

The AI Era Begins: The Biggest Technology Shift of Our Lifetime

Everyone remembers the first time that a smartphone app first wowed them, whether it was the way the interactive maps got you from A to B, or when Shazam identified the song that was playing in a bar or knowing that you could do your banking in the palm of your hand wherever and whenever you wanted. The smartphone app ecosystem revolutionised the way consumers and businesses interacted with technology and led to billions of dollars of investable opportunities to arise. It ultimately created the mobile era in computing and enabled Apple to become the biggest company in the world.

However, it was not the fact that Apple made the iPhone that grew their own business; it was the fact that they created a hardware and software model ultimately building an app ecosystem that helped them go from US\$50 billion to US\$3 trillion, making them one of the few big winners of the stock market.

At Munro Partners, we believe that we've exited the mobile era and are firmly in the next era of computing - the Artificial Intelligence era – the biggest technology shift of our lifetime.


The app ecosystem changed the way consumers interact with technology. However, many of the applications that are now 'front page apps' for the iPhone were created several years after the first iPhone was released, and we expect a similar pattern to occur with AI applications. We expect every business in the world to leverage AI to reduce their own cost structure or improve productivity for their customers. We believe there is potentially going to be an AI app or AI 'agent' for everything. These applications have to work in an ecosystem, they need both hardware and software infrastructure to exist, like the smartphone apps did back in the mid-2000s. Our job as a growth investor is to scour the world for the few winners that will enable the big structural change that is AI. The companies that truly win from enabling AI, in our view, will experience faster and more durable earnings growth than the market anticipates.

AI Infrastructure

When we think about the AI stack, we see lots of opportunities for investment. In many of the structural thematic we have invested in over the years, we aim to find the key enablers and beneficiaries of that structural change. Our approach to investing in artificial intelligence has been no different, and as such we have focused on the critical companies towards the bottom of the AI stack pyramid that enable all the AI applications that consumers and businesses will interact with in the future, that is the hyperscale companies (including Microsoft, Amazon, Alphabet & Oracle), and the semiconductor companies.

THE AI STACK – LOTS OF PLAYERS BUT ONLY A FEW ENABLERS

MUNRO



For illustrative purposes, the companies shown may or may not be held in the Munro Funds.

Similar to Apple, Nvidia has positioned itself as the key enabler of AI technology by being the hardware-software platform for AI. As the world's leading manufacturer of graphics processing units (GPUs) and the pioneer of accelerated computing, Nvidia produces both the chips that power AI and CUDA, the software that developers use to program them.

In our view, we are now approximately two years into the AI revolution, and all of Nvidia's customers, principally the hyperscalers, are still spending at breakneck speed in order to deploy large computing clusters for training large language models. The hyperscale customers, including Alphabet, Amazon, Microsoft and Oracle, are spending hundreds of billions of dollars each per year to build AI data centres, or AI compute infrastructure, because they can see tangible benefits for their customer bases. While not a pure hyperscaler, Meta provides a clear example of the benefits of this spending: Meta's performance advertising transitioned to GPUs which has led to more effective advertising, and in turn has driven higher revenues for the company.


The first iteration of Chat GPT, the large language model we are all familiar with, was built on a 'cluster' of approximately 10,000 GPUs, and was completed with a capex of approximately US\$150m. As the hyperscalers started seeing the benefit to their customers from the use of large language models, and the demand from corporates globally, they accelerated their investment and are now working on building 'clusters' of one million GPUs and beyond and spending upwards of US\$50bn to do so. Scaling up these large computing clusters allows for more powerful large language models to be trained and deployed in lots more use cases and industries around the world. As evidence of the scale of investment, Oracle and OpenAI are spending US\$500bn over 4 years to build Stargate, the largest AI training cluster in the world, and Meta is reportedly planning a US\$200bn AI data centre project in the Southeast of the US.

As companies rapidly scale up their AI computing resources, Nvidia is poised to capture the lion's share of this opportunity, leveraging its dominance, with over 90% market share in AI accelerators, and we believe this isn't slowing down in the near term. Overall, the company frames this as a trillion dollar opportunity. Hyperscalers continue to build out these clusters of compute, of which they can address approximately 70% of this trillion dollar opportunity, due to their focus on system-level solutions, that goes well beyond the GPU itself. However, the opportunity for Nvidia doesn't end with the hyperscaler capital expenditure. What the company talks about is several additional trillion dollars of opportunity in addition to the hyperscaler spending on AI. This includes government spending on AI, and to date we have seen announcements in both Europe and the Middle East where governments are wanting to invest and build their own AI infrastructure. The additional opportunity will also include new industries, such as robotics, which we believe will be a major industry in years to come, and finally, it also includes the concept of 'AI factories'. This involves industry players developing their own dedicated AI data centre and investing in their own AI infrastructure. As an example, a pharmaceutical company such as Eli Lilly may invest and build their own AI dedicated facility, exclusively for company use in drug discovery. What is clear, is that Nvidia has a long runway of growth ahead as the key enabler of artificial intelligence.


