Montaka Global Investments

WHITEPAPER

THE AMAZING HYPERSCALERS: Why epoch-defining AI is today's most important investment theme

And why the emerging 'AI flywheel' of the hyperscalers – Amazon, Microsoft and Alphabet (Google) – mean they are the surest way for investors to profit

Most investors are aware that artificial intelligence, or AI, is an important transformation. But many still underestimate the sheer world and life-changing power of this revolution and its investment potential.

AI will recast and refashion every aspect of life: how we consume, how we work, how we monitor industrial machines, how we deliver healthcare, how wars are won, and how we solve climate change.

> *"AI is the new electricity. Just as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don't think AI will transform in the next several years.*
> – Andrew Ng, co-founder of Google Brain

Not only is software eating the world, but it is becoming a lot smarter at an accelerating rate and will continue to do so. Why? Because of AI – and more specifically, machine learning (ML), the most common and practically applicable subset of AI today.

Over recent years, the combination of lower cost and scalable compute, with rapidly expanding data generation and capture, has allowed data scientists to successfully train specific ML models that are highly effective in making accurate predictions and interpretations. Think of Google's predictive text in Search or Gmail, or its facial recognition in Photos, for example. Or Amazon's voice recognition built into 'Alexa', Tesla's autonomous driving, or OpenAI's GPT-3 natural language model which is generating 4.5 billion words per day on average in human-like sentences and paragraphs.

Artificial Intelligence (AI) - Computer systems able to perform tasks that normally require human intelligence

Machine Learning (ML) - Computer systems that automatically learn and improve from input features without being explicitly programmed

Deep Learning - ML systems that use neural networks to identify complex patterns in data and automatically extract features

Source: Pradeep Menon, Data & AI Strategist, Microsoft (2018)

There are myriad ways investors can profit from this AI transformation, but Montaka believes the world's leading cloud computing providers, or 'hyperscalers', particularly Amazon, Microsoft and Alphabet (owner of Google) are one of the safest and surest ways to win.

Many investors perceive these mega-tech businesses as 'well-understood', 'mature' and sometimes even 'boring'. We disagree. We believe AI will spark a new phase of hypergrowth for these companies that will take investors by surprise and trigger a big re-rating of their shares.

The hyperscalers are uniquely positioned to create an 'AI fly wheel' – a virtuous cycle where they use their existing scale, technological advantages, huge R&D spending, and barriers-to-entry to both drive and dominate the surging demand for compute and storage that AI creates.

MONTAKA'S MISSION

Over the next decade, Montaka aims to generate superior compounding of investor capital. Our formula for success is to employ deep, multidisciplinary expertise to identify large, reliable structural transformations with favourable economics. We then seek to own the businesses which are high-probability long-term winners - if and when we assess they remain substantially undervalued by equity markets.

In addition to the hyperscalers described in this report, Montaka has identified a number of smaller companies leveraging the power of AI which are also high-probability long-term winners. The reward potential for these investment opportunities is extremely large but comes with less certainty, relative to the hyperscalers.

Montaka also draws on its expertise in other dominant and accelerating trends for the next decade and beyond, including enterprise digital transformation and the shifts in digital consumer behaviours, in assessing opportunities in the hyperscalers and elsewhere.

The hyperscalers are uniquely positioned to create an 'AI fly wheel'

