WALTER SCOTT

PODCAST TRANSCRIPT

Talking Research
The AI gold rush: infrastructure, chips and the race to scale

First published: 12 December 2025

Katie Boyce: What do a font designed for pixel clarity and the world's most advanced semiconductor chips have in common? In 1996, Matthew Carter designed the Verdana font to solve a specific problem: poor legibility of traditional fonts on low resolution computer screens. At the time, most monitors had coarse pixel grids, which made typeface details blur or break down at small sizes, but Carter's approach meant that every design choice was guided by the output – how the font would render on screen, rather than by classical print aesthetics.

In ASML's Silicon Valley lab, engineers do something very similar using a process called optical proximity correction to pre-distort patterns so that they create a perfect image on a wafer. A 30-year-old font and semiconductor chips may seem worlds apart, but in both cases, there's a surprising complexity in making tiny images as perfect as possible.

Hello and welcome to Talking Research, a podcast from Walter Scott. I'm Katie Boyce, an investment writer at the firm, and today I'm joined by investment manager Tom Miedema, who recently visited California and met this ASML team as well as more than 20 other tech and healthcare businesses.

Tom, thank you for joining me today.

Tom Miedema: A pleasure.

KB: Before we get started, a reminder that the podcast is intended for investment professionals only and should not be construed as investment advice or a recommendation. Any stock examples discussed are given in the context of the theme being explored, and the views expressed are those of the presenters at the time of the recording.

Tom, welcome back to the podcast. Plenty to talk about today, so let's dive straight into your trip. Why were you in California?

TM: It pretty much was a typical research trip for us. So, you know, we spent a week and a half in California, went up to Seattle as well, meeting with more than 20 companies across the spectrum, really to get a feel for what's happening with the companies that we own and to look formore information on certain topics. And we'll talk about that soon, I think, as well as a few prospective investments as well.

KB: And what kind of companies did you meet with?

TM: It was really broad. It's California so there was a lot of tech, tech hardware – the likes of Microsoft, Amazon, Arista technologies [sic], Broadcom etc; a lot of software companies – so your Adobes, Intuits, ServiceNows; and then a lot of healthcare as well. So on the healthcare side, it would be people like Intuitive Surgical, ResMed, Edwards Lifesciences that were amongst the companies we met.

KB: And there's obviously so much talk around AI right now and whether or not we're in a bubble. Given your recent conversations with the companies that you've mentioned, what's your take on this?

TM: We wish we had the answer. I think the reality is that most people would accept now that we're in an inflating bubble. And the question is, how large will it get before we get to some sort of change in that trajectory? I think whenever you have a hardware cycle, at some point it will slow and, at some point, it will reverse. The question is how far are we through that? I think there are various signs we can talk about through the trip, but it's not going to go yet. But I think clearly we are in an inflating bubble.

KB: And any particular areas of the industry?

TM: Yeah, I think it's most obvious in the kind of euphoria in the VC world. Funding in private businesses is going through the roof and they're funding a lot of what's going on. And then, on the other side, the infrastructure build is tremendous. Every day you've got news stories about another deal around buying compute power and another deal for funding another gigawatt or a larger data centre complex. So, clearly, those are the hottest areas today and the

areas that the market is fretting about the potential returns on those investments.

KB: That's a good place to start; on infrastructure. You met Microsoft, as you mentioned. How does a company like that balance the risk of overcapacity with the need to stay ahead in infrastructure? They've clearly spent a lot of money on this stuff.

TM: Yeah, I think it's really interesting if you can track what both Microsoft and Amazon have been doing in particular. When you look at it, they've been pretty careful to manage what they're doing on their books, so which data centres they're building, what kind of business they want to do, and actually allowing others to take on some of the more risky aspects of the investment.

KB: So what would those be?

TM: If you look at the deals, for example, that the neoclouds, the likes of CoreWeave and Lambda Labs but also you can think about Oracle a little bit in the same way; they are taking on the building of data centres mostly for AI training, mostly for single customers. You're saying to OpenAI or an Anthropic or someone else, I'm going to provide you compute capacity, and it mostly is just Nvidia chips that you're renting, for a given period of time, and it will dedicate all of that capacity to one customer.

Whereas a traditional cloud model would be much more multi-tenant. The classic AWS, which is Amazon's business, or Azure, Microsoft's business, they are building a data centre which can serve hundreds, thousands, tens of thousands of clients, and you're not dedicating that to anyone. Whether the clients are running standard compute capacity, they want to store things, they want to run software applications, they want to be a tenant for many, many different customers. And that's a much lower risk model than renting out compute capacity.