LARGE CLOUD SERVICE PROVIDERS ARE DRIVING AI-RELATED CAPEX

MUNRO

Cloud capex



Nvidia revenue



Source: Munro estimates, Google, Amazon, Microsoft, Meta, Nvidia as at September 13, 2024. Nvidia data is at January year end. Graphs' vertical axes are in millions (USD). Microsoft * – includes capital leases. For illustrative purposes only, companies may or may not be held by the Munro funds. Past performance is not an indication of future performance.

Further to Nvidia providing the hyperscale companies with the compute infrastructure required for building artificial intelligence capabilities, there are also alternative solutions that will play a role. Companies including Alphabet and Meta also rely on alternative chips, which are built with custom silicon. These are chips specifically designed by the hyperscale companies with the help of companies like Broadcom. Because a company like Meta or Alphabet have a vast array of needs for AI compute power, custom silicon chips are used for particular purposes. These solutions offer less performance but are a cheaper solution for certain use cases. We believe companies such as Broadcom, which help design the custom chips, have a long runway of earnings growth as this opportunity plays out.

Whether a hyperscaler is sourcing their semiconductor chips from either Nvidia under a GPU-based architecture or Broadcom under a custom architecture, Taiwan Semiconductor Manufacturing Company (TSMC) plays an important role in both cases. TSMC is what is known as a 'foundry', the company ultimately responsible for building the chips to be used to power AI applications. The company also manufactures chips for lots of other end markets, for example, chips designed for smartphones, industrial applications and a range of computers. However, as guided by management in January 2025, the portion of TSMC's revenues related to AI is set to grow at a mid-40% compound annual growth rate over the next 5 years.

These three companies, Nvidia, Broadcom and TSMC, all exhibit the characteristics that we look for in any growth investment. All three companies are clearly growing their revenue, significantly faster than the broader semiconductor market. For all three, this revenue growth can be leveraged, i.e., they can expand margins over time, and it is also durable, because the growth is backed by a significant structural change in the world that will take years to play out. Importantly, both Nvidia and TSMC are founder-led businesses, which in our view is important, because it focuses the company on the long-term opportunity, as opposed to any short-term performance metric. Finally, all three businesses have a very high market share in their respective areas. This is symptomatic of businesses that have excellent customer perception, which we believe is essential for any growth business.

Powering AI

In addition to artificial intelligence, we believe one of the biggest investment opportunities over the next several decades is decarbonisation and addressing climate change. The need to provide carbon-free power to businesses and homes all around the world presents a significant challenge, one which creates opportunities to invest in those companies enabling decarbonisation. However, the rapid rise of AI is placing extra pressure on energy use, compounding this already very large problem of climate change. Running large-scale AI models consumes a tremendous amount of power. For example, a 1 million GPU cluster could require 1.8 gigawatts of power to operate. Even a single ChatGPT query uses approximately 10 times more electricity than a traditional Google search. This places the critical piece of AI infrastructure, the data centre that houses the AI chips, as an incredibly power-intensive asset. As more companies scale up AI, the demand for data centre power is skyrocketing, creating real and increased pressure on energy grids globally. We expect these power pressures to continue as businesses globally adopt AI.


Back in 2022, before the current AI boom, data centres used around 2.5% of US electricity. It's estimated that this number could triple by the end of the decade. But instead of slowing down clean energy progress, AI is instead speeding it up. Tech giants like Microsoft and Amazon are doubling down on renewables and signing innovative nuclear power deals to meet their growing energy needs. In early 2024, we saw an important announcement for both the growth of AI and the energy demand it requires, when Amazon partnered with nuclear energy company, Talen Energy. The companies announced at the time that Amazon would be taking a portion of nuclear power directly from the Susquehanna plant in Pennsylvania into a data centre directly adjacent, designed specifically to be powered by the nearby power plant. The announcement was for 960 megawatts (MW) of power, which is an amount of power not easily attained through regular commercial means. This deal signalled both AI's thirst for power and the critical role that nuclear could play, given it is a form of baseload power that is carbon-free. Similarly, in September of 2024, another transformative nuclear power deal was announced, this time between Microsoft and Constellation Energy, the largest nuclear plant operator in the US. This deal was for over 800 MW of nuclear power, but instead of using an existing reactor, in this deal Constellation would restart an old nuclear power plant on Three Mile Island in the US, Microsoft would spend well in excess of US\$1bn to do so, and as a result Microsoft would pay a premium for that power over a 20 year period. This deal reinforced Microsoft's quest for not only computing power from companies such

as Nvidia, but also for reliable physical power, this time through nuclear generation. Lastly, in 2025, Constellation announced another 20-year deal, this time with Meta, to supply Meta with the power needed to run its compute requirements for AI, also at a considerable premium to power prices, highlighting the value of nuclear energy's carbon-free, baseload attributes.