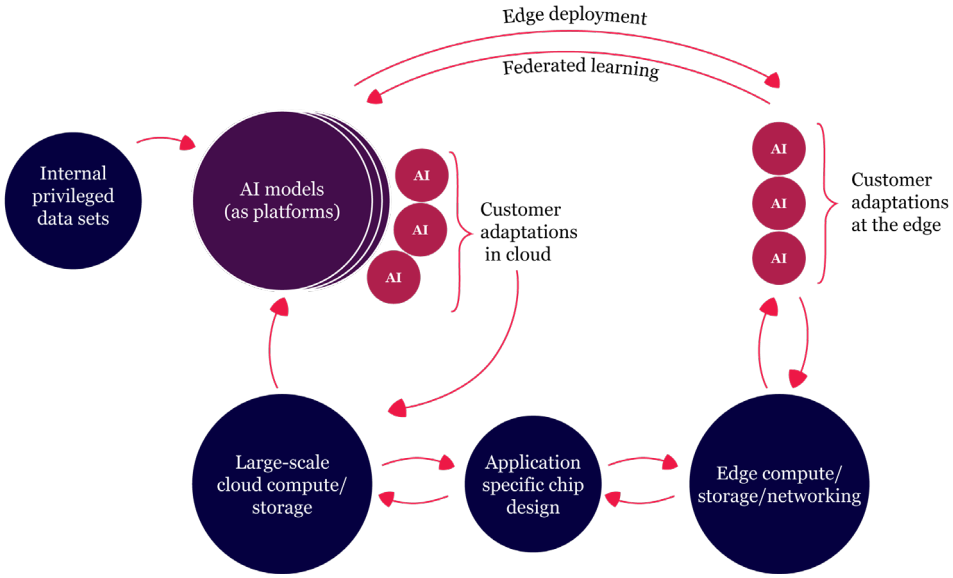

Source: Montaka

In this report we outline the 6 forces underpinning the AI revolution, how these benefit the hyperscalers, and why the market is significantly undervaluing the potential of AI to spark a new growth phase for Amazon, Microsoft and Alphabet, which creates a major opportunity for investors to safely profit from AI.

1. HYPERSCALERS ARE DEMOCRATISING AI

In 1908, Henry Ford launched the Model T as the first relatively cheap, easy-to-drive automobile. This was a highly effective strategy because it put the power of the automobile in the hands of the masses: by the 1920s, a majority of Americans had learned to drive the Model T.

Similarly, the world's hyperscalers are investing heavily in low code/no code interfaces (including, for example, drag and drop interfaces and natural-language-to-code translators) to democratise the development of AI-based enterprise applications so they are used by more and more employees of the cloud providers' customers. The more employees who can build, train and deploy AI-based applications, the more value AI can deliver these organisations (and the greater the demand for hyperscaler compute and storage services).

Imagine an alternative history where Ford owned all the world's oil refineries in 1908. Might Ford have adopted a different business model? Might Ford have given away a Model T to every driver to spark explosive demand for gasoline? This is somewhat analogous to what is taking place in global AI today.

Today's hyperscalers are strategically investing so their global customer bases accelerate the deployment of AI-based applications. And they are intelligently choosing to monetise AI in a way that continually strengthens their advantages over competitors and increases barriers to entry for others trying to compete against them.

The world's hyperscalers have developed – and continue to do so at remarkable pace – highly-complex ML models tailored for specific scenario-based business services, including personalisation, fraud detection, cognitive search, intelligent document processing, media intelligence and customer intelligence. They have important natural advantages in training these models because they have privileged internal datasets – that is, large and relevant datasets their core businesses have grown over the last two decades and that are near-impossible to recreate. They also have access to low-cost, large-scale internal compute, and access to the world's best engineering talent.

Most of these ML models are essentially 'given away' free. (And remember, unlike a Model T, the cost of incremental production and distribution of software is zero). Any customer can essentially take an ML model 'off the shelf' and relatively easily and cheaply adapt and customise it to their own needs using their own internal datasets. The quid pro quo? The customer must use the hyperscaler's compute and storage services.

> *"Large-scale AI models are becoming platforms in their own right.*
> – Satya Nadella, Microsoft Ignite 2021

Just as web development was in the late 1990s, ML is now a tool at the disposal of every software developer. In the future, no-code/low-code interfaces (which are currently being rolled out by hyperscalers) will enable non-developer employees to harness the power of AI. More than 500 million new applications are projected to be developed by 2023[1] – the equivalent to the total number of applications developed in the last 40 years.

ML is being incorporated into applications in just about every industry – not just the obvious technology industries. From banks, to insurers, to healthcare and pharmaceutical providers, to retailers and even asset managers, ML is being infused in a wide range of applications to improve customer value propositions as well as internal productivity. And the opportunities to leverage data in the enterprise are enormous and growing. Consider that only one-third of enterprise data is put to work today – and the volume of enterprise data itself is doubling every two years.[2]

[1] IDC FutureScape: Worldwide IT Industry 2020 Predictions, October 2019
[2] Seagate Technology Report, 2020

The opportunity that lies ahead for today's leading hyperscalers is nothing short of astonishing. Why couldn't every expense line-item on the income statement of every corporate, government and household be touched by AI in some way? Of course, all of these expense line items aggregate to nearly US$100 trillion per annum in global GDP, so the opportunity, through this lens, is staggeringly large.

Today, we have met only the tip of the AI iceberg. And we believe this long-term opportunity is currently being drastically underappreciated by equity markets.

AI - A QUICK PRIMER

"A computer would deserve to be called intelligent if it could deceive a human into believing that it was human."
- Alan Turing

Mathematicians and computer scientists have been in pursuit of the creation of a thinking machine, or Artificial Intelligence, since the 1950s. But it was a quantum leap in an area called Deep Learning in the 2010s which really put AI on the map. Deep Learning is based on training a multi-layer neural network which, while built in software, mimics the network structure of the neurons in a human brain. As a result of this technique, all of a sudden, machines were able to recognize objects in pictures and videos. Today, this form of AI underpins natural language processing (including Alexa, Siri and Google Translate), recommendation algorithms (including those found within Facebook, Instagram, Spotify and Google Search), medical diagnoses and predictions, warehousing and logistics, call centre assistances and autonomous vehicles.

