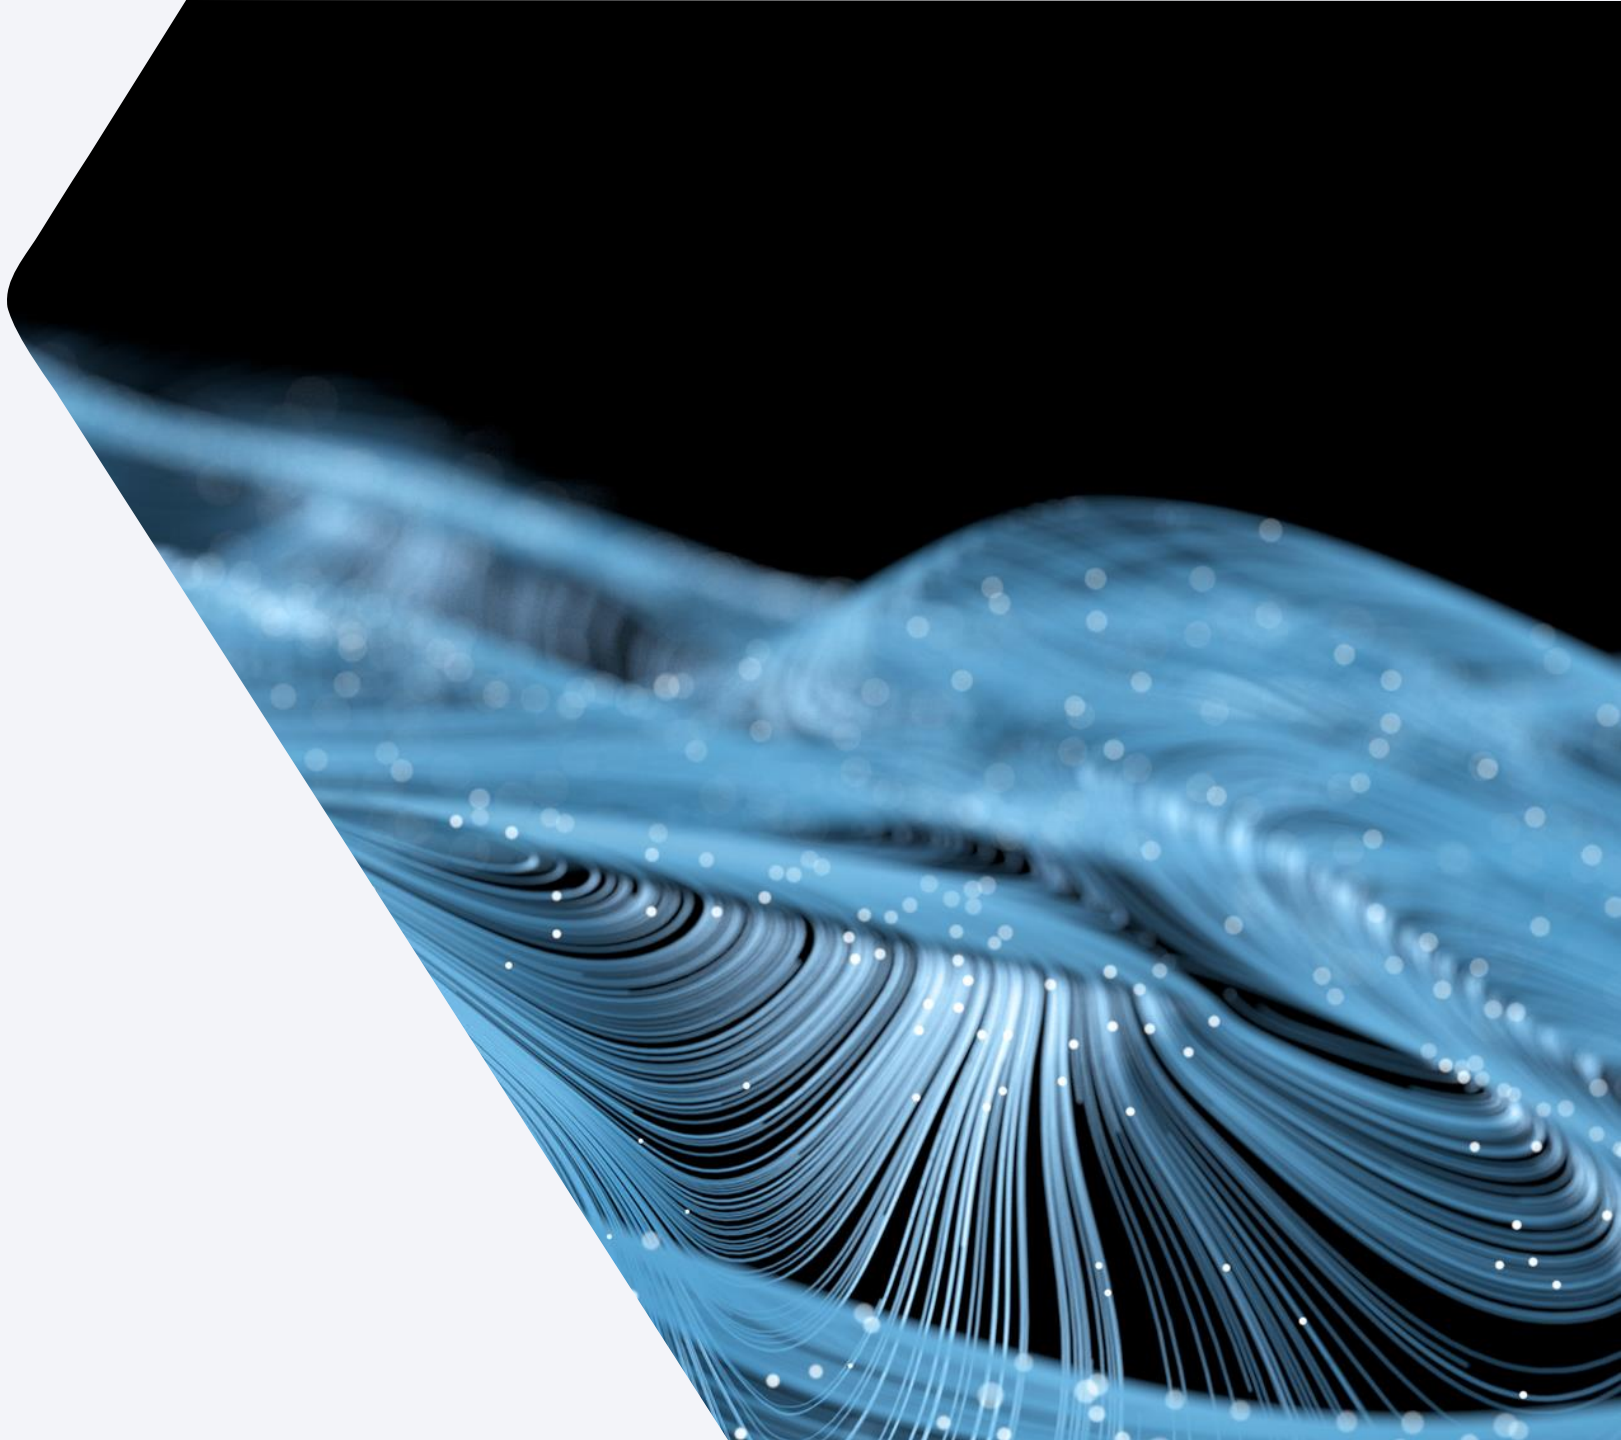
Access Partnership



From Cloud to Edge

The value of distributed AI in Australia

Commissioned by Qualcomm Technologies, Inc.

November 2026

Prepared by Access Partnership

www.accesspartnership.com

AccessPartnership

Disclaimer

All information in this report is derived or estimated by Access Partnership analysis using both non-proprietary and publicly available information. Qualcomm has not supplied any additional data, nor does it endorse any estimates made in the report. Where information has been obtained from third-party sources and proprietary research, this is clearly referenced in the endnotes.

With the exception of desktop researched claims which have been footnoted, all claims in the report have been derived based on Access Partnership modelling.

All calculations were done in USD and are based on the latest available data as of time of analysis in 2026.

Access Partnership (https://accesspartnership.com/), the frontier technology consultancy getting innovation into market faster, works with the public sector, corporations, NGOs, not- for- profits, and investors to identify opportunities and develop strategies to grow, improve service delivery, support economic prosperity, and generate tangible impact.

Access Partnership's Economics Strategy team is a leading economic and strategy consulting practice with deep experience across a number of topics in the digital economy and the future of skills, sustainability, and economic development.

