Secure your AI, GenAI and LLMs with Mindgard

As part of our quick founder questions series – or QFQs – we spoke to Peter Garraghan, CEO & Co-founder of Mindgard about adapting to the challenges in security within AI, investment and demystifying AI security.

Temps de lecture : minute

26 March 2024

Back in 2016, when I was leading my research lab, I started investigating the cybersecurity problems within AI and machine learning systems. It struck me how big of a challenge this was going to be, imagining AI and ML becoming as common as operating systems or virtual machines. They're everywhere, right? So, I postulated, if AI is going to be behind everything, making life better and all, it's also going to attract trouble from all corners – be it from countries, groups, or individuals. The more AI becomes integrated, the more risks we would be facing. And the cybersecurity tools we had? It is possible that they would severely struggle within that future.

I spent the next four years diving into research with a group of brilliant scientists and engineers I was lucky to work with. We had this big data centre within my lab where we were able to investigate some deep problems. It turns out, we were spot on about the challenges in security within AI. And then it occurred to us – we needed to build something, a product, to tackle this. But to get there, we needed solid tech and all that research we were doing. It's really blown up since then, in a good way. And we are sticking to one of our core principles that researching hard problems is critical in keeping us competitive.

How has the business evolved since its launch?

About two or three years ago, after making some real progress on the research side, we were in a position to say, "Okay, we're going to build a company." *Mindgard* was formed originally back in May 2022. I laid all the groundwork to raise the original capital, get the proposition in place, get the tech through the demos, etc. I basically did two full-time jobs. This entailed both operating a leading UK research lab and teaching the next generation of computer scientists, while getting a lot of the startup groundwork done over the next two years. I did this because I believe that we really need to get our technology into the world ahead of growing threats against AI. I had to come out to London non-stop from Lancaster.

Tell me about the business - what it is, what it aims to achieve, who you work with, how you reach customers and so on?

On one side, we have the decision-makers and on the other, the technical experts. The technical individuals, including data scientists and cybersecurity professionals, often lack a deep understanding of AI security. Data scientists focus on creating and applying machine learning, not necessarily with an eye on security, while cybersecurity folks might feel out of their depth with AI, not fully grasping its implications. We often find ourselves helping to guide them in understanding the significance and specifics of AI security.

A year ago, our conversations might have seemed more “future proofing” to them, intriguing but not urgent. Six months down the line, there was a shift towards acknowledging the need to gear up for AI's security challenges. Now, there's an immediate demand for our input as projects are about to go live. We explain to tech specialists that despite the

novelty of AI and LLMs, the principles of security management—like threat detection, data protection, and risk assessment—still apply. They just need to be adapted to the AI context.

Tell us about the working culture at Mindgard

I'd say it's really exciting. Everyone in the company is fully aware that we are working on unsolved problems with world-defining consequences. This motivates us in everything we do. When we create never-seen-before technology, helping businesses understand the cyber risks within AI, as well as supporting and mentoring each other. Everyone is passionate about the greater calling of Mindgard to secure AI on a global scale, and I am immensely proud of all my colleagues in what we have achieved so far, and will continue to do so in the future.

How are you funded?

As a Lancaster University spin-off, we were originally funding a lot of our initial research from grants and the like. We continue to keep a good presence of our activities within the academic realm, due to the huge amount of innovation.

On the business side, *we raised capital back in August 2023, and that was from IQ Capital and Lakestar*. Together, we announced a seed round back in September of £3M.

The aim for the near future is to continue to rapidly grow our company, get our product out there, and put this solution in the hands of as many people as possible to address real AI security challenges. Essentially, it's about expanding our focused team into a much larger entity. We're currently building a solid foundation; we just need more resources to broaden our reach by doubling down on product development and commercial deployment.

Many companies are on a similar journey, trying to come to terms with deploying AI, LLMs and GenAI. We anticipate that our fundraising efforts will align well with this shift, positioning us just ahead of the curve.

What has been your biggest challenge so far and how have you overcome this?

Many businesses are unaware of the cyber risks associated with AI. It is extremely difficult for non-specialists to understand how AI actually works, much less what are the security implications to their business. I spend a considerable amount of time demystifying AI security into a clear proposition. At the end of the day, AI is still essentially software and data running on hardware.

You don’t need to throw out your existing cyber security processes, playbooks, and tooling, you just need to update it or re-armor it for AI/GenAI/LLMs. This thinking has in recent months rapidly changed with the rise of GenerativeAI and LLMs, government attention, and working groups such as OWASP AI exchange that are really helping to actualize and ground cyber security risks within AI.

How does Mindgard answer an unmet need?

To date there is no major security platform for AI/GenAI/LLMs. This represents a fundamental problem for businesses looking to adopt AI (whether internally built or through a third-party app), as they are forced to do everything security-related by hand. They are probably very quickly realising that they don’t have the people or skills that can keep pace within the rapidly changing field. Mindgard answers this by providing a full-stack solution in AI security, including automated pen testing, threat detection, response, and Data Loss Prevention (DLP) so that businesses can use AI safely and securely.

What’s in store for the future?

Two things: The first is that the AI field will continue to innovate, allowing us to create things that appear to even technically minded folks as borderline sci-fi. The second is that as the adoption of AI accelerates over the coming years, cyber attacks against AI will in turn only grow in number and damage. This is the exact reason why Mindgard was formed, and is well positioned to help people to get ahead of the AI tidal wave securely.

What one piece of advice would you give other founders or future founders?

It’s a long road, so I would kindly suggest having your motivations in order. I still want to change the world. That's why I built this company. The whole point is to get this tech out there to actually solve real issues in security for AI and machine learning.

Generally, I've come out of my job as a professor, which I love, and still practise. It's a very similar skill set, and a similar motivation in that there is a greater purpose in what I do – I'm here to solve problems and define the future, because I can see over the next few years, whereby a huge amount of AI security issues are regular occurring. AI changes so quickly every week. To keep up and keep pace with this will be a huge undertaking. And if anyone's going to do it, it's going to be us. We live and breathe this stuff.

And finally, a more personal question! What’s your daily routine and the rules you’re living

by at the moment?

As CEO and CTO, my routine is quite varied in terms of day to day activities. I also still retain my professor position at Lancaster to develop the next generation of AI security researchers within the UK. Hence you can imagine (like all startups), we are kept rather busy! A guiding principle that I have stumbled upon very early in my career was to "identify and focus on the things that matter most" – whether that is day to day activities, or the strategic direction of the company.

Peter Garraghan is the CEO & Co-founder of Mindgard.

Article by Peter Garraghan