STRATEGY PERFORMANCE (%, NET)[1]

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Fund	Index
2015	+0.1	-0.3	+0.0	+0.7	+2.1	-1.4	+2.0	-0.2	+1.5	+1.7	-0.2	-0.9	+5.1	+9.8
2016	+0.0	+0.0	+2.9	+2.3	+11.8	-3.6	+5.3	-6.8	+0.6	+2.3	-3.5	-2.5	+7.7	+8.4
2017	+1.9	-2.6	-1.0	+3.5	+2.2	+3.5	+1.7	+3.0	+2.4	+9.5	+4.1	-1.1	+30.3	+14.8
2018	+6.5	-2.1	-0.7	+0.6	+9.1	+3.2	+5.6	+2.9	--	--	--	--	+27.4	+11.8
2019	--	--	--	--	--	--	--	--	+1.1	+0.8	+0.1	-1.6	+0.5	+6.6
2020	-1.3	-0.3	-5.2	+2.4	+7.9	+3.0	+9.5	+3.5	-1.4	-1.4	+4.8	+2.0	+25.2	+5.9
2021	+4.5	+6.9	+1.2	+3.0	-1.5	-0.7	-1.4	+2.6	+3.1	-0.6	-4.6	-0.4	+12.2	+25.8
2022	-4.1	-1.7	+1.5	+0.5	-2.0	-3.5	+1.4	+4.4	+5.8	+0.7	-0.4	-1.3	+0.8	-12.5
2023	+1.7	-1.5	+2.4	-2.7	+1.2	+0.8	+1.2						+3.0	+18.9

	Fund	Index
Since Inception	+175.1	+125.7
Geometrica p.a	+10.3	+10.5
Strategy p.a.	+14.3	+11.3

2015 – 2018: CVF (same portfolio managers and strategy)

2019 onwards: Geometrica.

Index = MSCI All Country World Index (AUD)

STRATEGY PERFORMANCE ASYMMETRY

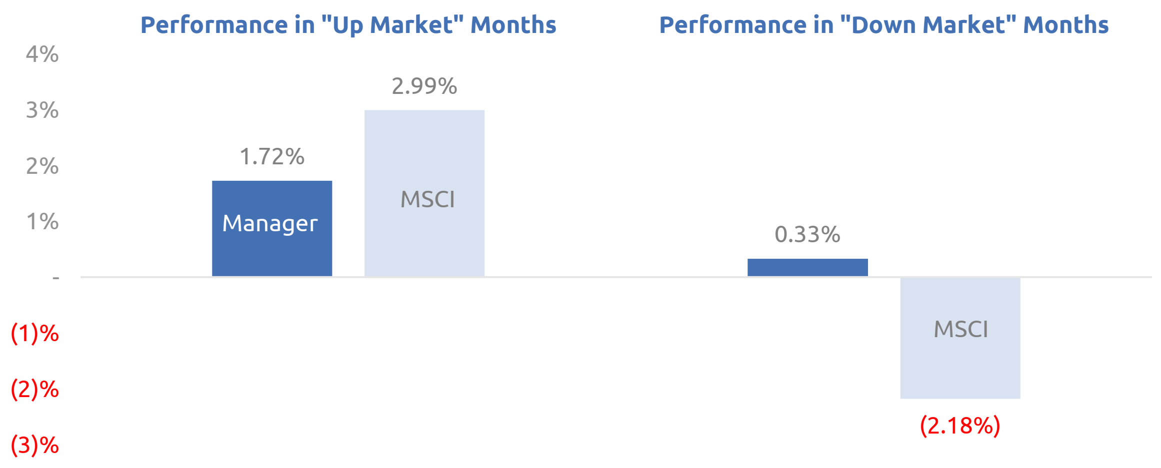

Source: Mainstream, ASX Announcements, Geometrica and Bloomberg. Performance is after all fees, from Jan 2015 (excluding the period of Sep 2018 – Aug 2019; Manager left CVF in Aug 2018 and began Geometrica in Sept 2019). MSCI = MSCI ACWI (AUD).

OVERVIEW

The Geometrica Fund returned **+1.21%** after costs in the month of July 2023.

Longs contributed +2.1% whilst shorts and currency detracted -0.2% and -0.6% respectively.

Exposures rose during the month, but we do this in step with the bottom-up process, not the top-down view. Keeping exposures low worked well in the turbulence of 2022 but it

[1] *Performance is after all fees, Founder Lead Series units.*

has been a handbrake on returns in 2023 in a market that has bounced savagely in the first half of the year.

As we find compelling opportunities, we deploy capital. For us, a compelling opportunity typically falls into either of the following two categories:

1. A high-quality business faces a short-term headwind, which is excessively discounted by the market. These tend to look cheap relative to their history, as they discount heightened investor fear.
2. A high-quality business has its forward earnings under-estimated by the market, resulting in undervaluation. These tend to look expensive on near term earnings but are cheaper than they look given under-appreciated growth.

You might say the first results from excessive extrapolation and the second from insufficient extrapolation. There are recurring behavioural biases that drive these situations.

Some examples, in the context of notable portfolio stocks this month (winners and losers):

Grifols (GRF.SM, mkt cap. EUR€8,0bn) rose +13.7% in July, contributing ~40bps to performance. Grifols collects, fractionates and sells blood plasma products such as albumin, pdFactorVIII and immunoglobulin (**IG**). Australia's CSL is the industry leader (CSL.AU, mkt cap A$127.3bn), but trades at eye watering valuations relative to Grifols.

Grifols and CSL were both adversely affected by covid, travel restrictions on the USA / Mexico border (CSL and Grifols are the #1 and #2 operators of US plasma collection centres) and government payments which restricted donation volumes – all due to covid.

We invested in Grifols because we thought the gross margin compression from covid would fully revert with time. Recently both companies indicated as much, with CSL projecting full gross margin recovery over 3-5 years.

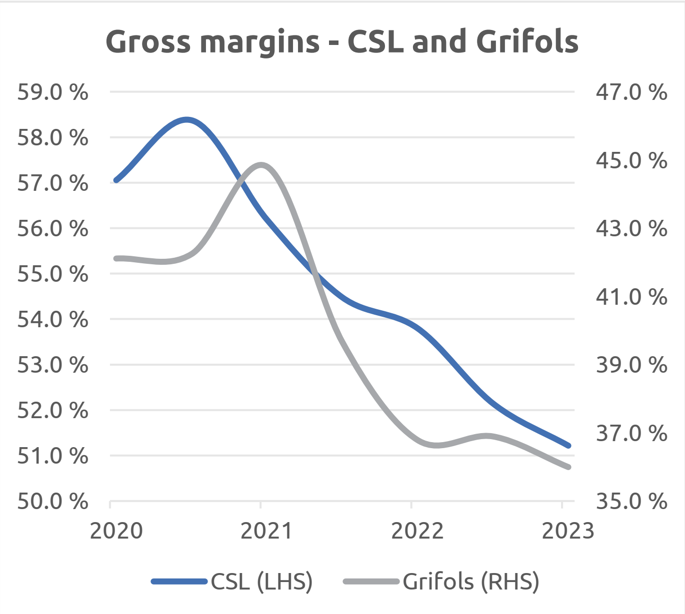

CSL have record volumes of plasma currently in inventory. Grifols is in a similar situation.

So, looking out we expect growing sales volumes, falling unit costs and modest price increases. This should act as a powerful fillip to earnings, particularly for Grifols with its higher operating leverage.

For CSL and Grifols there has also been market concern around novel therapies targeting indications or end uses already served by IG, with one such therapy reporting late-stage trial success[2]. We think these novel therapies will

[2] ADHERE trial efgartigimod targeting CIDP, an important market for Grifols and CSL.

result in increased segmentation of the market, not a winner take-all scenario. We note such concerns have dogged these stocks in the past. In the late 1990s the Cochrane report claimed albumin was responsible for more deaths than lives saved (the report was later heavily criticised). Similarly, fears over pdFactorVIII clotting factors being replaced by recombinant versions dogged the stocks early in the 2010s but the issue proved transitory from an earnings perspective.