Content

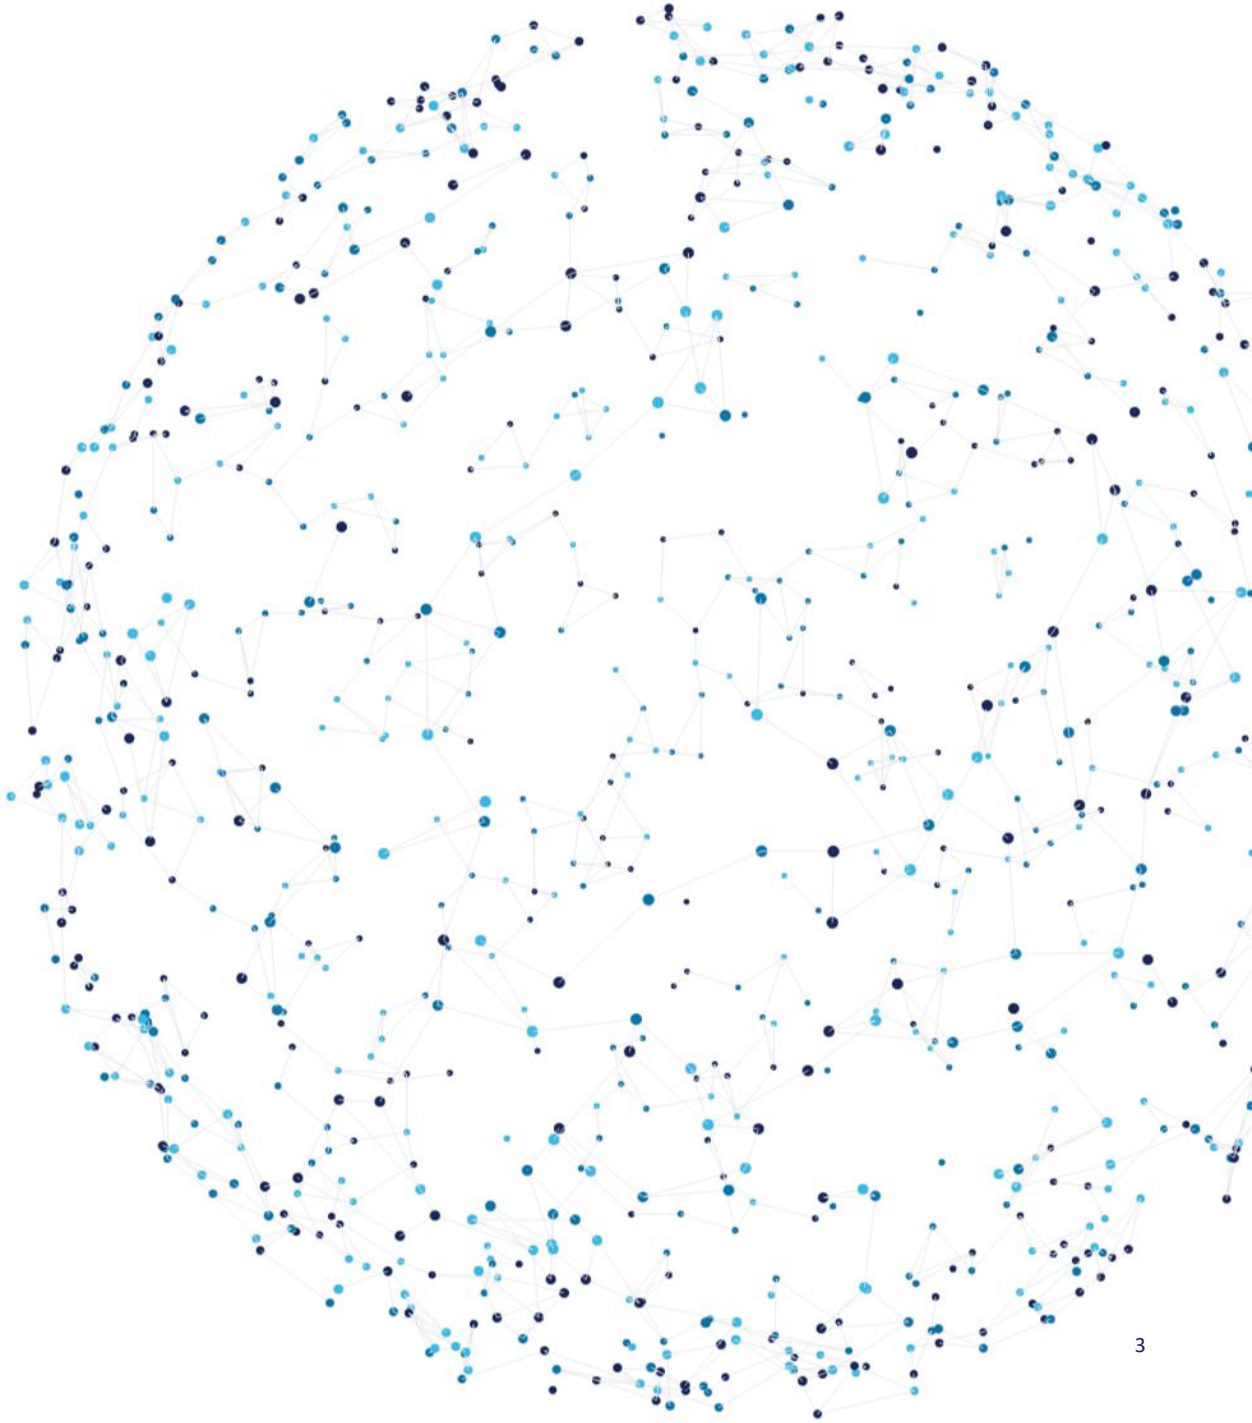
04 Executive summary and key highlights

05 **Chapter 1**
Australia's AI and digital landscape: An imperative for distributed AI

08 **Chapter 2**
Impact of distributed AI on Australia's economy and society

12 **Chapter 3**
Paving the way for distributed AI leadership with governmental support

16 Appendix



Access Partnership

Executive summary and key highlights

Australia's opportunity in AI will be shaped by how effectively it brings AI to where its economy works. Its largest sectors operate where networks are thinnest while current data centre investment trajectories mean Australia could face a 48% gap in AI inferencing capacity by 2030. Distributed AI can lift productivity, safety and resilience across its key sectors, generating more than $8.4B in economic value by 2035. Australia is already building AI compute at scale and formulating its AI policy. Next is making AI compute and adoption accessible throughout the economy through targeted policies, backed by around $373M in incremental investment over 2026-2035.

Australia can unlock opportunities with distributed AI…

Autonomous haulage and real-time hazard monitoring can make **mining more productive and safer**

>$2.8B
of economic value in 2035 from **connected and intelligent mining operations**

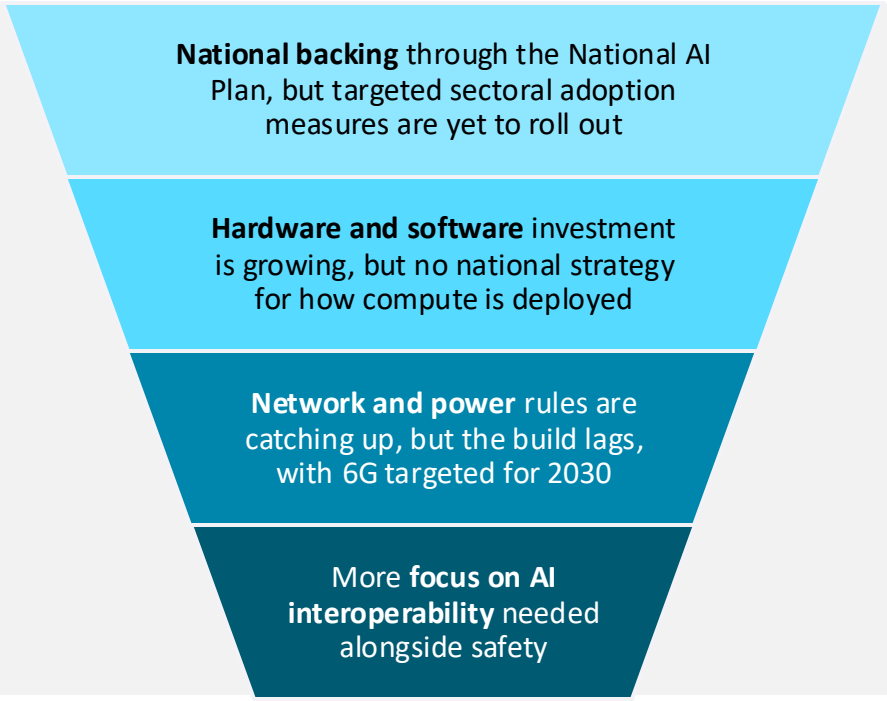
On-premise imaging and remote patient monitoring can deliver **faster, more accurate care**

>$2.9B
of economic value in 2035 from **AI-enhanced healthcare delivery**

Smart grids and remote asset inspection can make **energy infrastructure more efficient and resilient**

>$2.6B
of economic value in 2035 from **distributed energy and grid orchestrations**

…but must close gaps across the AI stack



Targeted policies will be key to unlock this

1. Recognise compute beyond data centres so edge and on-device AI can grow
2. Front-load 6G and device standards so distributed and physical AI can scale
3. Build open, interoperable AI device ecosystems into the national adoption strategy
4. Make vendor diversity and portability the default in government AI procurement

CHAPTER 1

Australia’s AI and technology landscape: An imperative for distributed AI



Access Partnership

AI is becoming a strategic anchor of Australia's economy – but how it will be deployed is just as important as how it will be adopted

Necessary action to expand AI is underway…

Productivity growth is an economic priority and AI can unlock it

Whole-economy labour productivity today is only around 1% above the 2015-2019 level, with multifactor productivity (MFP) levels below the 20-year average. The Productivity Commission estimates AI adoption could lift labour productivity by 4.3% over a decade.

