

The Outlook for NVIDIA & Generative AI;

An investor perspective from the BlueBox Global Technology Fund

There are two stages in the deployment of a major new technology:


Stage 1 – Building:

This is hardware intensive, involving enormous up-front investment to catch up with latent demand, closing a huge gap. This lasts for years, but not forever – at some point build slows, which means an actual decline in revenue for build-focused players, e.g. NVIDIA, as discussed below. We then move (probably gradually) to Stage 2 ...

Stage 2 – Operating:

Hardware spend drops to a combination of (a) maintenance of the installed base and (b) growth of the installed base in parallel with end demand, which eventually trends towards GDP growth. Running costs and software expenditure come to the fore. We work out which layers in the stack are the pinch points, which is where almost all the long-term profits will be made from the new technology.

For Stage 1 of GenAI, the numbers are staggering. A single world-scale data centre is now expected to cost around \$50bn. According to the House of Commons Library, UK defence spending was £62bn last year, or about \$85bn, so not even enough to build two of them. Not many of that size will be built, but plenty of “smaller” ones are, leading to this chart, with growth in business investment in IT hardware growing at 30% a year, an unprecedented level:



Source: US Commerce/BEA, Tech Dynamics August 2026

Hardware spend cannot grow at 30% for very long, as the entire industry appears to be constrained to around 15% overall growth:

Equal-weighted S&P Technology Index, log scale – 15.2% annual growth trend lines



Source: Bloomberg

The constraints include the following:

- Semiconductors
- Semiconductor capital equipment
- Raw materials
- Energy
- Staffing
- Permitting
- Subsidies

These seem to limit the IT industry to growth of about 15% a year overall, however much money is thrown at it. At the moment a vast amount of money is being thrown, and because the industry has probably already hit that 15% limit, each additional dollar just creates price inflation for the constraining items: processors, memory chips, networking equipment, energy infrastructure, generating equipment, and people with experience constructing data centres. Permitting is also an increasing issue, as it gets harder to get permission to build data centres in new locations – this may lead to promises of big financial incentives to local populations/authorities to accept them. Those are the pinch points today, creating huge short-term upside for everyone from memory companies to data centre engineers. It doesn't last.

Despite the price inflation, supply constraints cap the real growth rate of the infrastructure – it takes longer to build things than planned, and costs are rising fast. Amazon just increased its capex forecast for this year from \$200bn to \$220bn, but that upside wasn't to build more, it was just to cover the increasing cost of the memory chips required for the original plan! In addition, as hyperscalers and others increasingly fund their investment with debt rather than cash flow, their cost of capital is likely to rise. If projects take longer and cost more, with a higher cost of capital, the returns expected from them fall, making it less likely that the builders will eventually earn enough money to justify the investment, so funding starts to dry up, pulling forward the inevitable day when building slows.

At that point we are not just looking at revenue growth slowing for NVIDIA (and memory companies, etc.), but actually going into decline. That decline probably won't be huge (maybe 10-15%) and may not last very long (possibly a year or two), but it makes a complete mockery of all the financial models that assume a gradual deceleration of revenue growth from today's levels. At present consensus forecasts of NVIDIA's revenue growth for the next few years are as follows – a nice steady deceleration to single digits:

Current year FY27	FY28	FY29	FY30
+83%	+45%	+22%	+7%

In reality, when we move from Stage 1 to Stage 2, and data centre capex goes from “build to catch up with unserved latent demand” to “build in line with growth in served demand”, revenue growth is likely to turn negative: not only would capex decelerate sharply, but also recent price inflation is likely to unwind rapidly, while double-ordering and inventory building would reverse. If next FY (FY28, ends 31 January 2028) turns out to see a 10% decline in revenue, that's way below current expectations (almost 40% below consensus). It would be accompanied by compression of both margins and price/earnings multiples, which together create most of the downside for the stock price.

The good thing is that everyone in their heart of hearts knows that this is what will happen, so it's already partly in expectations: NVIDIA, currently the most important company in the world, in middle of the biggest investment boom history has ever seen, is at just 19x next year's earnings. At some point – maybe next FY, maybe the one after – that will turn out in hindsight to have been 38x earnings, so not that cheap after all. That's also why you're not seeing TSMC go flat out for capacity expansion, as management knows from experience that it is better to cut the top off a demand spike by rationing through price, rather than be left with massive new capacity just as demand (temporarily) slows sharply.

The only place you don't see this is in published sell-side forecasts, because no one knows when to put the down year. So they don't put one in their models at all, just a comfortable glide down to 7% growth.

We may also have seen the market acknowledging the approach of the slowdown with a change in tech industry leadership: since late June, semiconductors (of which NVIDIA is the most valuable) have been underperforming, and the recent new highs for the sector have been driven instead by software names, the businesses that are likely to thrive when GenAI moves to Stage 2, “Operating”. Semiconductor stocks don’t wait for a slowdown to become obvious – they tend to underperform for months in advance: there is no leading indicator for the semiconductor industry!

This doesn’t mean that NVIDIA is a bad investment: in the short run (e.g. Wednesday night’s results) it is probably absolutely fine and possibly very good. In the long run it is a key player at the centre of technology, and probably a great investment. However, in the middle there is likely to be a period of serious disappointment ... but at least that is already partly priced into the stock. We just don’t know when that disappointment will arrive: it could be this week, it might not be for 18 months; but no industry of this size can grow at 30% for very long. If the industry slows to a more sustainable rate, NVIDIA cannot escape a sharp, but temporary, correction.

That is not a reason to sell the Technology sector today. GenAI is a transformational change, powered by the tech sector, and the industry is likely to remain on its 15% long-term trend for a long time into the future, albeit with plenty of volatility. The correction, when it comes, may start at a much higher level – booms generally go on longer than one expects – and probably doesn’t last that long. However, investors should be aware that the current semiconductor leadership will not last, and indeed it may already have been handed on to the software industry.

William de Gale

26 August 2026

This document is intended for professional investors only and does not constitute investment advice or a solicitation to invest. Past performance is not a reliable indicator of future results. The BlueBox Precision Medicine Fund is managed and distributed in accordance with applicable UK and Swiss financial services legislation. All clinical data referenced herein is drawn from publicly available conference presentations and company disclosures at the time of writing.