As Grifols earnings recover, its cash generation capacity should rapidly increase, which we think will drive a positive change in the market valuation of the company.

Uber (UBER.US, mkt cap. US$91.2bn) rose 14.5% in July and contributed 45bps to performance.

> *"The probability of this business going bankrupt in a decade is 99%".* 2017

Uber has always been a polarising stock, as the quote above suggests. It is unlikely to go bankrupt; Uber is now generating cash (and not the "adjusted" kind).

Uber's historical gargantuan aggregate losses were driven by a land grab mentality enabled by near free capital.

But the kernel of truth at the heart of the model is that provided market share is high in a given location, unit economics are attractive and defensible against competition.

What has changed now is capital discipline. Just like in streaming.

As a general observation, the most aggressive of growth investors have pulled back from funding multiple competitors in the same market racing each other to grow revenue by ramping cash losses. Cost of capital matters.

Uber is now focused on generating cash and profits. As Uber has slowed early-stage investment, the profitability of the more mature businesses has come to the fore. Uber is dominant in mobility and has significant margin and profit upside from here.

Universal Music (UMG.NA, mkt cap €41.6bn) rose 14.6% during the month. Music under monetises versus video-streaming. Yet music catalogues tend to have far greater commercial longevity.

There are moves afoot to recast streaming economics to more fully reward talent and we think these are not in the price, leading to a reasonable degree of upside asymmetry in a business that should continue to grow its revenues and earnings over time. There's also an element of the inexorable here given the link between talent and streaming audience share.

First Solar (FSLR.US, mkt cap US$18.9bn) was up 9% in July, contributing 0.3% to performance. Solar is a market currently on the nose. Sector companies such as Enphase (inverters) and SunRun (system installation) have been poleaxed, due to the US housing slowdown and in part due to changes to California's electricity market. None of this affects First Solar, who operate exclusively in utility scale solar. Baby with the bathwater.

Detractors included Netflix, Renesas, Daiichi Sankyo and Pilbara Minerals. We exited Pilbara due to observing lacklustre end-demand indicators, but continue to own Netflix, Renesas and Daiichi Sankyo.

No investor letter would be complete without some mention of Artificial Intelligence or AI. Is it a bubble or is it real?

We own a modest stake in **Nvidia** (NVDA.US, mkt cap US$1.1tn). Given the controversial nature of the sector, we have put down a bit more detail than usual below why we hold it.

There's a lot of trash around the periphery of AI that's gone up a lot. Nvidia, however, doesn't strike us as trash. It may be expensive today, but that perspective on today's price could change given time.

There are very few companies in the technology world that have the potential to dominate industry defining profit pools. Apple for smartphone, Intel/Microsoft for PC.

The risk return set up here strikes us as decidedly asymmetric, especially on a multi-year timeframe. The catch is the volatility which in turn constrains sizing capacity.

NVIDIA

Background

Nvidia is the dominant player in the Graphics Processing Unit (GPU) market with exceptionally high shipment share.

These days, a GPU is somewhat similar to a Central Processing Unit (CPU) in that it can be used for general computing tasks. The difference lies in the architecture; a CPU is built for breadth and versatility while a GPU is built for throughput.

GPU vs CPU

A CPU has:

- A small number of very powerful "cores" / ALUs (arithmetic logic units) – think of a core like a brain, it is responsible for taking input and converting it to an action.
- Large inbuilt storage / memory called a "cache" which allows the CPU to store information close by that it knows it will need soon rather than sending it all the way to the computer's separate memory drive (hard drive or RAM).
- A control unit which helps the CPU interface with other devices so that it can interpret all the different kinds of inputs it receives and correctly translates / directs its outputs.

By contrast, a GPU has:

- A large number of relatively low powered cores.
- A tiered cache system that is typically smaller than what a CPU has but has higher bandwidth allowing for fast but limited information storage during calculations.
- Less "complex" control units.

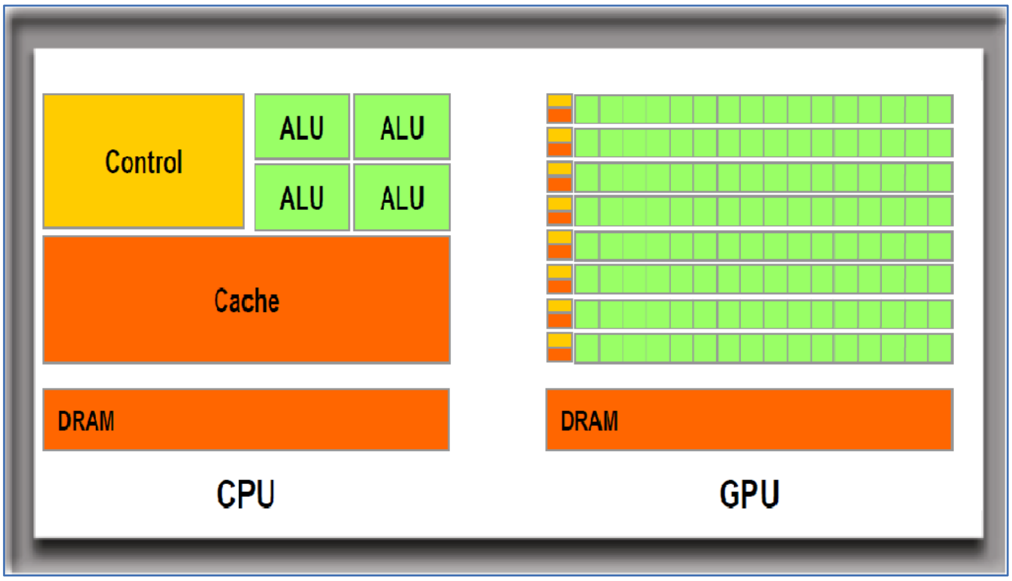

Source: Nvidia developer blog

An example of a task that a CPU would be good at is calculating a Fibonacci sequence: 1, 1, 2, 3, 5, 8, … n. This is because calculating a number in the Fibonacci sequence requires knowledge of the two previous numbers, and as those numbers become larger, they take up more memory to store. The CPU's higher on-chip memory and faster single-core speeds mean that it will outperform the GPU in in summing the previous two numbers and storing them. The GPU would be forced to use only one of its thousands of slow cores for each calculation and may face memory issues when the numbers became large enough.

Conversely, a GPU would outperform a CPU on a task such as matrix multiplication:

$$\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \times \begin{bmatrix} 5 & 6 \\ 7 & 8 \end{bmatrix}$$

This is because matrix multiplication is done number by number and the result of one calculation does not depend on the result of another. In the example above, you don't need to know the answer to 1 x 5 to determine the answer of 3 x 7. This means that the GPU could allocate each of these calculations to a different core and calculate them *at the same time* (that is, parallel processing) whereas the CPU would have to do them one by one (that is, serial processing).

This advantage scales exponentially as the size of the matrices / data becomes larger; the GPU can just keep assigning tasks to new cores whereas the CPU still has to do the calculation piece by piece.

History of GPUs

The first GPUs were designed purely for image rendering either for use in computer games or other graphical software. At the time computers were only just introducing graphical user interface (GUI) based software as opposed to text-based commands and there was an inherent bifurcation between what the existing CPU and the purpose-built GPU were trying to do in your PC.

In a highly stylised sense, images on a computer are smaller squares (pixels) inside one big square[3]; each pixel is independent of one another in the sense that the one hundredth pixel doesn't care what colour the first pixel is. This independence allows pixels to be calculated in parallel so a chip that has a dozen extremely high-powered cores will be outclassed by a chip that has millions of low powered cores because it's a simple throughput exercise.

The introduction of graphics intensive software in the early 70's and subsequent explosion in commercial demand in the 80's created an intensely competitive environment of ~40 producers by the end of the following decade. Many were knocked out by the rapid product cycles. By the 2000s there were only three players left: Nvidia, AMD (through the acquisition of ATI) and Intel.