Today, AI models are typically trained in one of three ways: (i) supervised learning; (ii) unsupervised learning; and (iii) reinforcement learning.

Supervised learning employs labelled datasets and are used to help the AI associate the dataset features with its appropriate label. In unsupervised learning, the training data contains only inputs and the AI simply seeks to identify patterns - in many case patterns which humans have missed. In reinforcement learning, the AI is an 'agent' in a controlled environment and can perceive and interpret its environment and take actions. By rewarding 'desirable' actions and punishing 'undesirable' actions, the AI learns through trial and error.

Solving the problem of artificial 'general' intelligence (AGI) - that is, software capable of human-level performance of any intellectual task and capable of relating tasks and concepts to others across disciplines - will ultimately require machines to train themselves in complex unsupervised ways, which, practically, remains a long way off.

"For general AI, almost all learning will have to be unsupervised, but no one has yet come up with the kinds of algorithms needed to perform successful unsupervised learning."
- Melanie Mitchell, Davis Professor of Complexity at the Santa Fe Institute

2. THE INTERNET OF THINGS (IOT)/'EDGE' REACHES A TIPPING POINT

The second driver of massive demand for compute and storage services delivered by hyperscalers is that AI is increasingly being incorporated into every software application in every device. This includes IoT for which we are about to see an explosion in smart devices, deployed at the 'edge', built to be used across all aspects of life.

We are now at a tipping point where the number of connected, intelligent devices and sensors will explode. This is known as the 'Internet of Things' or IoT. These devices will not only continually measure, capture and communicate data, they will also be armed with ML models that continually relearn and improve over time. Industrial sectors, in particular, will drive the lion's share of the growth, as illustrated by the chart of expected future data generation below. No longer can humans keep track of today's increasingly complex industrial processes – but AI can.

Data Generation - By Category

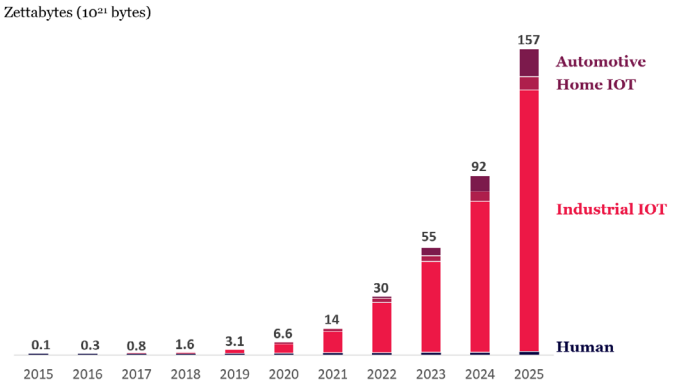

Source: Applied Materials (May 2021)

An important characteristic of IoT is that these devices physically reside near the 'edge' of the network where data is created, not in some centralised datacentre. Edge devices include, for example, industrial plant sensors, autonomous vehicle cameras, and even our smartphones.

This has important practical implications for ML-based applications. Given the extreme explosion in data generation, and bandwidth and latency constraints, it is practically unfeasible to continually transmit this data back to a centralised datacentre for interpretation by the ML model (known as 'inferencing') and secondary model training.

Instead, ML models – initially trained in the centralised 'cloud'— will be moved to edge devices, or an edge server, for localised application and ongoing retraining of the localised ML model[3]. This explains why the world's leading hyperscalers are rapidly building out their edge infrastructure (largely through partnerships with mobile networks) as well as mission-critical edge services (e.g. edge security; edge-friendly model design/ compression; federated learning), many of which are not yet being monetised today.

[3] Interestingly, the original "mother model" in the cloud still benefits from this localised secondary training via a bandwidth-friendly solve called "federated learning" in which only the required model augmentations are transmitted back to the centralised cloud

The impending explosion of edge inferencing will be unprecedented. Each time an application seeks an 'interpretation' from the deployed ML model, it 'inferences' the model, that is, the application presents some new data to the pretrained ML model and asks it to make a prediction. Consider that Amazon expects the future split in infrastructure costs between ML model training and inferencing to be 10:90 – despite the compute intensity of ML model training to be many multiples that of model inferencing. The sheer volume of edge devices, as alluded to by the chart above, will ensure that AI on the edge is going to be a much bigger business than AI in the cloud.

3. ML DRIVES SURGING DEMAND FOR COMPUTE AND STORAGE

The effect of incorporating ML into applications is a significant increase in compute and storage intensity, which will benefit hyperscalers.

