 (0)21 863-6200

New York
Tel: +1(646) 786 1879

Philadelphia
Tel: +1 (610) 995 9600

Boston
Tel: +1 (617) 206 6112

Chicago
Tel: +1 (847) 993-3140

Los Angeles
Tel: +1 (424) 303 3449

Florida
Tel: +1 (954) 916 5373

Toronto
Tel: +1 (905) 361 6573

Mexico City
Tel: +52 155 5419-4173

Singapore
Tel: +65 6850 0110

Hong Kong
Tel: +852 3974 8842

Tokyo
Tel: +81-(0)3-4588-2798

Dubai
Tel: +971 56 396 7998

Email:
info@oxfordeconomics.com

Website:
www.oxfordeconomics.com.au