Image generation seems highly abstracted from machine learning and deep learning which GPUs are known for today, but they are actually very similar because for both the underlying workload is linear algebra[4] which is algebra but using matrices of numbers rather than individual numbers. You can see the similarity then: a matrix (a big square) full of numbers (little squares). You can also see the fortunate accident of history, GPUs were not created for machine learning, but they are better suited to the task than CPUs.

The issue was that GPUs were not "programmable" and hardware was configured with image rendering in mind: hardware was limited to specific size / dimensions (of images), there were limited shader outputs, instruction sets needed to compute integer and bit operations, communication between pixels was limited and, perhaps most important to general purpose uptake, there was no GPU computing software ecosystem.

Enter two cards from Nvidia in the mid-2000s: the GeForce 7800 and the now iconic GeForce 8800.

The GPU software revolution

The GeForce 7800 launched in 2005 and solved many of the hardware issues listed above with thread programs, global memory and shared memory but Microsoft's DirectX 9 SDK (a software package that governed how developers could interface with hardware) was not accommodating to the complexities required for general purpose compute instructions.

In 2006 it was the confluence of Microsoft's DirectX 10, four years of development and $475m that resulted in Nvidia's 8800 chip and a general-purpose software solution around the architecture of the GPU processor. This software was called CUDA, and it was a simple, explicit programming language that allowed developers to directly program the GPU for general purpose computing.

For the next 5 years, Nvidia's GPU market share was not positively impacted by this revolution, probably because the market was still dominated by graphics demand. By the

3 Pre-1987 GPUs did not render using pixels but calligraphic displays instead. Today though, the barebones (i.e. no shaders) process for generating pixels involves grouping vertices into primitives (triangles), rasterizing into pixel fragments and then combining those fragments into single pixels.
4 Calculus is required for stochastic gradient descent, the most common back-propagation algorithm in training, but this is still CPU bound in many cases.

end of 2012 the discrete GPU market was split: 1) Nvidia: 65.7%, 2) AMD: 34.3% and 3) Others: 0.1% and in unit terms wasn't growing.

In the following years, GPUs would become the most important (and usually the most expensive) component that could be found in any computing device and Nvidia's choice to develop CUDA along with AMD's financial struggles would be key drivers of their success.

Fast forward to today and the market is split: 1) Nvidia: 85%, 2) AMD: 9% and 3) Intel 6% and in unit terms, total market shipments were growing at c.30% y/y pre-COVID. Nvidia's market share is even higher when you drill down to the fast-growing server segment within discrete GPUs with shipment share estimates ranging between 90 – 95%, which in a sense is understated because at times Nvidia cannot supply demand, meaning AMD gets spillover demand.

Why software is an important moat and how it drives Nvidia's market share

To highlight how powerful Nvidia's software edge is when it comes to today's datacentre applications, we can compare Nvidia's A100 to AMD's MI250 chip; both the A100 and the MI250 are pitched at high-end datacentres. On hardware specs alone, the AMD MI250 has a slight edge in FLOPS/s, HBM memory and memory bandwidth. However, we can see that the Nvidia A100 outperforms in training throughput and performance by 20 – 50% on the MPT benchmark, a public large language model that mostly outperforms GPT-3. This is a huge performance difference from software, especially when you consider that the A100 is not even Nvidia's current leading generation chip…

To be fair, this is just one benchmark and there are many others where AMD would beat Nvidia but given the current surge in end-demand is driven by machine and deep learning applications for LLMs, we think it's important to compare the two on this metric.

There's also a price trade-off of around 20% per chip but AMD typically ships in smaller system configurations, so you must buy more chips to reach a given compute target. Also, if you train quicker per chip, there's likely a benefit to energy consumption and floorspace required. Nvidia is effectively printing money for their customers.

Source: MosaicML

CUDA, and other Nvidia software stacks, are closed source and available exclusively on Nvidia GPUs. This remains an enormous moat for competitors to cross as seen above. This is not to say that it's insurmountable, and AMD has made significant strides over the past 5 years plus their ability to control CPU architecture is powerful, but the data suggests they are still 2 to 3 generations behind.

This explains why Nvidia has gained market share.

Why the GPU market is growing and taking share from CPUs

The overarching driver of growth is that people need more computing power and within that, GPUs are better suited to the tasks they want to do so are taking share away from CPUs at an accelerating pace. We have read many articles stating that this is because a perceived death of both Moore's Law and Dennard Scaling have driven demand toward massively parallel compute solutions, but we think this is putting the cart in front of the horse.

When Moore postulated that the capability of computers can be expected to double every two years *as a result of increases in the number of transistors per chip*, CPUs were not designed with the same principals in mind as they are today. This is simply a function of end-demand, which today skews increasingly towards tasks that are better suited to GPUs.

As workloads shift to be more GPU bound, there is increasing pressure placed on CPUs to interact well will the GPU.

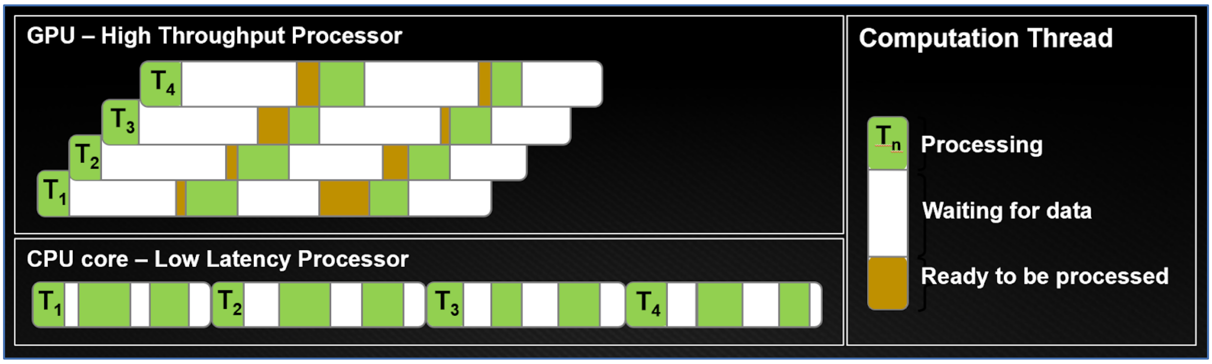

Source: Nvidia Developer Blog

Originally, CPUs were focussed on improving the clock speed of a low number of cores but have shifted to preference latency reduction - reducing the white lines in the chart above - by increasing core count per chip and introducing threads. This puts pressure on die size which is inextricably linked with heat absorption and electricity consumption and likely has a lot to do with the technical death of Moore's Law and Dennard scaling as it concerns CPUs. In addition, as die size increases, manufacturing yields decline as a constant incidence of manufacturing errors will have a greater impact on the number of successfully produced chips as you can now cut fewer chips from a wafer.

In its strictest sense, Moore's law links growth in computing performance to the number of transistors on a chip. It's true that chips have physically begun to deviate from this definition but in performance terms, or what we would term the "intended outcome" of the law, they are actually outperforming Moore's expectations.