It's also much higher margin, typically. And I think what you can see in the activities of those companies is that most of what they're focusing on is: can we reuse this for our core cloud businesses? And we'll focus on the growth there. I think it's a very sensible risk management strategy from our perspective.

KB: That does all sound very sensible. But what are some of the biggest challenges these companies, such as Amazon, are facing?

TM: I think it's fascinating when you're speaking to people who are closer to the data centre buildout – definitely your Amazons, your Microsofts, but many, many others – that is when you're looking at what's happening on the ground in some of those states, because it's mostly in the US, where the big data centres build is going on. It is turning into that Covid-era reference that I heard talked about many times: "a daily knife fight" in the supply chain.

KB: So what does that mean?

TM: What that means is trying to get the stuff that you need, get the people that you need in the right place to actually get these data centres stood up, and you're seeing that now really creak for certain companies. You've seen a bunch of neocloud vendors report results and guidance that disappointed the market, because they just cannot get this physical infrastructure built. I think it's very likely that we start to see differentiation between those who have experience, have very strong relationships and supply chains, really know how to execute on this stuff – which I think the big guys do – versus those people that are less experienced and I guess newer to the business. It's going to be fascinating to see how that pans out.

KB: What happens next?

TM: Great question. So I think what happens next is that it becomes very stark, the differences there. You obviously have the demand questions, that looks very solid at the moment but we'll see how that tracks going forward. But I definitely think you'll see even more disconnect. We've seen it multiple times over in different bubbles. Oil and gas is maybe a very good proxy, but when you really get into an up cycle, when everyone is trying to fight for the same resources, typically what that means is prices go through the roof and delays start happening and they can get really extreme. That would be my expectation across the board in a lot of these places.

KB: And Amazon's ramping up its own AI chips, investing heavily in GenAI infrastructure. What's the significance of these moves?

TM: Yeah. A lot of people are talking about custom chips, custom ASICs, the best known…

KB: And just to clarify, ASICs are "application specific integrated circuits", but what are these?

TM: The best known of those are the TPUs from Google. Google really started this in an AI context in 2013/2014 but ASICs have been around forever. Whenever an application gets big enough, people start looking to make custom chips specifically optimised for that single application and then, typically, over time,…

KB: Are they more expensive?

TM: On an individual basis, at small scale, they're more expensive. But, if you have big enough scale and you get enough efficiency from those chips,

they can be much, much cheaper than buying the alternative. And the alternative would be known as merchant silicon, so, in the AI world, that would be buying an Nvidia chip or an AMD chip, those would be the merchant vendors. And the alternative would be that I go out and I design my own chip, which will cost me multiple hundreds of millions of dollars per design, and then I have to go out to a TSMC to manufacture that. So, it's a massive advantage if you have scale and you can really optimise your entire stack, your entire software and your hardware around your own infrastructure.

So, Google does that really effectively and has done for years. Amazon has done that really effectively for years as well, so their cloud business has always, always been focused on custom ASICs, typically CPUs called Graviton, and now they are just about to release their third generation of their own custom silicon, which is called Trainium, which should be announced in December. This is going to be really important for cost going forward and for their ability to differentiate their business from others.

KB: I mentioned ASML in the introduction, it's clearly a linchpin in the whole infrastructure story for semiconductors. And possibly I'm not sure I quite got across the excitement of the team that you met in San Jose, perhaps you could do a better job and tell me what kind of things they were telling you?

TM: I think you did alright. Most of ASML is based in the Netherlands in Veldhoven, but they have this team that have been based in Silicon Valley for a very long time. I actually had the opportunity to meet with them a number of years back, so this was kind of a revisit. And what they do is use a lot of algorithms and a lot of compute, and try and optimise the manufacturing process of semiconductors based around ASML's tools.

KB: So, to that example I used earlier, the kind of almost reverse engineering, you explained that very well to me before if you could try that again?

TM: Where they are working on this process is thinking about TSMC's leading edge process. They're thinking about the process in four years' time, what TSMC is going to be doing in 2029-2030 and so you're working at incredibly low – sub two – nanometres, which is really tiny. And you're trying to optimise the creation of patterns through a series of mirrors but using light onto the wafer. Light has all these wonderful properties and the materials that you're trying to create a pattern on have all these different properties, and you are trying to use all sorts of algorithmic techniques to finely optimise how many good copies of the pattern that you're going to get at the end of the day. It's incredibly complicated stuff. These guys are absolutely world leading in what they do so it's a real privilege to spend an hour with them.