Australia is taking action through initiatives like the National AI Plan

The Plan commits to (1) capturing the opportunity, (2) spreading the benefits, and (3) keeping Australians safe. Additionally, the government's *Expectations of data centres and AI infrastructure developer* outlines best practices for infrastructure investments.

AI infrastructure is arriving at scale – most in urban areas

Australia already has roughly 160 operational data centres. Adding to this is the development of "AI factories" with up to 2GW of capacity by 2027, with local data centre operators in tow.

… but constraints may hinder progress

- Australia's largest sectors, mining and construction (17% of GDP) **have the largest to gain from AI** (mining recorded its fifth annual decline in productivity; construction is 12% below its 1994 level) – however they mainly operate in **connectivity-constrained areas.**
- **AI needs more than compute to scale.** A survey by Ericsson measures only 32% of businesses in Australia having the necessary network infrastructure to fully realise the benefits of AI which requires reliable, high-capacity connectivity and the ability to process data close to (or at) where work is done.
- **Compute is concentrated in high-density cities and providing the necessary connectivity will slow adoption.** Half of Australia's data centres are in Sydney. Initiatives to expand connectivity require years of investment (e.g., NBN Co's fixed-wireless access, rollout of 6G services).

Missing from discussion is **how to bring AI to businesses and individuals at scale and at speed** to ensure Australia can capture the broad benefits of AI

A distributed AI approach will be necessary to fill the gaps left by cloud-only deployment

Complementing cloud with distributed AI enables rapid AI rollout at scale, with lower latency, stronger security and greater resilience, while hitting Australia's policy objectives



Distributed AI enables the AI ecosystem to meet Australia's resource goals

Shifting AI processing to on-device could reduce a query's total power consumption by around **90%** and increase water efficiency by **96%** as compared to having it processed centrally in data centres.



Lower latency enables always-on physical AI applications that require near-instant processing

Processing AI on or near the device **removes the round trip to data centres**, keeping physical AI (drones, robotics) responsive in connectivity downtimes and time-critical conditions



Stronger security and sovereignty through on-premise, on-device deployment

On-device AI processing keeps sensitive data local, **strengthening privacy and security and supporting Australia's AI sovereignty considerations**, while easing load on connectivity networks.

Cloud-based AI alone will not be sustainable to deploy AI at scale – given the following AI inference trends

40x Demand for AI inference compute in Australia could multiply by **more than 40 times** by 2030…

10x Meeting this demand is equivalent to multiplying existing data centre capacity in Australia by **10x**…

48% … and based on current data centre investment trajectories, Australia could have a **48% gap** in AI inferencing capacity in 2030

CHAPTER 2

Impact of distributed AI on Australia's economy and society



Access Partnership

Specific distributed AI use cases can unlock transformative applications across Australia's priority sectors, boosting productivity

Distributed AI use cases and solutions that are significant to Australia *(and scope of potential market)*

Connected and intelligent mining operations
- Autonomous haulage & rail systems
- Remote fleet management
- IoT equipment & predictive maintenance
- Real-time safety & hazard monitoring through computer vision

Smart agriculture & precision farming
- On-device crop & pest detection
- Precision irrigation & variable-rate spraying
- Soil & nutrient optimisation through real-time environment tracking
- Autonomous field machinery

Autonomous defence & critical security capabilities
- On-board AI and low-power edge compute within drones and autonomous vehicles
- Computer vision threat detection & situational awareness
- Secure and resilient communications AI-enabled encryption

Smart public service delivery and safety
- On-site computer vision for public safety and crowd management
- Welfare support and service coordination in satellite government offices
- On-premise AI processing for sensitive data

AI-enhanced healthcare delivery
- AI-powered assistant and robots for doctors, nurses
- On-premise diagnostic imaging
- Edge-powered processing of patient data for rapid intervention
- AI-enabled patient tracking (cameras and wearables)

Intelligent construction & site management
- On-site hazard and near-instant BIM clash detection through cameras for safety
- Real-time equipment & telemetry monitoring for digital twins
- Connected machinery & computer vision for asset inspection

Distributed energy & grid orchestration
- Intelligent grids and power distribution
- AI-enabled asset monitoring and inspection
- Camera and sensor-based based monitoring for remote assets

AI-enabled consumer & commercial devices
- On-device assistants for personal devices
- Private inferences using personal data and calendars
- Real-time translation & productivity tools for everyday use