US POWER DEMAND FOR THE FIRST TIME IN 20 YEARS IS BEING DRIVEN BY DATA CENTRES

MUNRO


Chat GPT vs. Google search query (Wh)



US data centre power consumption (TWh)



US total power consumption (TWh)




Source: Munro Partners estimates and industry research as at 30 September 2024. The companies mentioned may or may not be held in the Munro funds.

This situation has created what some are calling the “AI paradox”, that is AI is driving up energy use, but it is also helping manage and reduce it. AI is already being used to forecast electricity demand more accurately, boost battery performance, and improve the efficiency of buildings and industrial systems. So, while AI increases the need for power, it is also pushing companies to use energy more intelligently. As a result of this, we have chosen to invest in several of the companies that provide solutions to this ‘power AI’ requirement. Our focus is on the companies that can potentially provide a solution to this paradox. We have identified the data centre as the critical piece of infrastructure for AI, and have therefore invested in several companies that play a role to help reduce this demand on physical power, whether it be companies that provide liquid cooling solutions for semiconductors, mechanical and electrical work for the data centre, component infrastructure or the physical power such as nuclear energy. As AI continues to grow, we expect the companies that provide clean power solutions to continue to grow their earnings.

COMPANIES ENABLING ENERGY EFFICIENT DATA CENTRES

MUNRO



Source: Schneider Electric, Nvidia, Vertiv, Constellation Energy, Comfort Systems, Industry Research, World Economic Forum, slide prepared December 2023. Companies mentioned may or may not be held by Munro funds and are used for illustrative purposes.

AI Applications – the ‘Total Addressable Market’

In the years to come, we believe more investment opportunities will become apparent at the application layer of the AI stack. Apps like ChatGPT and Perplexity are just the beginning, but we believe there's a lot more coming, which is exciting both from a productivity and investment point of view. What's particularly exciting about AI is that it differs from past tech revolutions like mobile or cloud because it can be applied across almost every industry. Instead of writing code, people can now teach machines using language, text, or video, and AI can learn as they go.

It can already be seen that there are digital agents and chatbots that outperform humans in some tasks. Physical robots are an example of what is possible, and training them may be as simple as showing them a video. Autonomous vehicles are becoming safer and more commonplace on roads. We believe AI has enormous potential in the medical industry, whether that be through drug discovery or in more targeted and effective medical imaging. In our view, it is impossible to put a number on what the total addressable market is for AI or to quantitatively define the return on investment. We believe applications such as these will increasingly show productivity gains, efficiency improvements, cost or time savings, and should be measured as such. One example of a live AI productivity improvement is through a company called Axon Enterprise. This company produces the bodycam that police officers wear during their shifts. The company has developed AI capabilities for this bodycam, transforming it into a virtual assistant for the officer while they are on the job. Harnessing the data this bodycam captures allows for a police report to be automatically generated through integrating AI into the bodycam software. Instead of spending 1-2 hours per day writing police reports, as the average US officer does, AI can generate that police report directly from the bodycam. This is one small example where AI is already creating a productivity improvement, and we expect many more examples like this one to emerge over time. Similarly, Spotify is testing AI-curated content for their platform. The company are not doing this to simply charge more for this content, rather they are doing this to create better engagement with the platform among their customer base, and better engagement should lead to improved monetisation over time.

For companies that adopt this technology early, the benefits could be huge - cost savings, better customer experiences, and stronger performance. And as leaders push ahead, their competitors will inevitably follow. This could lead to a major boost in productivity across the economy and strong earnings growth for the companies leading the charge.

IMPORTANT INFORMATION: *The information contained in this document reflects, as of the date of publication, the views of Munro Partners. There can be no guarantee that any projection, forecast or opinion in these materials will be realised. The views expressed in this document may change at any time subsequent to the date of issue. The material contained in this publication is for general information purposes only as is not investment advice of any nature. Past performance is not a reliable indicator of future performance. This information has been prepared without taking account of the objectives, financial situation or needs of individuals. While every effort has been made to ensure the sources of this information are reliable, no representation or warranty is made concerning the accuracy of any data contained in this document. This document is issued on 16 July 2025.*