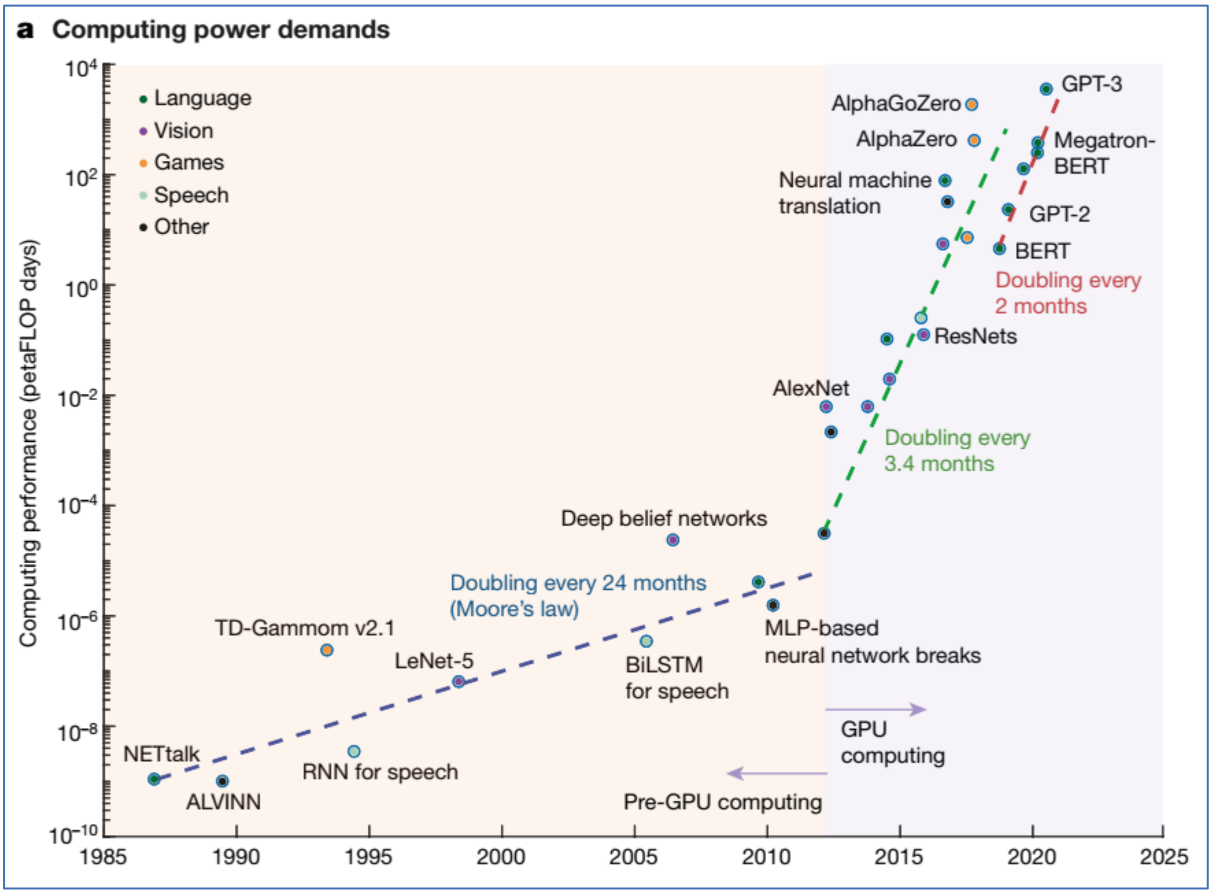

Source: Mehonic, Adnan & Kenyon, Anthony: Brain-inspired computing: We need a master plan, 2021/04/29

The above chart shows the computing power demands of popular software models through time. Before Nvidia released Cuda in 2012, we were roughly doubling our processing power every two years. *Then from 2012 to 2020, we were doubling every 3.4 months*. **Today, we are doubling every 2 months**.

Clearly there is an incentive to buy these chips and ***we should be able to see a bifurcation between CPU and GPU producer revenue***.

Summing the datacentre CPU revenue of AMD and Intel and comparing them to Nvidia's GPU revenue shows that Nvidia holds ~35% of the total processer spend as at the end of 2022 but has been taking share rapidly. Using consensus estimates, we can see the market expects the CPU to GPU shift to hold but not continue in any meaningful way. We're not so sure.

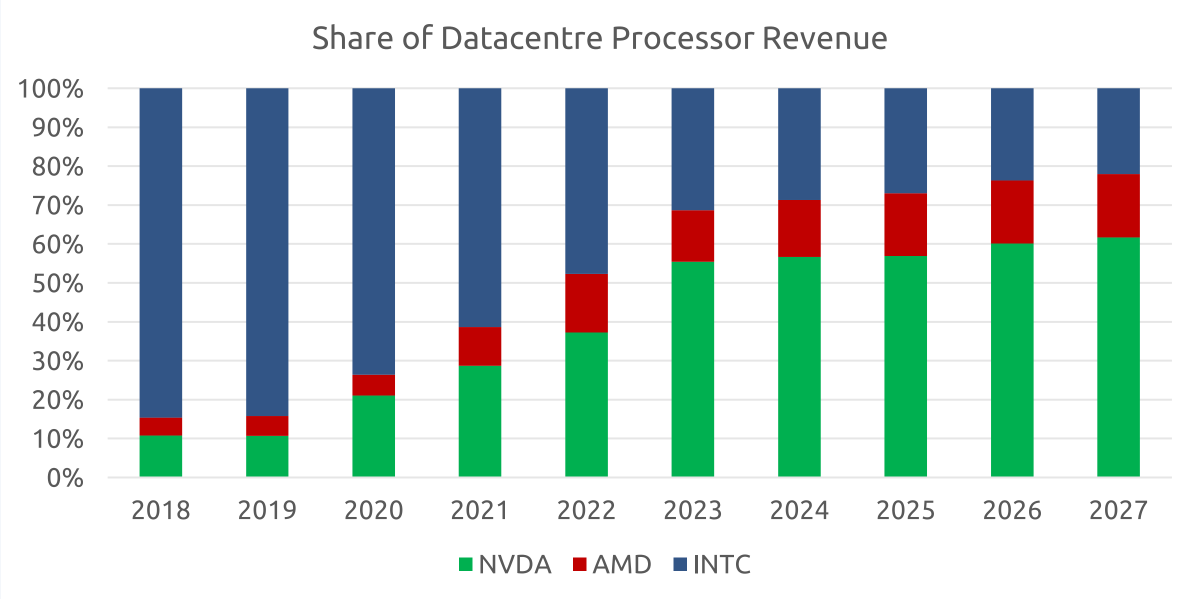

Source: AMD, INTC & NVDA

That GPUs are taking share from CPUs is self-evident. The question becomes: what is driving this end demand and how fast is it growing? To see this, we can take one more step through the supply chain to the datacentre's revenue.

End Demand Growth

Datacentre demand comes from five (to be six) buckets as outlined below. The nomenclature is not standardised and there is some overlap but, in any case, this is how Nvidia has defined them.

1. Supercomputers: Universities, research labs, etc.
2. Enterprise computing: IBM, etc.
3. Hyperscalers: Meta, Amazon Web Services (AWS), etc.
4. Cloud computing: AWS, Azure, Google Cloud Platform (GCP), Oracle Cloud Infrastructure (OCI), etc.
5. AI factories: TikTok, Meta, etc.
6. (Not yet out) Edge computing: Low number of concrete datapoints as yet.

For public cloud providers like AWS, Azure, GCP, OCI, etc. their revenue should indicate end customer demand as they charge on a $/hour basis. So, we can infer total system growth from the sum of their revenues, which are run rate ~20% y/y

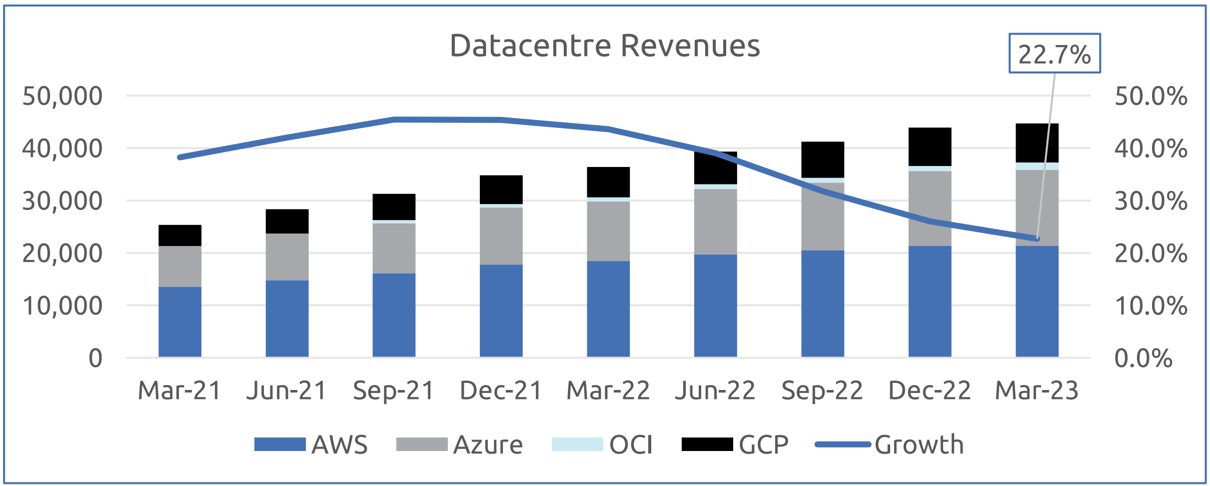

Source: Geometrica, Bloomberg, Amazon 10Qs, Google 10Qs, Oracle qtly 8Ks and Microsoft 10Qs

Who will generate high returns on capital within the AI ecosystem?