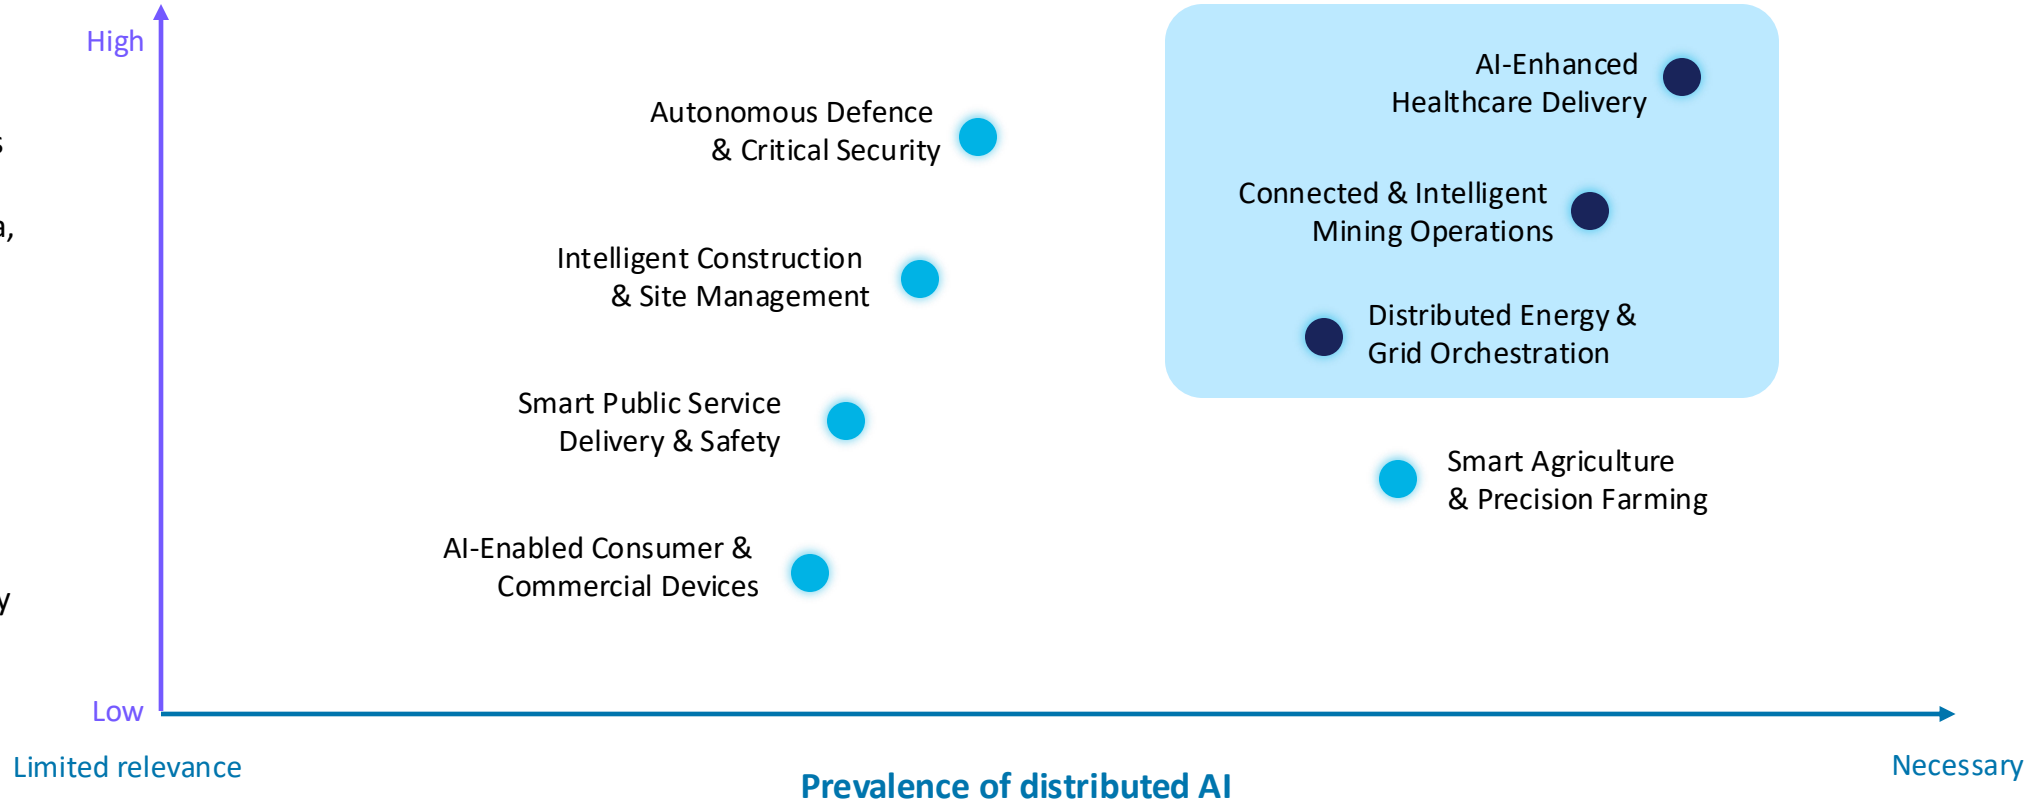
Distributed AI brings AI to businesses and individuals at scale, ensuring that Australia can capture the broad benefits of AI

Impactful use cases for Australia that will require distributed AI

Impact on economy

This dimension analyses the scale of a use case's impact using public data, such as:

- Impact on the socio-economic challenge
- Market size
- Breadth of impact
- Quality of current infrastructure
- Governmental priority



This dimension relies on assessments of the scale of distributed AI deployment due to factors such as real-time or near-instant processing, reliance on sensitive data, level of personalisation needed, or ability to leverage lightweight AI models

Source: Access Partnership analysis

Australia can derive more than $8.4B worth of economic value from distributed AI applications

Impact of **three use-cases** – driving improved productivity in the mining sector, improving healthcare outcomes and service efficiency, and enabling more efficient, resilient energy infrastructure.

	>$2.8B **Connected & Intelligent Mining Operations**	>$2.9B **AI-Enhanced Healthcare Delivery**	>$2.6B **Distributed Energy & Grid Orchestration**
Equivalent to:	**1.1%** of the mining sector's gross value added (GVA) in 2035	**1.1%** of healthcare expenditure in 2035	**4.2%** of Electricity, Gas, Water and Waste Services' GVA in 2035
Examples:	Applications such as autonomous haulage & rail systems, IoT equipment & predictive maintenance, and real-time safety & hazard monitoring through computer vision	Applications such as ambient clinical documentation tools, robotics-assisted surgeries, remote patient monitoring systems, and fall prevention camera systems	Applications such as smart grid technology, IoT equipment & predictive maintenance, and AI-enabled and remote inspection tools
Impact created by:	• Greater productivity • Greater asset reliability • Fewer safety incidents	• Greater productivity • Faster and more accurate diagnostics • Improved patient monitoring	• Energy efficiency gains • Greater asset reliability • Reduction in annual disaster cost and greater worker safety

Note: Economic Value refers to the increased revenue or cost savings that businesses generate from increased market reach or increased productivity. It is not directly attributable to GDP.
Source: Access Partnership analysis

CHAPTER 3

Paving the way for distributed AI leadership with governmental support



Access Partnership

Australia can pave the way for distributed AI leadership with governmental support

We conservatively estimate that the incremental investment required is around $373M over 2026-2035.