KB: Another element of the AI debate is the suggestion that software-as-a-service applications are dead, given the rise in AI-based apps. As you mentioned already, you met with a lot of software companies. What did you hear from them?

TM: Yes, I think this is one of the most interesting questions right now. If I can stand back and go, okay, people are very excited about AI infrastructure so all this spending, all this money on AI, that only makes sense if you're going to use that eventually. How will you use that? You'll stand up data centres and then use inference, so you'll run the models. I spend all of my investment on training models, then I'm going to run the models therefore the business of running models and making use of the models has to be much larger than what I spend on training them, otherwise AI is not going to work.

And the secondary part of that is that the applications that I end up building and I sell to an end customer – whether they be enterprise or consumer – again, have to be bigger than the cost of running those applications. I'm going to pay a cost to run an application or provide a service using AI, maybe it's a kind of medical advice, maybe it's a chatbot, whatever that might be, and my cost for running those things is going to be inference, so actually the biggest market of the three is going to be the application layer – and that's where we're going next.

The huge question is who wins and who loses? Is it going to be the incumbents that are able to adopt AI, bring it to their customers and leverage that and accelerate their businesses? Or is it going to be this whole sea of AI startups that are able to do the same thing? But somebody is going to win that and it's going to be a huge, huge market, and we're spending so much time on that right now.

KB: And in your view, which companies are going to be the winners?

TM: Clearly, this is a huge, huge question so I wish I had a perfect answer for that but I think we've got quite a lot of clues in terms of what we think is going to be really important. One of the things that's going to differentiate players is just the basics of competition. How quickly can you iterate? AI in a software world speeds up innovation, speeds up your ability to bring products to market, speeds up your engineering teams, so the AI startups are really good at doing that but some of the incumbents are becoming also really, really good at doing that. They know they need to do that to survive. We met the CTO of Intuit, who was very impressive.

KB: Intuit is a tax and accounting service.

TM: Absolutely. They are moving at a real pace for a very large business. And that's the thing we need to test – are you innovating fast enough to win that competitive race?

KB: Within the space, there's also a lot of talk about vibe coding, the idea that some of these software agents can replicate the likes of, say, Salesforce. What are your thoughts on this idea?

TM: Maybe I should first explain what vibe coding is. It's the idea that you can use plain English or something close to code so that you can instruct the model. Say, if I want a CRM for my business, I want to replicate Salesforce, and actually the AI model would go in and spit that out for you.

I think we would say vibe coding is massively overhyped today. It has some great use cases and some really cool stuff you can do with vibe coding, and you can create some interesting things. I think the extrapolation of that to actually building credible enterprise software is massively overhyped. I also think that maybe that narrative is already fading a bit, because I think people recognise that that's an overstretch. To maybe explain that a little bit better, I think, Andrej Karpathy… There's actually a recent podcast with him that he does with Dwarkesh (Patel), which is well worth the listen.

KB: We'll put it in the show notes.

TM: And he [Karpathy] led Tesla's autonomous driving efforts for quite a number of years. And he talks about "the march of the nines".

KB: Which is?

TM: Which is this idea that you can get to a product once you get 90% of good outcomes, you can do a demo. We had autonomous driving demos in the 90s. To get to 99% accuracy, it takes the same effort to get from 0 to 90 as to get from 90 to 99. And then every nine beyond that, so going to 99.9 actually takes the same effort that went in before. So it rises exponentially, the difficulty of every next nine that you go to. Now the problem with something like self-driving, and the same is true for enterprise software, is the cost of failure is incredibly high. It's very logical. The cost of failure in self-driving is self-evident or...

KB: Medical devices or…

TM: Yeah, exactly. And, in enterprise software, things like security are incredibly difficult. If you are running your whole business on a piece of enterprise software, and there is one vulnerability that allows a hacker to get into your system, it's a massive problem. There are many things that can happen and so you need an incredibly robust enterprise software. The problem with the march of the nines is it looks like you're very close very early, and then actually it might take another decade before you actually make reasonable progress. And that's what we saw in autonomous driving, so I think the same is true for enterprise software, and the hype is overplayed.

KB: Another question that often comes up with some of these software-as-a-service applications is monetisation. This is all great, everyone's doing great stuff, but are they actually making any money from it?