Developing an ML model is a particularly compute-intensive process. Very large quantities of raw data are structured and prepared for model 'training'. During the training process, billions of model parameters are continually updated and tested until such time as these parameters can make predictions above some threshold accuracy.

The compute-intensity of ML model development has increased sharply over the last decade, an increase that reflects the growing complexity of the underlying models (with today's largest models surpassing one trillion parameters).

As you can see in the graph below, compute requirements now double every 3-4 months. In the 50 years prior, this doubling was taking place every two years on average.

Once the ML model has been trained for long enough (usually many weeks or months), a copy of the model can be deployed – often to an edge device – as part of a broader software application designed to add personalised, predictive value for customers, employees or other stakeholders. Each time the application seeks a 'prediction' from the deployed ML model, it 'inferences' the model, that is, the application presents some new data to the pretrained ML model and asks it to make a prediction. Inferencing is also compute-intensive, though materially less than for model-training.

Importantly, the most successful AI applications are those which integrate continual secondary data-collection to enable ML model 'relearning' into their use-cases. In a consumer-facing application, for example, the more customers that engage with the application more frequently, the greater the volume of data that is captured and the more accurate the embedded ML models with the application become. This is the 'AI fly wheel' that strengthens existing business advantages over time.

Naturally, the more successful an AI application is at continually extracting additional relevant data to virtuously improve the accuracy of the embedded ML models, the vastly more compute and storage that is required. This is great news for the global hyperscaler oligopolies that supply the world's compute and storage, in particular Amazon, Microsoft and Alphabet (Google). Itis similarly positive for China's cloud oligopoly, Alibaba, Tencent and Huawei, though the Chinese addressable cloud market is considerably smaller and some years behind in development.

Compute Usage in Training AI Systems

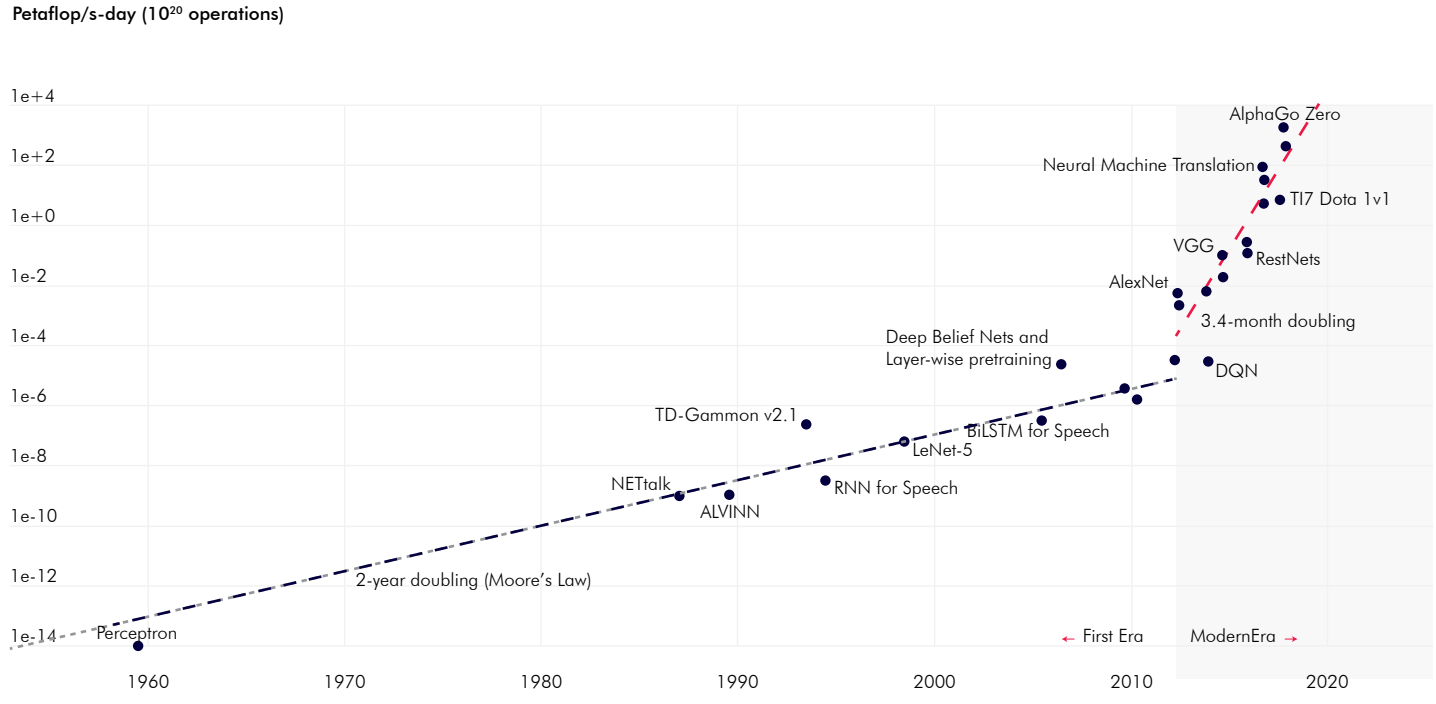

Source: OpenAI (2019)