We have established that the market for GPUs is growing but understanding where value accrues is crucial when estimating long run returns and thus shareholder value creation.

One way to infer this to understand the datacentre's unit economics. A rational datacentre operator is likely solving for some IRR before spending billions of dollars on GPUs. Amazon releases profit, capex and asset granularity on AWS, so we can infer unit dynamics from here.

We can see that AWS's non-lease related CAPEX[5] burden as a % of sales was initially decreasing but started to rise in 2019, exactly when NVDA's market share as a % datacentre processor revenues started to inflect.

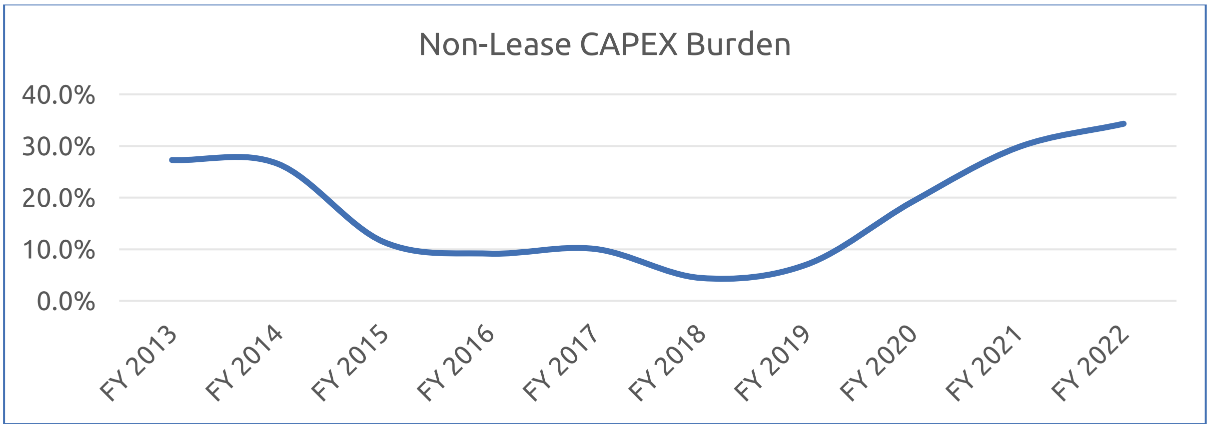

Source: Geometrica, Amazon

[5] *Amazon states that non-lease CAPEX includes "…investment in technology infrastructure, the majority of which is to support AWS…". This definition is slightly vague but focussing on non-lease related CAPEX should isolate the spend on chips better than looking at the total.*

Our proxy for return on capital, implicitly assuming no debt, which would decrease ROIC anyway, shows that AWS' cash ROIC has been declining. Coming at it a different way, incremental ROIC, whilst higher than total ROIC, is similarly trending down, even using total CAPEX.

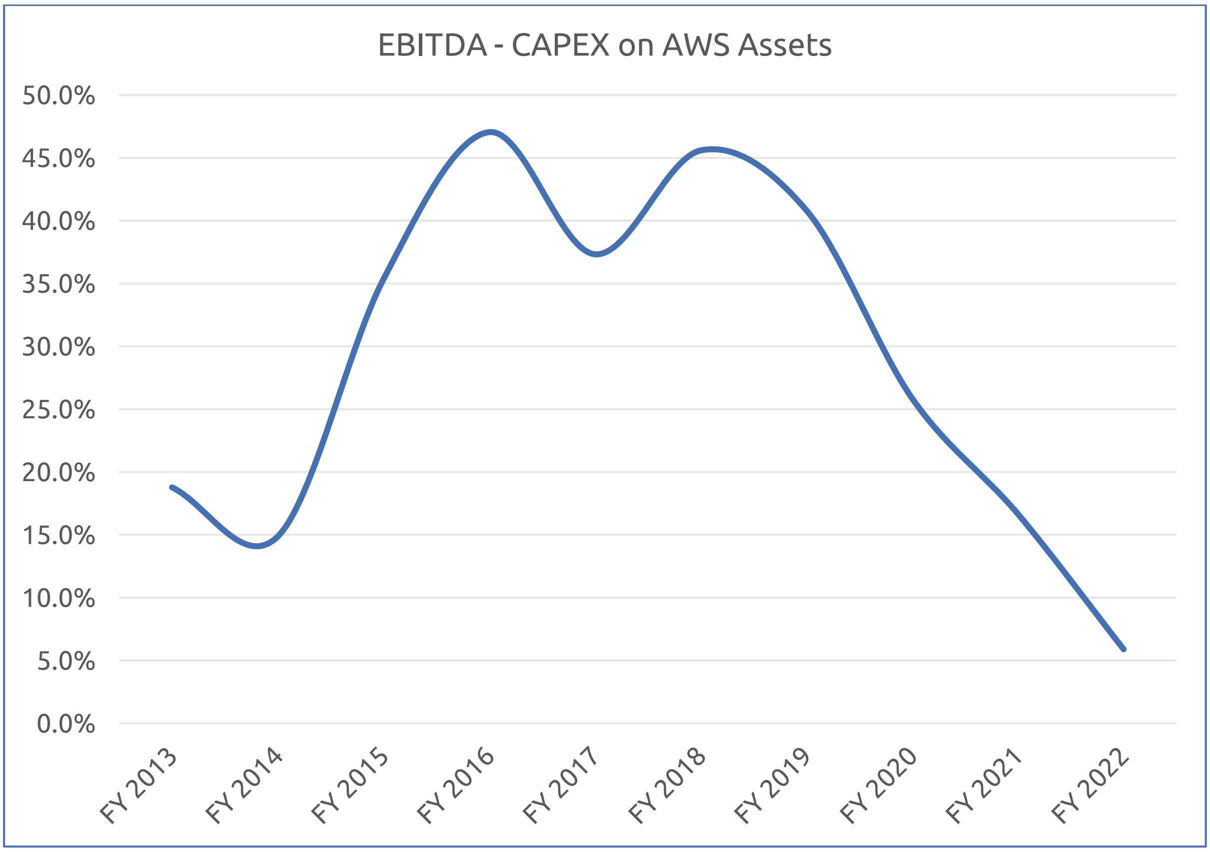

Source: Geometrica, Amazon

The key inference from this is that Nvidia is extracting more rent from the datacentres through time.

We think this is occurring because the datacentre is increasingly akin to a flat cost-curve commodity industry which logically should cede margin to the asset with real competitive advantages, which in this case is the Nvidia GPU ecosystem.

We hypothesise that the datacentre market likely shakes out to very few players over time and that their long-run returns are likely much lower than Nvidia's. This intuitively makes sense if the key differentiating asset in the offering is the chip, not the services.