Estimated investment required for effective distributed AI deployment (2026-2035)

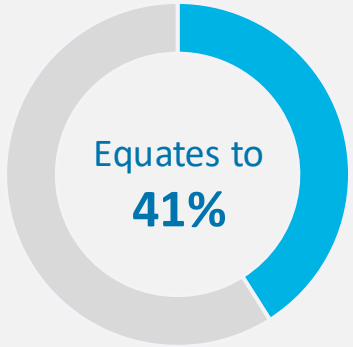
~$373M

(or ~$37M per annum over a decade)

Includes funding for:

- Sovereign AI capability & defence
- Industrial AI adoption
- Edge AI infrastructure & connectivity
- Semiconductor & edge hardware ecosystem
- Workforce & skills for distributed AI
- Regulation & digital governance

If the required investment were spread over a decade (~$37M per annum), the annual investment…



of Australia's annualised AI funding commitments to date (~$90M per annum)[1]

[1] AI funding commitments refer the National AI Plan investments and AI Accelerator programs annualised over their commitment periods (~$90M per annum in total). Figures are correct as at November 2026 and do not reflect any subsequent cancellation or deferral of announced data centre plans.
Source: Access Partnership analysis

Australia is formulating its AI policy, but missing the pathways to make AI compute and adoption easily accessible

Initiatives already underway can be expanded to bring distributed AI into each policy foundation

	Australia's policy foundations	The distributed AI stack	How distributed AI builds onto these foundations
AI deployment	**AI has national backing – next is implementation.** The National Robotics Strategy (2024) and National AI Plan (2026) provide high level support for AI development, but targeted sectoral measures have not been rolled out, limiting AI adoption.	On-device & physical AI devices Industrial AI deployment	**Taking AI to the point of work** by extending sectoral adoption support to on-device and edge deployment in key sectors such as mining, agriculture, construction and health
Hardware & software	**Backing wider adoption of AI.** The AUD 1 billion-National Reconstruction Fund supports critical tech including AI and robotics; MOUs with Anthropic and Microsoft provide infrastructure backing; Buy Australian AI Partnership (2026) lays out domestic AI capabilities – but no national strategy to how compute is deployed.	Data & Models (*Software*) Compute processing & chips (*Hardware*)	**Promoting competition and vendor diversity** by recognising edge and on-device compute as part of national AI capacity; keep the AI stack open to multinational investment so buyers retain genuine choice
Infra-structure & utilities	**Rules are catching up, but the build lags.** New data centre energy rules could take up to three years to implement; 6G spectrum is under discussion with potential commercial rollout targeted by 2030; expansion of mobile networks into harder-to-reach areas is still underway.	Network infrastructure Power infrastructure	**Tying infrastructure investment to outcomes** so AI-ready networks reach where economic activity is; backed by active Australian participation in discussions on 6G standardisation to keep pace with standards
Regulation	**Agile, technology-neutral governance.** Existing initiatives emphasise safety – an AI Safety Institute, targeted rules on smart device security, and globally aligned safety standards – but little effort on interoperability.	Standard & governance[1]	**Extending governance from safety to interoperability:** open ecosystems and data portability let end-users stay vendor-agnostic and minimise lock-in

1. Cuts across all layers

Source: Access Partnership analysis

Australia can capture more of the AI boom it is already funding by deciding where AI runs