4. HYPERSCALERS SET TO DOMINATE 'ASICS' – THE NEW WAVE OF CHIPS POWERING AI

Another reason that hyperscalers are positioned to win from AI is that they are at the cutting edge of developing new chips – Application Specific Integrated Circuits (ASICs) – that will dominate in coming years.

AI applications are simply becoming so large and complex that traditional chips are increasingly too slow, energy intensive and expensive for these purposes. So increasingly compute is becoming application-specific – both in the cloud and at the edge. This form of compute is being delivered by ASICs, which are designed on a more bespoke basis so they are more economical for the specific nature of the task to be undertaken.

And with the ongoing acceleration in model complexity and associated compute intensity, the imperative to continually improve ASIC performance is only growing stronger.

ASICs outperform other chips because they are designed to perform very specific computational processes required by the underlying application. Interestingly, Bitcoin mining is among the most mature of these applications.

It is not surprising that the hyperscalers are all pursuing the design of new ASICs tailored for specific ML-based applications. Amazon Web Services (AWS) commenced in 2015 after it acquired Israeli chip designer, Annapurna Labs. Since then, AWS has demonstrated a 40% reduction in the cost of running its online photo sharing service by shifting to its in-house chip, called Graviton.[4] Already, AWS is up to the third generation of its Graviton chip designed specifically to train ML models – at a cost that is 40% lower than leading third-party alternatives (such as Nvidia's flagship chip).[5]

Similarly, Alphabet launched its Tensor Processing Chip (TPU) in 2016, designed specifically for superior ML model training and inferencing, which powers several of Google's major products including Translate, Photos, Search Assistant and Gmail. It was also included in the Pixel 6 smartphone for the first time in recent months. Others are following down a similar road including Microsoft, (Meta Platforms for its internal ML models), Alibaba and Tencent.

The logic for this vertical integration is compelling. Who better to more efficiently uncover optimal chip design specifications than the designers of the underlying ML models these ASICs will be training and inferencing? And with experts that Montaka interviewed expecting that ASIC designs will need to be reoptimized every nine months, tight and timely integration between model designers and chip designers is imperative.

Furthermore, designing in-house means that hyperscale customers no longer need to pay for the significant gross margins achieved by existing designers, such as Nvidia (around 64%) and AMD (around 47%) – instead they buy directly from the semiconductor fabrication plant. Of course, there will be some incremental R&D expense associated with design, but the sheer scale of the hyperscaler R&D budgets can easily absorb this. Consider for context that, in just the next eight months, Alphabet will spend more in R&D than Nvidia has spent over the last decade.

Finally, more efficient compute for ML model training and inferencing can be sold at much higher prices than traditional forms of compute (CPU or GPU, for example) and still deliver superior value for customers. Consider that the employment of Alphabet's TPUs for ML model training is 40 percent cheaper relative to using GPUs, despite being priced at approximately 5x the level of GPU compute.[6] This is due to the TPU's superiority in speed and energy efficiency.

5. HYPERSCALERS LEAD RACE TO DEVELOP ENERGY EFFICIENT COMPUTE, AN IMPORTANT DRIVER OF OUR PLANET OVERCOMING CLIMATE CHANGE CHALLENGES

Successful innovation on chip energy efficiency is an imperative for the long-term proliferation of AI. Given the enormous economic incentives at play, there is a high-probability the hyperscalers will lead this innovation. Not only will this have a positive direct impact on the world's physical environment, but the enabling of increasingly powerful AI will likely, itself, resolve many of the challenges the world needs to overcome to mitigate climate change risks.

Today's chips will not be able to handle the compute requirements of the future. A major reason for this is energy efficiency (and associated cost). Researchers, for example, have estimated that extrapolating the current rate of improvement in ML model error rate reduction for, say, image recognition would require as much as $100 billion for model training (as soon as 2025) and generate as much carbon emissions as New York City does in a month.[7]

Similarly, analysts have noted that, while only around 1% of the world's electricity goes to cloud computing today, by the end of the decade this will be north of 8%[8]. And even this is likely to be a conservative projection, given the impending explosion of IoT.

It does not take much to realise that either: (i) the energy efficiency of compute needs to drastically improve; or (ii) the progress of AI will be drastically constrained.

