



THE INFLUENCERS: DIGITAL TRANSFORMATION

TRANSCRIPT Stefan Schiefer

<p>Leo von Gerlach</p>	<p>Hello, everybody, and welcome to another edition of The Influencers, our podcast conversation on digital transformation and law. I'm Leo von Gerlach and with me today is Stefan Schiefer. Stefan is CEO of Genesis Cloud. Genesis Cloud is a leading company in the emerging field of cloud infrastructure services for AI applications. Stefan is not only extremely technology savvy and business smart. He's also one of the most ardent promoters of the blessings that AI will bring for all of us. It is always most inspiring to speak with Stefan, and surely this conversation will be no exception. So before looking closer into your vision for Genesis Cloud, Stefan, and its place in business, a question to you and your personal story. With your background in semiconductor engineering, how did you end up driving Genesis Cloud? What brought you there?</p>
<p>Stefan Schiefer</p>	<p>Thanks Leo, for the warm welcome. Very appreciated. Actually, from the background. I'm a physicist, I studied semiconductor physics in Munich. And with my last startup that I built up, we were already building analytical equipment for the semiconductor industry. And we were selling to companies like ASML, for example, or TSMC. And those are the companies actually that actually build those high end GPU chips that are now needed for AI. So I was already in that industry already with my last startup. And I got to know the founders of Genesis Mining, founded in 2013, and one of the largest mining companies in the crypto space. And in crypto, you also need GPUs to mine. And that was actually their idea in 2018, to build a company that not only mines, but also uses GPUs for AI, machine learning workloads. At that time, generative AI was not a big deal. So it was a quite early vision from them to build such a company. And they called this company Genesis Cloud. And then in 2022, when generated AI, they really took off - we all know the Chat GPT moment. And then they change the strategy a bit from Genesis Cloud, it was on consumer GPUs before that. And then they change to those professional high end GPUs that are needed for those AI workloads. And that was the moment where I joined. So I know them personally. And it was that big strategic shift towards generating AI. And that's when we decided that I will become the CEO of Genesis Cloud. Also, because I had AI consulting background, I was consulting for more than two years AI startups, and I'm very familiar with the generative AI space.</p>
<p>Leo von Gerlach</p>	<p>So that brings us already close to the business model of Genesis Cloud. So you spoke about GPUs meaning graphic processor units. And there is obviously that difference between GPUs and those traditional process CPU central processing units. Perhaps you share with us what those specific GPUs mean, for Genesis Clouds business model, and why that is so tightly knit with one another.</p>
<p>Stefan Schiefer</p>	<p>So maybe a little bit technical background -- why GPUs and why not CPU. So we all know those traditional cloud services, everyone has an</p>

	<p>email account in a cloud, and all those clouds from Google, Microsoft, and so on AWS, they are all CPU centric clouds, that means all those services, all those applications that are running in this cloud, use a CPU, then you sort of type something into your app, you want to actually read your email, then the app sends a command to the database and that CPU based cloud and you get your data, you get your email. With AI, it's totally different. You have AI models and those AI models are very compute intensive. When an AI model is in production, it's called inference when you train a model is just AI training. And when you actually ask the AI something, then this is called inference. And this inference task is very compute intensive and GPUs are very simple processing units that can do one task, but they can do one task very fast. And they are work highly parallel. So that is the difference between a CPU and a GPU or CPU, so to say, a general manager that can do everything, but there's only one manager, a CPU is 1000s and 1000s of very small circuits that are very specialized, and but then they can be highly parallel. And this is what AI actually needs to be fast. So you have to build an AI data center on GPUs and not CPUs.</p>
Leo von Gerlach	<p>So I understand you really focus on the use of GPUs, because they are very good at matrix calculation. And matrix calculation is kind of at the heart what AI is doing. And when you now just think about better sharpen the business proposition of Genesis Cloud to optimize its infrastructure, ever further to improve these cloud services for AI applications, is it that this also hinges on improvements or enlargements of the processors, or what other things come into play?</p>
Stefan Schiefer	<p>Yeah, so it's two things actually, of course, you need the GPUs, because with CPUs, it's simply not enough compute there to do the AI models. But you also need the software back end, our customers, they develop AI models. And to develop AI models, you need specialized software, for example, ML Ops platform, which means machine learning operations software, this is more or less a software where you can put in your personal training data, or enterprise training data, then you can pick a certain model that you want to train. And then you can put that model also into production, those kinds of applications, we are integrating also into our backends and you can put certain AI application on top of that basis. And that back end actually then distributes all those necessary calculations to the 1000s and 1000s, of GPUs that you need to train and put a model into production. So we are also building up the software stack that makes AI models over time easier and easier to use, with the end goal that some employee and a company can use and program its own AI model.</p>
Leo von Gerlach	<p>So that's interesting. You spoke about the data that go into the models that are then onboard process that immediately leads to a question for a cloud service provider. How do you cater for the security for the protection of those data that go in? How do the firewalls, how do the rail guard look like?</p>
Stefan Schiefer	<p>Yeah, so security, of course, is a big issue in general in cloud business, so that that was already part of the CPU centric clouds. And with GPU, it's even further because AI is something that most enterprises will use in the future very intensively. So it has to be especially safe. So two things in the modern GPU cloud, like we use, and also our competitors. Encryption is really key. So we encrypt everything, this is one part. And then if you're an enterprise, there are certain models, for example, RAC models, so retrieval augmented models means that the data that is</p>