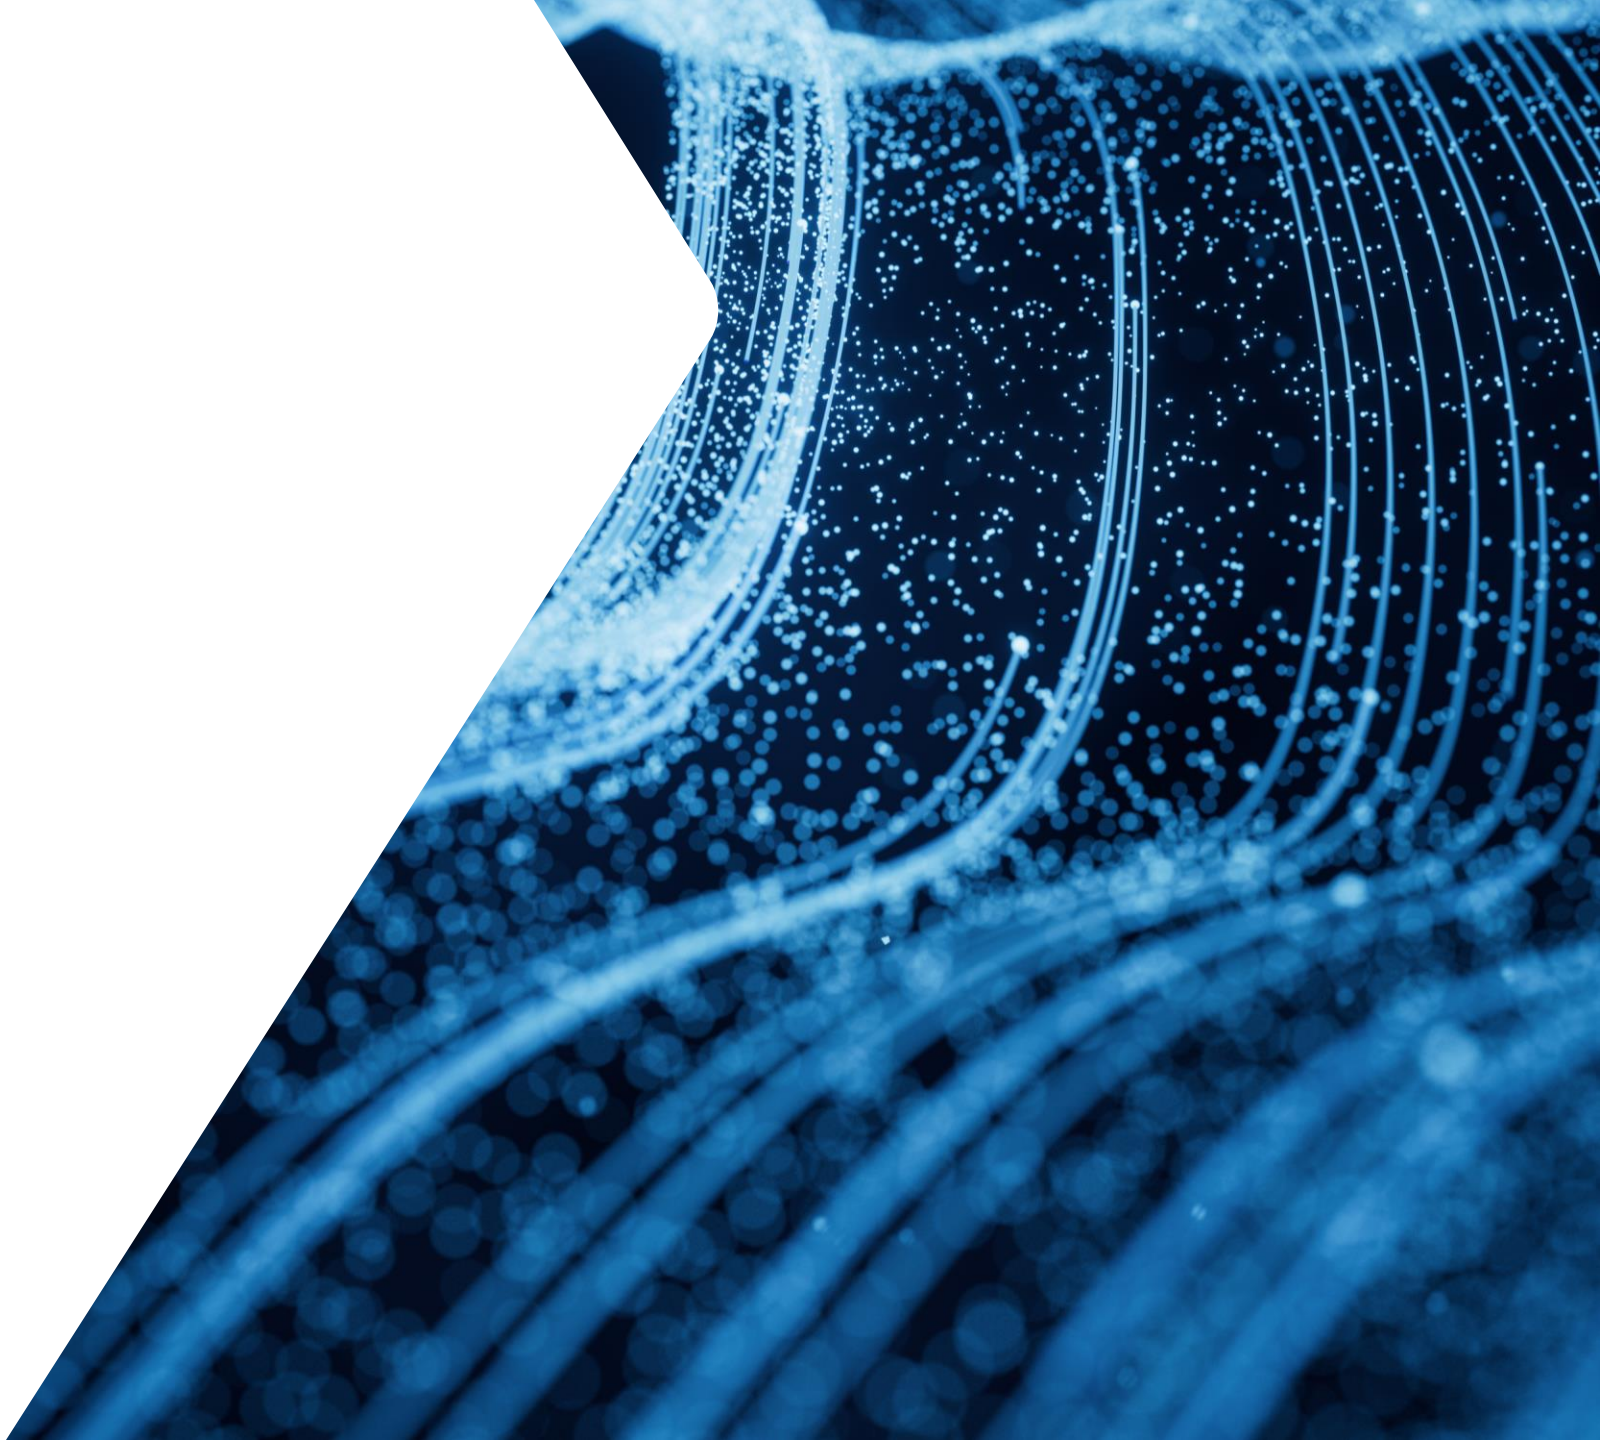
Four priorities, spanning compute beyond data centres, 6G and device standards, consumer adoption and procurement, would convert Australia's AI investment into productivity and local value

1. Recognise compute beyond data centres so edge and on-device AI can grow	The government is currently exploring tiered implementation of AI standards based on data centre thresholds at the facility level (primarily based on power use). Australia should recognise that large, centralised data centres are not the only path to expanding AI compute, and create a lighter regulatory pathway and targeted investment incentives for enterprise and edge AI, giving businesses a faster, lower-cost option to deploy AI closer to where data is generated.
2. Front-load 6G and device standards so distributed and physical AI can scale	Australia should commit to a pathway for adding mobile data to the universal coverage obligation (currently just up to voice and SMS), set a decision date for commercialisation of 6G so the ecosystem can anticipate scenarios for ubiquitous 6G connectivity, AI, and communication & sensing. Alongside this, regulatory initiatives should be extended from safety testing into device-level interoperability and certification.
3. Build open, interoperable AI device ecosystems into the national adoption strategy	Almost half of Australians have already used generative AI but affordability constraints have led to one-fifth of the population being digitally excluded. Government digital-inclusion funding and device programs should explicitly recognise devices with on-device AI compute (smartphones, PCs, wearables), which run AI privately without constant connectivity and create more dynamic AI-device ecosystems by promoting vendor diversity, application interoperability and data portability as key consumer requirements.
4. Make vendor diversity and portability the default in government AI procurement	Future whole-of-government AI and cloud arrangements should require data portability, open interfaces and the ability to run more than one provider, so agencies can switch without re-platforming. On-device AI capability should be treated as a specified government requirement, assessed against the whole-of-life value-for-money test already in procurement rules, including privacy and security benefits. With two-thirds of businesses using AI but fewer than one in ten using it significantly, Australia should promote adoption in depth and direct adoption support toward the physical, dispersed sectors where AI runs at the point of work.

Access Partnership

Source: Access Partnership analysis

Appendix



Access Partnership

Key definitions: the AI computing continuum



Cloud AI

AI training and inference executed in hyperscale data centres. Cloud AI provides the capacity needed for model training, global model hosting, orchestration, and the heaviest workloads, ensuring consistency across large user populations, but requires data to travel to and from centralised servers.



Edge AI

AI inference executed at or near the network edge (i.e., factory gateways, hospital edge servers, or micro-data centres) rather than in centralised cloud. Edge AI delivers real-time responsiveness, reduces backhaul costs, and keeps data within local or organisational boundaries.