In addition to rapidly accelerating in-house efficient ASIC design and optimisation, hyperscalers are also exploring alternative forms of compute, such as photonic processors [9], which use light frequency and intensity to perform calculations – not transistors – and are substantially more energy efficient as a result.

The imperative to tailor ASICs to unlock the potential of AI is very strong – both for today's hyperscalers, given the enormous commercial opportunities that lie ahead; and today's nation states – given the geopolitical power that will stem from AI superiority.

Montaka expects today's hyperscalers to continue to invest heavily in ASIC design at a scale that will not be matched even by today's leading chip designers.

Given the strength of this imperative, we believe there is a high probability that the energy constraints on AI will be resolved over time. This is good news for investors in AI, and even better news for the world's climate.

[4] (Bloomberg) Why Amazon, Google, and Microsoft Are Designing Their Own Chips, March 2021
[5] (Reuters) Amazon's cloud unit launches new chips to take on Intel, Nvidia, November 2021
[6] Stone, Now you can train TensorFlow machine learning models faster and at lower cost on Cloud TPU Pods, 2018
[7] (IEEE Spectrum) Deep Learning's Diminishing Returns, September 2021
[8] (IEEE Spectrum) Cloud Computing's Coming Energy Crisis, July 2021
[9] For example: Lightmatter, a private company owned in part by GV, Alphabet's venture arm

Not only will ASIC improvements reduce the energy intensity of compute – clearly a first-order positive for the world's climate – but the enabling of increasingly powerful AI will likely, itself, resolve many of the challenges the world needs to overcome to mitigate climate change risks. That will include climate risk detection models to grid scheduling algorithms, new fuel-material discovery, waste-reduction algorithms, supply chain optimisations, shared mobility, precision agriculture, and infrastructure design.[10]

> *"I think machine learning so far has realised only a small fraction of its potential. For machine learning to reach its full potential, it's going to be a multi-decade process and over time, it will make its way into every problem.*[11]
> - Francois Chollet, Google AI researcher

And this is not simply conjecture. In 2016, Alphabet's DeepMind employed ML to reduce Google data centre cooling expenses by 40 percent.[12] Through this lens, the valuable investment opportunity we see in today's hyperscalers is also one of the most environmentally supportive investments we can think of.

6. HYPERSCALERS WILL BENEFIT FROM GROWING BARRIERS-TO-ENTRY THAT MAKES THEM LONG-TERM WINNERS FROM AI

Today's hyperscalers have a significant lead over competitors and it is highly likely this lead will only extend over time. The 'barriers to entry' in the space are already very high – and rapidly growing higher, all but eliminating the realistic prospects of a major new competitor materialising.

Here are some of the significant contributors to these growing barriers to entry.

First is the **rapidly strengthening 'customer captivity' enjoyed by today's hyperscalers.** Most of the world's large corporates and governments already rely on these businesses for their cloud infrastructure and many of their mission-critical IT platform and enterprise application services. But as powerful ML models are increasingly built into services, these services become 'must haves' for customers, driving greater adoption.

The more adoption, the more relevant data that is generated to retrain and continually enhance the ML models, making them even more essential to customers – and so on as the 'AI fly wheel' gathers steam. It is exceedingly likely that early winners in AI will keep on winning. And the early winners in AI today unquestionably include the hyperscalers: Amazon, Alphabet, Microsoft, Alibaba and Tencent.

[10] Climatechange.ai
[11] Google Cloud Applied ML Summit 2021
[12] (DeepMind) DeepMind AI Reduces Google Data Centre Cooling Bill by 40%, July 2016

2020 Market Share

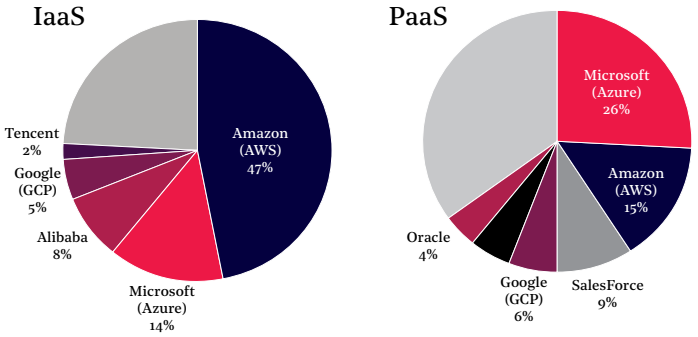

Bernstein, Montaka Global

Second, **the scale advantages enjoyed by today's hyperscalers are much more significant than many investors appreciate and will likely continue to strengthen over time.** These represent enormous barriers to entry by enabling investments in capex and R&D (such as the expensive training of powerful ML models, for example) that cannot be rivalled by competitors. It is these investments that have created the early lead in AI for today's hyperscalers. And this lead is growing rapidly thanks to new investments in the space each year.