	<p>precious, so the enterprise data remains within the enterprise, and the AI model is in the cloud. So that means the AI model, you can actually look like a reasoning engine. So the reasoning engine is outside of the enterprise, and the precious enterprise data like customer contracts, whatever, whatever it is, is in a database inside of the enterprise. And if the employee asks some questions, so can you please check if this incident is according to some customer contract, for example, then the reasoning engine outside of the enterprise in the cloud in our cloud, processes this data anonymously, and then sends it back to the retrieval model and the retrieval model inside of the enterprise loads the data from the enterprise database, and actually present that data to the employee that has prompted this question. So in that case, the data is as safe as the enterprise makes their databases safe. And there is actually no really valuable information sent over to the cloud. So that security is not a big issue there. But there are also other models, of course, where the data has to be put in the cloud, because the amount of data is so large that you want to train the model with that data. And they are we have, of course, industry standards, security measures like encryption, ISO certifications, for example, there is the ISO 2701, which is security, ISO certification, all our data center have that our back end technology is also done in a way that it's very secure.</p>
Leo von Gerlach	<p>So that's very reassuring for your customers to hear that there is kind of a minimal data flow going into the system. Just leading me to the next question about scaling. With all we read about the enormous demand and processing power for AI applications, these comments about investments into data centers into cloud service provision, how do you make sure that Genesis Cloud keeps pace with this growing demand for computing and cloud services in the AI space?</p>
Stefan Schiefer	<p>Yes, that's absolutely true. The pace of how fast this infrastructure scales is really amazing. I think I've never seen such a speed before. And there are always bottlenecks that you need to take care for. In the beginning, we have bottlenecks, for example, with the high end NVIDIA GPUs, because they are limited through TSMC is production capacity. As we started with Genesis Mining in 2013, we have very, very good supplier relationships with Nvidia. And we are actually favored in their supply chain. So it's really, really key to have good supplier relationships with those companies that make those precious GPUs. And then also network chips were a big issue after that. So it's even harder to get network chips, because all those GPUs are connected with a very fast network. And those chips also have become a bottleneck over time. And we also have very good supplier relationships with those companies that produce those network chips. And nowadays, it's even the transformers that are used in data centers. So they also were a bottleneck. Elon Musk phrased it very nicely. He said, You need transformers to train transformers. So the transformers in data centers, they actually take care of that the high voltage, the scale down to the very low voltage that those chips need. And those were also bottlenecks. So now everyone is fighting actually, for data center capacity. And yeah, with Genesis Mining, we are very familiar how to build data centers, we have very long lasting relationships with data center providers. And yeah, this has become a bottleneck. And ultimately, there is a discussion right now ongoing if the electricity is enough to build up AI. And it goes even that far that one of our competitors, thinks about building his own nuclear power plant to actually make sure that they have enough electricity for their data centers to deploy enough GPUs. So we're talking here about gigawatts. And,</p>