If this is true, then competition should be increasingly price driven. And indeed, when we look at datacentre pricing trends through time, it appears that price is the largest driver of market share shifts. **As an example, Microsoft's Azure has consistently priced below AWS over the years and has been the largest share gainer over this time, with AWS the largest cedent.**

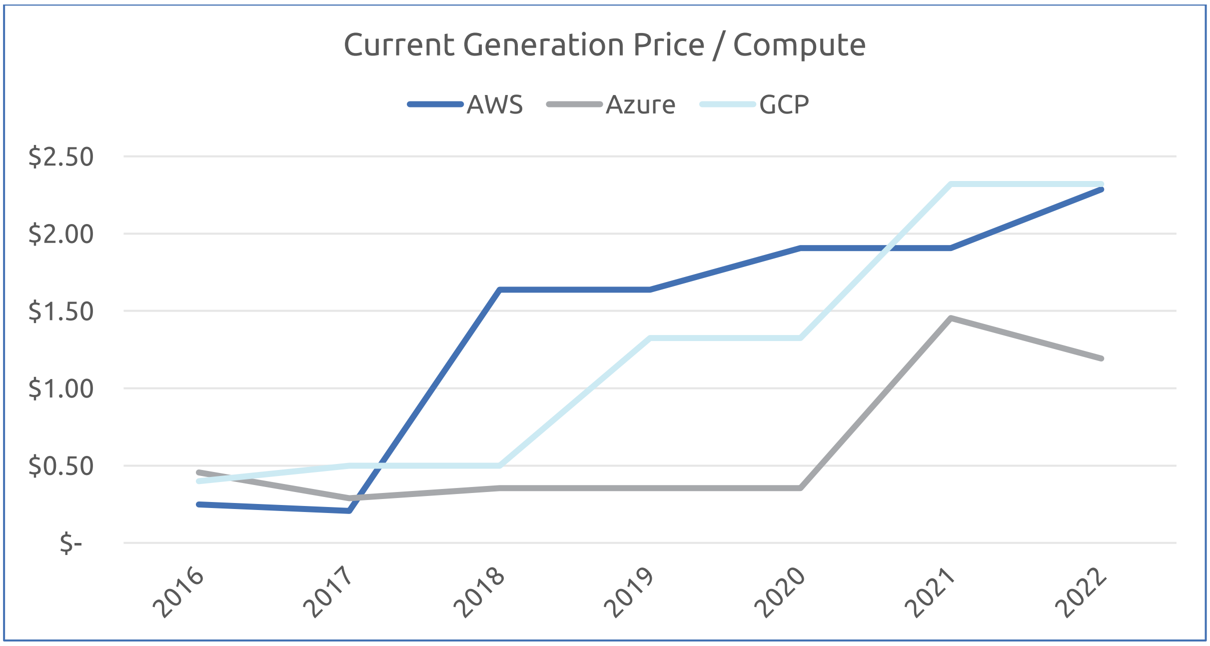

Source: Geometrica, Redmonk

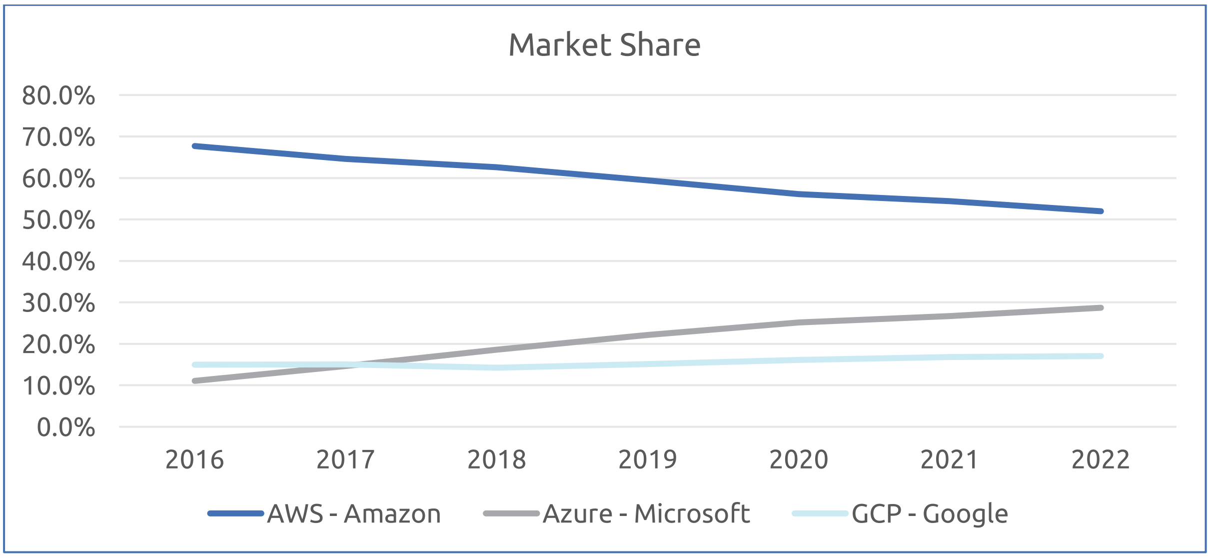

Source: Geometrica, Amazon, Microsoft, Coogle

Price per unit of compute has also become increasingly homogenous across operators which adds to the suggestion that price is a key battle ground rather than any differentiated offering.

The charts below plot price on the x-axis and processing power on the y-axis. A steep line indicates that a datacentre charges very little for increasing units of processing power versus a line that is flat which indicates the datacentre charges comparatively more for that extra power.

Note that today, the lines are all linear and the highest point of price variance for a given level of compute is a 10th of what is was in 2016.