It is well known that Alphabet, Amazon and Microsoft, for example, outspend every other technology business each year in R&D, representing a large-scale advantage and barrier for would-be competitors. But Montaka believes this cursory comparison significantly understates this relative scale advantage.

Consider the extraordinary fact that around 70 percent of hyperscaler cloud compute capacity is used for internal purposes. Capacity utilised internally is expensed at cost, while the true 'market value' of this compute and storage is around 3x higher. Upon adjusting this component of R&D expenditure, we illustrate how today's hyperscalers – particularly the American leaders – are actually much further ahead of the competition in R&D than many likely appreciate. (Meta Platforms, given its enormous internal cloud infrastructure and AI capabilities is included in this category). This is a near-insurmountable barrier for potential competitors to overcome.

Comparison of R&D + Capex

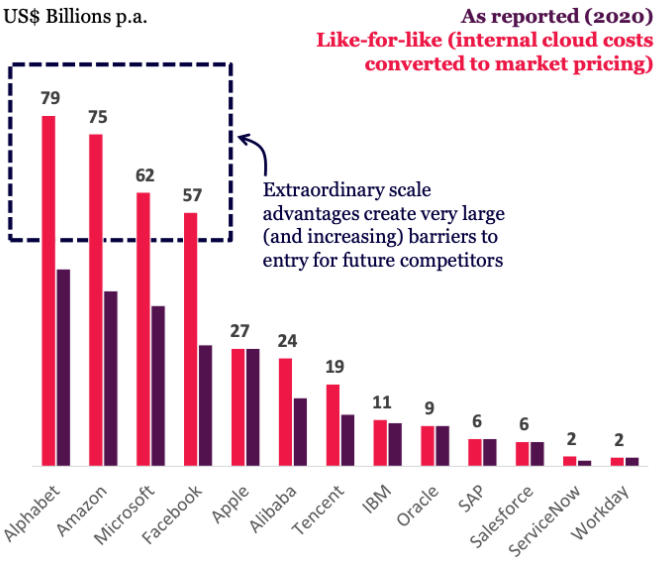

Source: Company Filings, Montaka Estimates

Now, what on earth could so much money be spent on? Increasingly, the answer is: ML models (and the optimised chips required for efficient training and inferencing).

> *We have made progress in understanding different modalities, be it text, images, voice, vision, etcetera...I think we are at an inflection point, and we are investing to build better models and deepen our understanding and do it in a more generalized way and when we do that, it will apply horizontally across our products.*[13]
> - Sundar Pichai, Alphabet CEO, 2021

Today's leading cloud providers are rapidly building out their 'ML as a platform' offerings. By investing tens of billions of dollars each year to internally develop and offer what would otherwise be prohibitively expensive scenario-based ML models to developers (mostly for free!); and by creating the low code/no code interfaces to democratise the development of AI-based application across the much larger base of non-software-developer employees, the cloud providers are drastically accelerating demand for their compute and storage.

Through this lens, today's hyperscalers can in effect 'amortise' these very high upfront ML model development costs across their very large customer bases – while also driving greater adoption. In doing so, they are extending their already-strong advantages in scale, customer captivity, and data/ML, thereby increasing further the enormous barriers to entry to this highly-attractive global market.

AI WILL CREATE A CONCENTRATED HANDFUL OF WINNERS

Even with substantial and sustained government intervention, most-countries - even technologically advanced ones - will not give rise to companies that produce or maintain an advanced "national" version of each globally influential network platform (such as those used for social media, web search, and so on). The pace of technological changes is too rapid, and the number of knowledgeable programmers, engineers, and product design and development professionals too few, for such broad coverage. The global demand for talent is too high, the local markets for most services too small, and the product and service costs too substantial to maintain an independent version of each network platform.

To stay at the evolving forefront of technological development requires intellectual and financial capital beyond what most companies possess - and beyond what most governments are willing or able to provide. But even in such a scenario, many users, if given the choice, would rather not be limited to a network platform that hosts only their compatriots and the software offerings and content they produce. Instead, the **dynamics of positive network effects will tend to support only a handful of participants who are leading the technology and the market** for their particular product or service.

- The Age of AI (2021), by Kissinger, Schmidt, & Huttenlocher

OWN TODAY'S HYPERSCALERS TO WIN IN AI

Montaka believes AI will drive much, much more demand for compute and storage – both in the cloud and at the edge – than many are expecting today. Furthermore, this growth is largely assured, the long-term winners are already known today with a high degree of certainty, and we believe the current stock prices of these businesses are failing to adequately reflect what Montaka's sees on the horizon.

There are a number of supportive reasons for this view.