On-device AI

A subset of edge AI, on-device AI operates inference directly on end-user hardware (smartphones, PCs, robots, industrial endpoints) without connectivity, which enhances privacy and offline resilience with low latency.



Distributed AI

The hybrid computing continuum in which AI workloads are dynamically partitioned across on-device, edge, on-premises, and cloud tiers based on latency, privacy, bandwidth, energy, and cost requirements. No single tier can meet the demands of real-time, multimodal AI at scale; distributed AI treats the tiers as one tightly integrated system, with each handling the workloads it serves best.



Sovereign AI

A nation's ability to independently build, deploy, and govern AI using its own infrastructure, data, and talent that advances its laws, cultures, and values.



Physical AI

Unlike the compute tiers above, physical AI describes a class of applications: AI systems that perceive, reason, and act in the physical world in real time – robots, autonomous machines, drones, and connected vehicles. Physical AI must operate, learn, and scale reliably under real-world conditions, on-device and at the edge, coordinated with cloud-based learning across fleets.

Access Partnership

Source: BCG (2026), 6G: The Network for the Future of AI and Immersive Connectivity, and Access Partnership analysis

Top-down framework for estimating AI compute inference

Overall description of approach taken

- To describe **total compute available**, the total compute capacity[1] available in data centres (public and on-premise) dedicated to AI up to 2030 as proxy (including compute used for training and inference) was modeled based on publicly available data.
- To measure the **potential on-device AI compute capacity, the overall AI compute capacity demanded** was also estimated – the "gap" between demand and current supply represents the total compute that will happen away from data centres.



Compute capacity
(AI DC supply)

Key assumptions for projecting supply of data centre (DC) compute capacity

- Projected compute capacity of DCs dedicated to AI workloads using: (a) total DC power capacity in gigawatts (GW) and; (b) total AI compute[2] to estimate a *compute-per-GW ratio* that changes over time
- **Assumption**: Of total cloud workloads. the share dedicated to AI would grow from 14% in 2025, estimated by Goldman Sachs, to 70% by 2030 based on McKinsey's projections. Additionally, around 40% of workloads are dedicated to training-related compute based on analysts from Centre for the Governance of AI



Compute demanded
(AI demand)

Key assumptions for projecting demand for AI compute

- Estimated total demand for AI by leveraging estimates by McKinsey that projects "annual FLOP demand" from use of **generative AI** to be 0.2 QFLOPs in 2024 – growing to be 25 QFLOPs in 2030
- McKinsey also estimates that by 2030, 40% of total compute demanded will be for generative AI applications and use cases. Additionally, we assume that of total demand by businesses and consumers, demand for training applications will range between 10% and 25% of total compute

1. Given limited research on TOPs (tera operations per second) used to describe neural processing units, compute is measured in terms of FLOPs (or floating point operations per second) as a proxy of potential capacity to complete calculations.
2. Estimates on the total AI compute were estimated in April 2025 by the AI Futures Project, a team of AI developers and researchers. See AI Futures Project (2025), *AI 2027*. Available at: https://ai-2027.com/about

Top-down framework for country-specific AI compute projections

Overall description of approach taken

- Existing modelled figures on global compute capacity and demand were apportioned into country-specific figures. This illustrates the relative levels of installed DC capacity (AI DC supply) and AI demand for each country, with the assumption that DC capacity available within each country is dedicated to its domestic users.



Country-level compute capacity (AI DC supply)

Key assumptions for projecting country-level supply of data centre (DC) compute capacity

- Existing modelled figures were apportioned into country-level figures based on a count of DCs (Data Centre Map) and DC energy usage (IEA). These two factors served as proxies for supply of DC compute capacity.
- Growth in Australia's supply of DC compute capacity is based on growth of domestic DC capacity (BCG).



Country-level compute demanded (AI demand)

Key assumptions for projecting country-level demand for AI compute

- Existing modelled figures were apportioned into countries based on the relative usage of AI tools (aitools.xyz). This is measured by utilising web traffic analysis software to track the number of average visits[1] to a database of over 10,500 AI tools, applications, and platforms. This serves as a proxy for demand from both businesses and consumers.
- Growth in Australia's demand for AI compute is based on global growth rates (McKinsey). Global growth rate was further adjusted based on the respective CAGR of AI spending as measured by market research firms, including IDC.

1. The number of average visits refers to the monthly average visits for each month between January and December 2025.

Access Partnership

accesspartnership.com

Need a problem solved?

Our dedicated experts, located around the world, are here to help.

Europe

London

Thomas House, 84 Eccleston Square
Office 5.06, 5th floor,
SW1V 1PX, London
United Kingdom

+44 20 3143 4900
london@accesspartnership.com

Brussels

8th Floor, Silversquare Europe
Square de Meeûs 35
B-1000 Brussels
Belgium

brussels@accesspartnership.com

North America

Washington DC

1300 Connecticut Avenue NW,
Suite 250
Washington, DC 20036
USA

+1 202 503 1570
washingtondc@accesspartnership.com

Asia

Bangkok

188 Spring Tower
11th Floor, Unit 106, Phayathai Road
Thung Phayathai, Ratchathewi, 10400
Thailand

+ 66 (2)-8216148

Hanoi

19th floor, Tower 1
Capital Place Building
No 29 Lieu Giai Street
Ngoc Khanh Ward, Ba Dinh District
Vietnam

Jakarta

Revenue Tower 21st Floor
Unit 104 SCBD Lot 13, Jl. Jend. Sudirman Kav. 52-53
Provinsi DKI Jakarta, 12190
Indonesia

+62 21 5020 0949

Kuala Lumpur

Common Ground Q Sentral
Level 39, Unit 39-02 (East Wing), 2A,
Jalan Stesen Sentral 2, Kuala Lumpur Sentral,
50470
Malaysia

Manila

CG8ROCKWELL Level 21, 8 Rockwell, Hidalgo Dr.,
Rockwell Center, Bgy. Poblacion, 1210
Makati City, Metro Manila
Philippines

Singapore

12 Marina View
#11-20, Asia Square Tower 2
Singapore 018961

+65 8323 7855
singapore@accesspartnership.com

Follow us