TM: You're starting to see that. Monetisation is definitely something the market is hyper focused on right now. What you've typically seen is that there's certain applications where you're seeing real monetisation very quickly – so coding agents have been a huge success case. Anthropic, Codex, Cursor and a number of others are really ramping up revenues. So, that's a great use case with real revenues from enterprise customers, but they're very early adopters because they're selling to coders.

Enterprise customers have been slow. So, generally, outside of Microsoft, which has done pretty well with some of their tools (Copilot etc) and some of the coding agents, enterprise has been slow to move. And so we're waiting for that inflection but, actually, when you can sell it like when a consumer can sign up for a subscription, like a ChatGPT type model, where people can do it individually or as part of small teams with an enterprise and have it on the side, that's actually been quite successful. You've already seen that pick up in legal settings and healthcare. So, you are definitely starting to see the ramp of monetisation, both in private companies but also public companies as well. But everyone is watching it so carefully.

KB: Yeah, because a company like Adobe is one that often comes in under the spotlight for this. And I think you met with them.

TM: Yeah, we had a really good meeting with Adobe. That's another example where the core of the business is enterprise, and so that's actually a comfortable place, or a good place to start, but they have to move fast and they have to innovate and they have to really use these tools. On the monetisation side, there's an advantage in deferral as well. If you can follow a freemium type model, where you offer a lot of what you're doing for free initially and really drive adoption, drive usage, gain the customer and then ramp monetisation later – that's a classic Adobe model. Now we clearly have to watch it very closely because it has challenges in the market.

KB: And I think we're probably running out of time. But before we go, you mentioned that you met some medical device companies whilst you were out in California – another obvious area of expertise out there. Could you give us a quick sense of who you met and what you were hearing?

TM: We met a range of companies that I mentioned earlier – Intuitive Surgical, Edwards Lifesciences, ResMed. From robotic surgery, heart valves to sleep.

KB: And the visit to Edwards' heart valve facility sounded really fascinating.

TM: Yeah, it was at the very end of the day when we went into the factory, and it's actually a surprisingly manual process. You're stitching these tiny, lifesaving heart valves. It's a long process, taking animal tissue, treating it through many various different stages to then stitching it very carefully into the correct shape. And this is done by hand, so we saw a whole host of people stitching – it takes roughly 1,200 stitches per valve. We saw people looking through microscopes doing this incredibly detailed work, so it's quite humbling to be honest.

KB: Now, before we wrap up, there's clearly no shortage of content on AI, and the story keeps evolving. As someone who spends a lot of time thinking, reading and listening to this topic, do you have any particular articles or podcasts that you'd recommend, beyond the one you referenced earlier?

TM: So yeah, definitely listen to the Dwarkesh podcast. He is doing a whole series of podcasts around that are looking to understand the nuance of AI development. I could recommend probably the last five or six podcasts he's done, because it's really been an interesting journey, the interviews that he's done. Also listen to Satya Nadella's last couple of interviews. He did one with Sam Altman after they closed the OpenAI deal, where they really firmed up that relationship. That's really interesting because it helps you understand Microsoft's very strong position, but it also helps you understand the thinking behind the strategy, so that was very revealing. I would start there but there are many more we could list as well.

KB: Tom, thank you very much for sharing your thoughts today. AI clearly remains one of the most fascinating and transformative forces shaping the world today, and I'm sure we will continue this conversation on future episodes. We will also include links to the podcast you mentioned in the show notes. And to our listeners as we head towards the end of the year, we hope you have a happy festive season and a prosperous year ahead. Thank you very much again. Goodbye.

This podcast transcript has been edited for clarity.

WALTER SCOTT

Important information

This article is provided for general information only and should not be construed as investment advice or a recommendation. This information does not represent and must not be construed as an offer or a solicitation of an offer to buy or sell securities, commodities and/or any other financial instruments or products. This document may not be used for the purpose of an offer or solicitation in any jurisdiction or in any circumstances in which such an offer or solicitation is unlawful or not authorised.

Stock examples

The information provided in this article relating to stock examples should not be considered a recommendation to buy or sell any particular security. Any examples discussed are given in the context of the theme being explored.

1.0_COM002666_12/12/2025

Walter Scott & Partners Limited, One Charlotte Square, Edinburgh EH2 4DR
Tel: +44 (0)131 225 1357 · Fax: +44 (0)131 225 7997 · www.walterscott.com