First, consider that the cloud businesses of the top three American hyperscalers (Amazon, Microsoft and Alphabet) aggregated to 0.10% of world GDP (ex-China) in 2020. Current stock price expectations see this ratio increasing to only 0.55% by 2030. Straight away, this appears conservative for an epoch-defining technology, housed on the platforms owned by these hyperscalers, with applications across every sector and in every region.

For another reference point, Microsoft's Satya Nadella has forecast global IT spend to increase from approximately 5% in 2020 to 10% in 2030. This implies around US$6 trillion in incremental annual IT spend by 2030. By 2030, our three American hyperscalers are expected by the market to have added less than U$0.5 trillion in annual aggregate revenues. Again, this sounds like an extreme mismatch and points to the possibility of hyperscaler revenue forecasts to be unreasonably conservative.

Thirdly, one can compare the expected 2030 aggregate American hyperscaler revenues of approximately US$600B to the consensus cloud total addressable market (TAM) of approximately US$1.2 trillion. At first glance, this ~50% TAM penetration might sound reasonable based on current market positioning. But this 'TAM' has been constructed based on the opportunity to move current on-premise IT workloads to the cloud. It includes little, if any, for the new workloads that will be created by AI applications – particularly at the edge.

> *It's the need for distributed cloud infrastructure. It's both needed for modernizing existing applications you have (and so that's, by the way, 20% penetrated. So, there is 80% more that needs to move). But more importantly, there is going to be new application starts, which need infrastructure. And so, if you sort of add those up, I think that we are still in early innings...*[14]
> - Microsoft CEO, Satya Nadella, 2021

It appears highly plausible that hyperscaler revenue expectations are far too low in the context of the scale of the AI-based opportunity that lies ahead. If so, then Amazon, Microsoft and Alphabet – as well as Alibaba and Tencent – will likely surprise investors substantially to the upside over the coming years.

Finally, investors should also remember that the enormous R&D being incurred by the hyperscalers, while expensed fully each period to satisfy accounting rules, represents an economic investment in future earnings power. Through this lens, hyperscaler earnings power today is 'artificially' understated – and valuation multiples, therefore, overstated.

[13] Alphabet Q4 2020 conference call

[14] Microsoft Q3 2020 conference call

History has shown that prior long-term winners that leverage their scale to continually grow and improve their customer value proposition have remained undervalued by equity markets for decades. Investors such as Nicholas Sleep[15], for example, pointed this out for businesses like Wal-Mart and Costco. And these were largely physical rollouts.

The current AI 'rollout' is, of course, largely a software story. And that means the economics of this transformation are generally far superior to the expansion of physical retail stores and logistics infrastructure. One would hardly be surprised to look back at this moment from the year 2030 or 2040 and remark how undervalued the world's leading hyperscalers were in 2022.

The conclusion is clear: to win from the AI revolution, own today's hyperscalers, Amazon, Microsoft and Alphabet (Google).

[15] Nomad Investment Partnership Letters

NOT ALL EXPENSES ARE CREATED EQUAL - WHY R&D MASKS EARNINGS & VALUE IN TODAY'S HYPERSCALERS

When a business reports its operating profits, it wants to give investors a picture of the revenues generated, less the costs incurred to generate those specific revenues. For most expenses, such as 'cost of goods sold' (that is, the direct costs incurred to create a product or service which has been sold) there is a clear link between the cost and revenue.

But for R&D, the link between cost and revenue is less clear.

Under US accounting rules, if a company spends money on R&D the costs are immediately and fully 'expensed' (that is, today's profits are reduced by that full amount) because it is uncertain whether the company will derive any future economic benefits from said R&D.

This convention is very conservative and the implicit meaning of R&D expense in the US is that a company will not realise any future economic benefits. But for many businesses - particularly today's hyperscalers with powerful market positions, scales, and data advantages - R&D does often lead to huge future economic benefits in the form of massive cash flows, profits and more entrenched market positions.

"We've routinely seen stronger returns on our [AI] investments over time than we've expected." - Mark Zuckerberg, Meta Platforms, 2022

So, for businesses like the world's leading hyperscalers, R&D is an 'investment', not an expense. And by implication, their earnings power is understated and, therefore, their valuation multiples overstated.

-Montaka

CONTACT DETAILS

PRIVATE CLIENTS

Giles Goodwill
Private Client Manager
E: ggoodwill@montaka.com
P: 02 7202 0100

INTERMEDIARIES

Matthew Briggs
Investment Specialist
E: mbriggs@montaka.com
P: 02 7202 0093

OFFICE

Montaka Global
www.montaka.com
E: office@montaka.com
P: 02 7202 0100

Visit montaka.com and subscribe to the Montaka Monocle - our monthly newsletter curated with our latest insights.