	<p>yeah, it's crazy. I mean, if you see the, the investment from Microsoft, last days, I think they announced \$100 billion investment in the project Stargate, or general in GPUs, or AI infrastructure. It is a race that we've never seen before about resources, capacities, data center electricity, but I think Genesis is a very good positioned here, because of their history and mining.</p>
Leo von Gerlach	<p>I mean, all the more important just to understand that you have those very good supplier relationships, which can allows you to maneuver around those bottlenecks as best as you can in order to get just capacity and computing power. Which leads me to a slightly different field, namely, just as we see these advancements in semiconductor technology, there are so many other advancements just new models, new applications, new devices, new combinations. What do you see as those most relevant trends for Genesis Cloud where you say, we look closer into them, we may invest into them, we follow up on them.</p>
Stefan Schiefer	<p>I just came back from the GTC which is the most important AI conference that is hosted by Nvidia in Silicon Valley. You of course have to keep very close relationships with suppliers like Nvidia, Intel and AMD. Those are the ones that produce the GPUs. And this is the basis of AI. All AI's are built on those GPUs. So you have to learn understand what is their roadmap? What are their next generations of chips look like? What can you do with them? How can you deploy them in the data center. And then, of course, also our software back end. I mean, we've been talking about those additional kind of managed services. So we keep also constant and regular contact, our customer success department has constant contact with our customers. And those customers are developing the next generation AI applications. And we need to understand what are their needs, and they are actually telling us which managed services we are going to integrate next. So there is a constant development on the on the hardware side on the GPU side. And there's a constant development actually on the software back end, where you need to integrate useful services so that AI systems can be developed faster and be more powerful and more helpful.</p>
Leo von Gerlach	<p>Taking this one step further, are there any business relationships, cooperations perhaps even beyond the sphere of hardware suppliers, software providers, where you say they are particularly relevant to accomplish Genesis Cloud services?</p>
Stefan Schiefer	<p>So for example, we are teaming up with all kinds of specialized AI companies, for example, ML Ops platforms, we are collaborating with clear ML, then you also need specialized databases for AI, you're not storing in a classical database, you have to vectorize your data. So everyone has heard probably heard meanwhile, the term of a token, a token is, let's say, the bit of AI. And to store and process all those AI's, you need to store in a vector database that is very beneficial. So we are also partnering with, with companies that have very fast vector storage databases. And of course, we also like to collaborate with academia. I mean, there's also a lot of projects going on and research going on in academia. And we also would like to hear their opinion, what they are working on, because this, this might be the next big model that's coming out, and you need to prepare for that as well. So those are more or less the relevant players, we're in constant contact.</p>
Leo von Gerlach	<p>And then you're ideally located in Munich with the Technical University that place and put so much emphasis now on artificial intelligence, and</p>

	<p>everything that comes along with it. Perhaps going even one step further beyond Genesis Cloud and the impact, the most immediate impacts of AI generative AI on society, I'm not here to speak with you about doomsday scenario, you would not be the right person, because like myself, you, I think you strongly believe in the benefits that the new technology can bring, perhaps your idea, your vision, how you see this playing out in the short term.</p>
Stefan Schiefer	<p>Yeah, I'm quite optimistic actually, about AI. Of course, there are always those doomsday guys, right, whenever there is a great technology coming along, you have always both sides, right? What I think that the immediate near term effects with AI for enterprise customers is that you probably get rid of all those boring work that that employees anyway don't want to do, right? If you if you look at law companies, and you get a new case, and you have to read 1000s of pages, just to extract some information out there, which is relevant for you. You can either ask like a couple of lawyers to do that for weeks, or you just ask the AI, and it does it for you. And you get that output immediately. I think it doesn't make on the near term. At least it doesn't make us all redundant. It just makes us more effective. So I think everyone that embraces AI and collaborates with the AI will just be a 10x more productive, right? You will just be faster in what you've already done before. Because you get rid of those long term boring tasks that anyway, no one wants to do. Right. So I think that is quite interesting. And I see a more productive working environment and a more fun and working environment. Actually, if you collaborate with AI in the right way.</p>
Leo von Gerlach	<p>Stefan that has been most exciting speaking with you. And that was certainly no disappointment to my great expectation to this conversation in this interview. So thank you so much for giving us the opportunity to listen to this just very insightful remarks about you and your company. It's very encouraging.</p>
Stefan Schiefer	<p>Thank you, Leo, for having me and goodbye everyone.</p>
Leo von Gerlach	<p>Thank you all for tuning in for listening and I hope you'll join us again for the next session of the Influencers which will be coming up soon. For now, take care goodbye everybody.</p>
	<p>Visit our digital assets and blockchain hub at engage premium Hogan lovells.com for more podcasts and other resources.</p>