Compute price – 2016

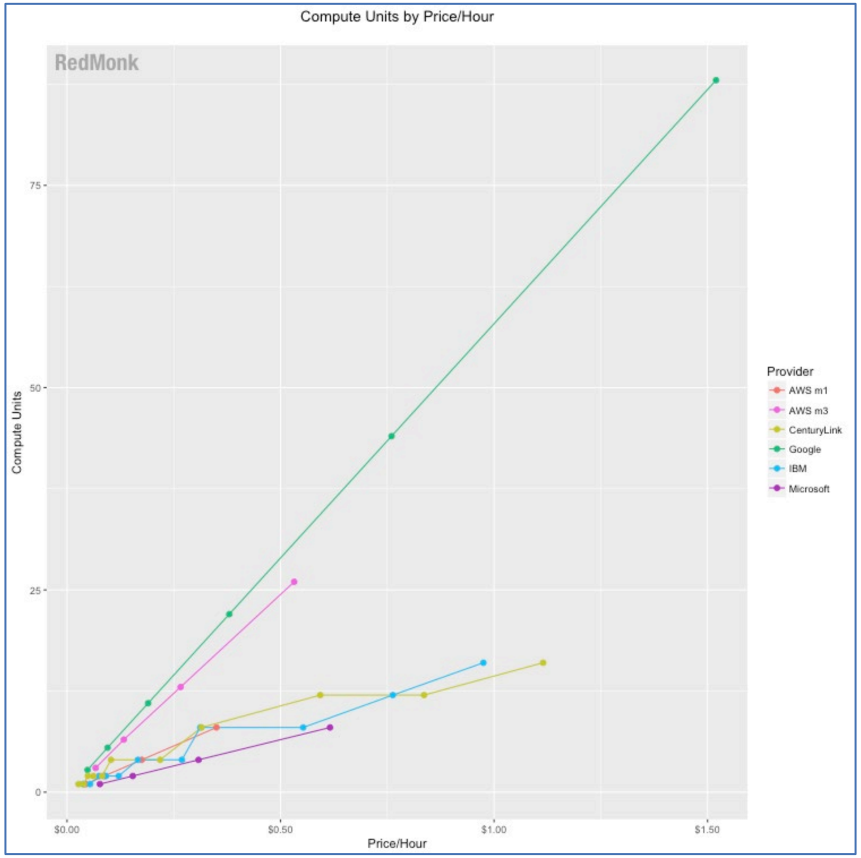

Compute price – 2022

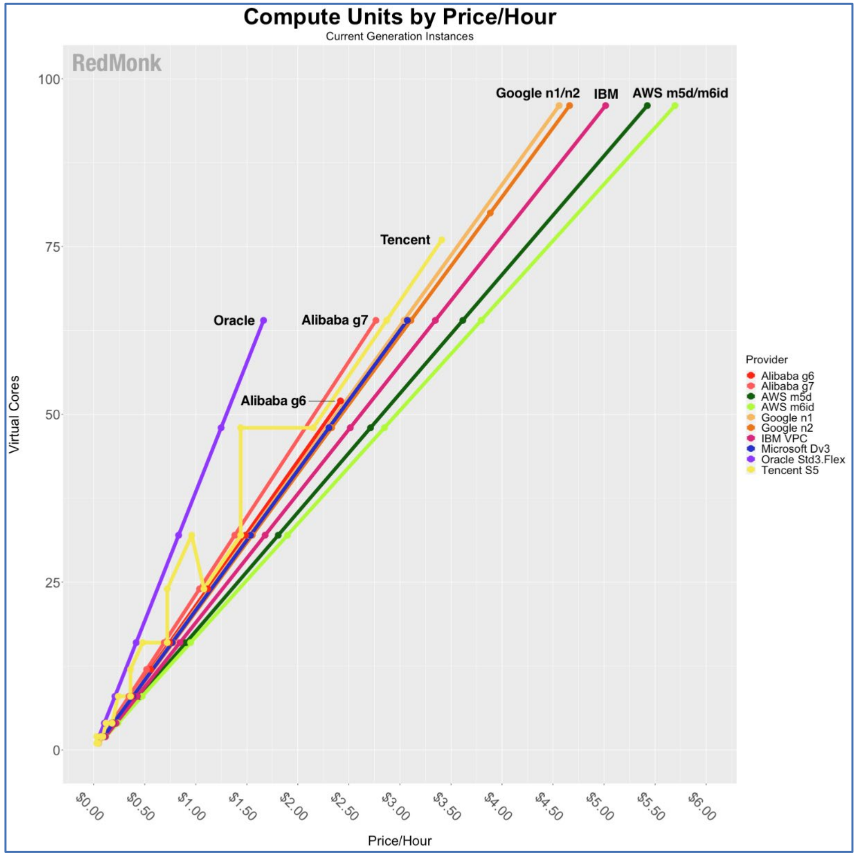

Source: Redmonk

The nature of cloud competition has clearly changed. In the early days, there was high variance between pricing, indicating differentiation between products. Now, there is very little indicating that price is an increasingly important point of competition. Indicative of commoditisation.

Even outside of price, competition has shifted to parts of the cloud offering that indicate decreasing asset differentiation:

- Some providers have started to design their own chips in an effort to optimise price to performance without having to change the price points of their instances.
- Wider infrastructure capability, including ARM CPUs / software support, GPUs and bare metal[††], despite being more expensive to provide, are important points of differentiation.
- Egress costs[‡‡] are an emerging battleground.

All of this suggests that pricing on standard instances has finally reached a commodity status. So, **there is a massive financial incentive to own Nvidia GPUs because without them, they only have highly commoditised offerings that don't appear to have long-term pricing power.**

The final question: how big is this opportunity?

Again, there are a few ways to do this. Most simply, in June this year, AMD stated that the market for accelerated compute will be ~US$150bn p.a. by 2027 and that triangulated closely to the figures we got to.

Consensus research estimates are forecasting that Nvidia's total datacentre revenue, which also includes their Networking solutions which we estimate at ~$4.5bn annually growing at ~30%, reaches US$58.1bn by 2027. For this to be true, one or more of the following must also be true:

- US$150bn is the wrong number; and / or
- Nvidia rapidly loses more than 50% market share over the next 5 years; and / or
- The CPU to GPU shift inverses; and / or
- The price of accelerated compute must fall rapidly.

It certainly won't be a straight line, Nvidia's revenues are still influenced by the product cycle, and recent stock price action likely reflects a positive deviation from the long run trend.

Regardless, **we think the end point is still underappreciated by the market.**

[††] Bear metal refers to a single-tenant server which gives users a higher level of control over their infrastructure (e.g. they can chose operating system, tune hardware, etc.). It effectively mimics a local server. This compares to a standard cloud server / instance that has multiple users sharing the same physical computing resources which allows for full utilisation of the computing hardware but can result in the "noisy neighbour" problem where one user monopolises the shared bandwidth / compute resources.

[‡‡] Egress fees are incurred when a user transfers their data out of the cloud server to which it was originally uploaded. This could occur if you choose to change operators and want to take your data with you. Fees typically range between 5 – 20 cents per gigabyte of data and are incurred each time you make a transfer.

There are analogues: IBM's Mainframe, Intel's PC CPU and Apple's smartphone. Like these, the adoption curve, we think, is exponential, yet the market persistently projects linear development.

The smartphone market in particular illustrates the importance of delivering a wholistic platform to capture the industry's profit margin. Nokia originally dominated but by 2008, Apple's iPhone platform blew them out of the water. Apple owned the chip (AX SoCs), the operating system (iOS), the hardware (iPhone) and the end software (App store). The adoption curve was exponential and cased a growth inflection for the entire industry.

Chart 19: Cellphone / Server Operating Profits ($B)

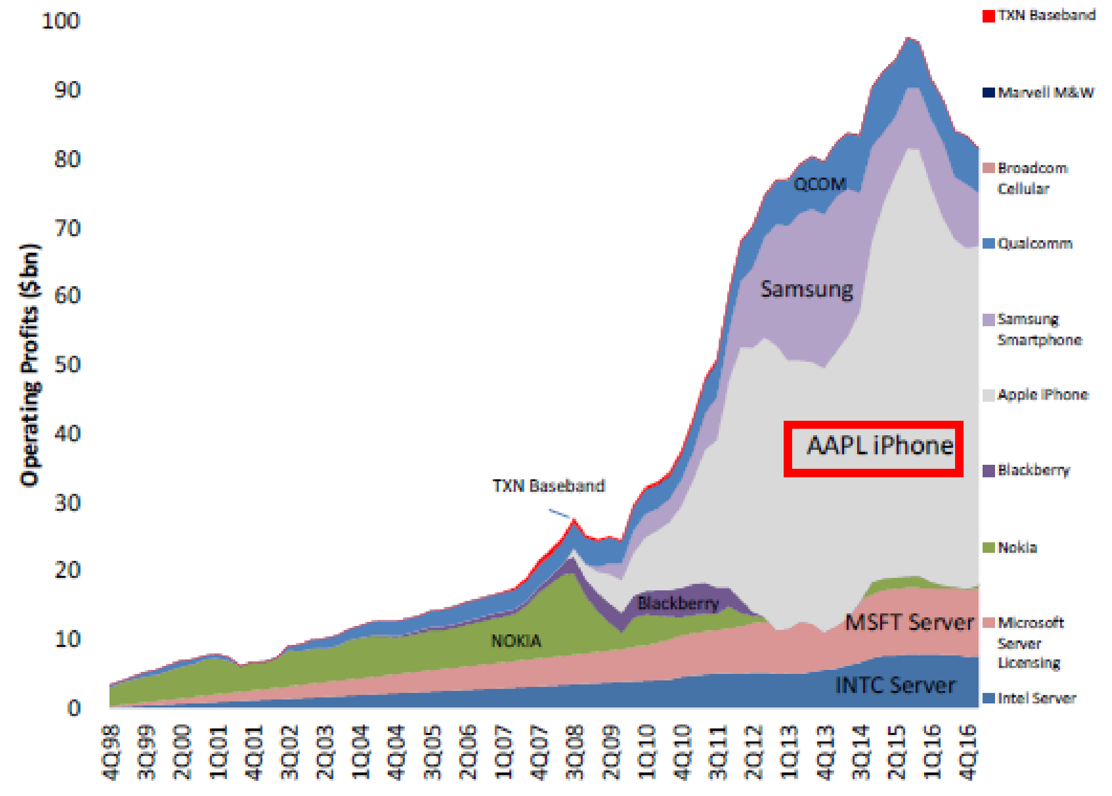

Source: Jefferies, company data

In summary

1. End demand for compute is growing at ~20%.
2. But within that, GPUs are taking share away from CPUs.
3. Nvidia is the largest share gainer so will grow well above this system demand.
4. Nvidia's software lead is a significant moat and a sustainable driver of market share growth going forward.
5. Datacentres operate in an increasingly commoditised market and there is a sustainable financial incentive for them to cede margin to Nvidia long-term.
6. The core adoption trend of accelerated compute is likely to be exponential to start with, noting that there will be perturbations along the way.
7. **The financial outcome of this is CAGR profit growth of 2x consensus estimates over the next 5 years.**

MANAGER PERFORMANCE HISTORY§§

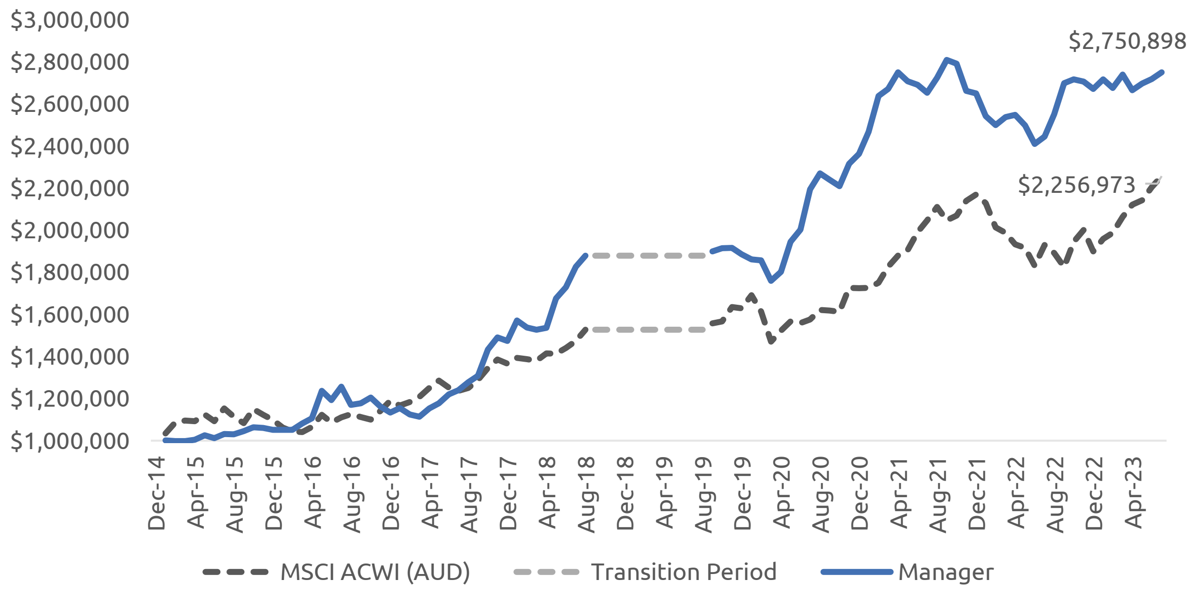

31 Jul 2023	Strategy Inception	Strategy inception pa	Geometrica inception pa	CYTD	1 year	1 month
Founder***	175.1%	+14.27%	+10.25%	+2.98%	+12.54%	+1.21%

* Manager left CVF in Sept 2018 and began Geometrica in Sept 2019 NB: Performance period is from 5 Jan 2015. Performance is net of all fees.

ASSET ALLOCATION

Country	Long	Short	Gross	Net
Australia	2.8%	0.0%	2.8%	2.8%
Americas	43.9%	(0.7)%	44.5%	43.2%
Asia	9.2%	0.0%	9.2%	9.2%
Europe	16.5%	0.0%	16.5%	16.5%
Total	72.4%	(0.7)%	73.0%	71.7%

§§ Manager left CVF in September 2018 and began Geometrica in September 2019. Performance period is from 5 January 2015. Performance is net of all fees.

*** Founders Class units – Lead Series. Small variations will occur between unit classes and series based on differences in timing and terms. Source: Mainstream Fund Services, the Fund's external administrator and calculation agent.*

GROSS & NET EXPOSURE

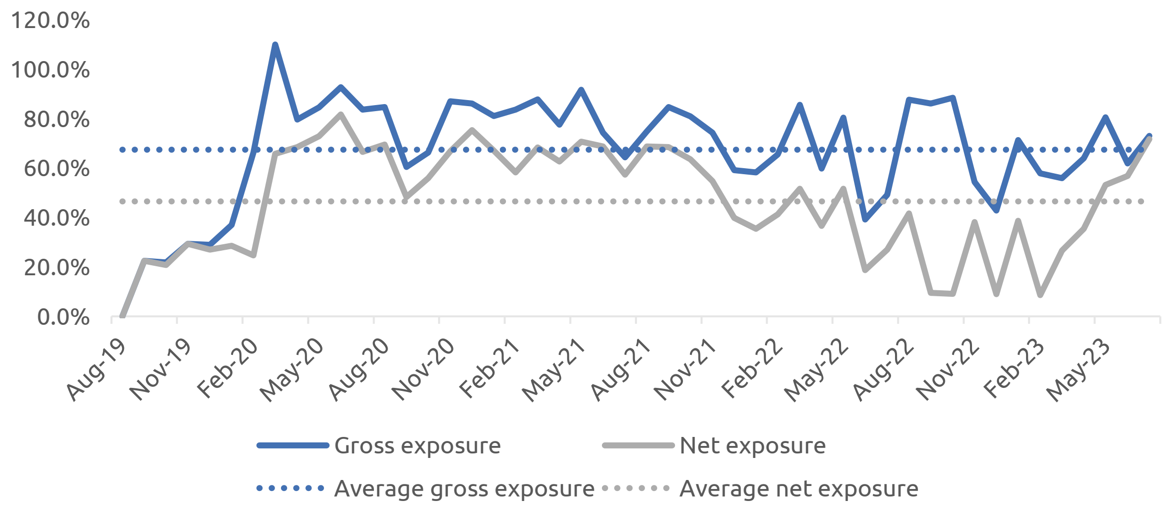

GROSS EXPOSURE BY MARKET CAPITALISATION

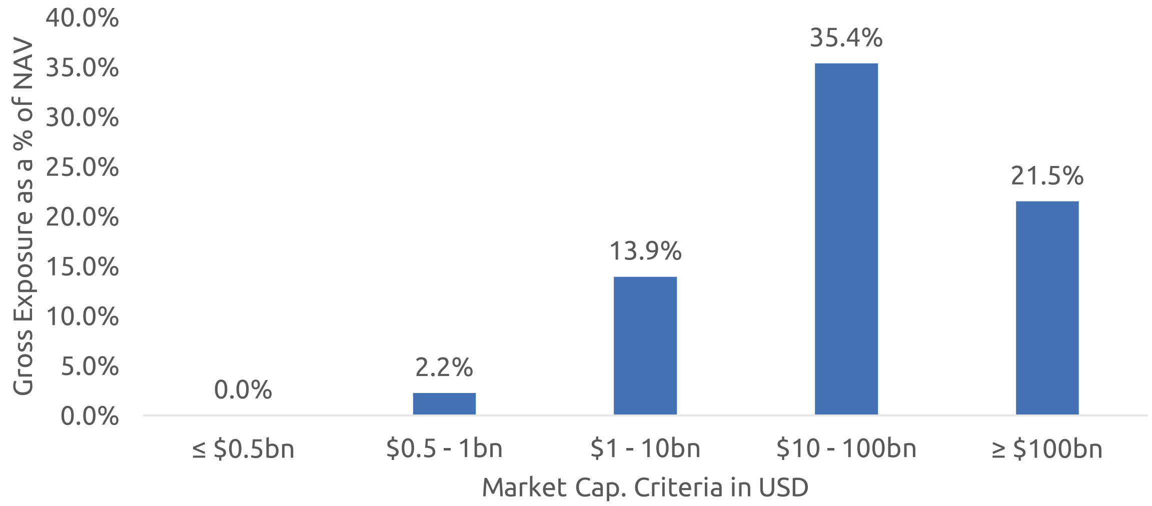

FUND OVERVIEW (ALPHA UNITS)

Fund	Geometrica Fund
Structure	Wholesale unit trust
Mandate	Global long short Mid-cap focus
Gross exposure range	0 - 200%
Net exposure range	up to 100%
Single stock long limit	15% at cost
Single stock short limit	5% at cost
Buy / Sell Spread	Nil / 0.25%
Investor Eligibility	Wholesale only
Platforms	Ausmaq, Hub24, Powerwrap, Netwealth
Fees	1.5% management (+GST) 20% performance (+GST)
Benchmark	RBA Cash Rate
High water mark	Yes
Liquidity	Monthly
Administration & custody	Apex