The Cloudcast (00:01.526)
And we're back and we have a topic today that actually we've been meaning to talk about for a little bit now. And actually Brian talked about on a Sunday perspectives. I'll put a link to the show in the show notes previously, but we're going be talking about shadow AI and the implications of it today. And for that we have Rohan Sathe, co-founder and CEO at Nightfall AI. So first of all, welcome to the show Rohan and give everyone a brief introduction and especially about your time at

Uber Eats and kind of founding engineering at Uber Eats. Tell us a little bit about all.

Rohan Sathe (00:35.897)
Yeah, awesome. Thanks. Thanks for having me. I'm happy to share. Yeah, the Uber Eats story was interesting. When I joined Uber, I knew we'd have kind of the opportunity to build a business within a business. So at that point, it was still very much a small company and we had built out this kind of ride sharing platform. And we thought of that more as a logistics platform where we weren't just delivering people to places, but we could also deliver other things like food, right? And so

I got a great opportunity to kind of start that Uber Eats business with a few engineering counterparts and the rest is kind of history. It's a pretty massive business now and we scaled that actually close to 10 billion in bookings before I left in just about two years. So phenomenal kind of ride.

The Cloudcast (01:18.83)
Absolutely. Those are, those are definitely the kind of rides that you see. There's, they're very hectic, but very fun kind of all at the same time. So let's talk about, um, shadow AI. Let's kind of start at the beginning here. How would you define shadow AI?

Rohan Sathe (01:25.615)
Yes.

Rohan Sathe (01:37.808)
Yeah, so the way I define shadow AI is, you of course there's been this rapid push from CEOs or boards amongst their kind of employee base to use and leverage as much AI as possible. We see that, we see all the benefits with of course things like ChatGPT already, but I think businesses feel like they can be a lot more productive, right? And so shadow AI is just this concept that

your employees might be spinning up different AI tools or ways to use AI and the organization may not know about that. And so what is the implication of that? Well, there's a security, both a security and a compliance implication. The security implication is, hey, somebody's taking my corporate data potentially and then sharing it with an AI application and who knows, are they training with that data and improving their AI models with our corporate data? And that's not something that we've allowed them to do, right? And the compliance version of that also is like,

certain data just can't be shared with anyone, right? So if you're a healthcare company, you deal with protected health information, you can't just willy nilly share your customers like social security numbers and stuff like that, right? So that's just the general concept of Shadow AI.

The Cloudcast (02:49.646)
Yeah, yeah. And maybe kind of a follow on on how I've been thinking about this because you hear shadow AI and shadow AI is based off of the original shadow IT term back in the kind of the start of cloud, But I feel like this looks a little different. And here's why I think that because the learning curve is much smaller. And what I mean is like for shadow IT to happen, you had to know IT. You had to know servers and you had to know storage and

And with shadow AI, mean, it's kind of everywhere. And the reason why it's everywhere is because it's literally like, hey, if you know how to type, you can go use shadow AI, right? And so like, is that a correct way to think about the problem? It's like shadow IT on steroids.

Rohan Sathe (03:28.475)
Yeah.

Rohan Sathe (03:33.999)
Yeah, I think so. think it's very similar. Perhaps the only difference is, as you mentioned, the adoption curve. the time for somebody to adopt an AI application is probably a lot smaller than it is just any generic IT application, right? So I think that's been one big piece. The second is there's an extreme push from CEOs or boards of companies for the employee base to use AI applications. So that push, though,

perhaps strong during the era of cloud migrations and stuff like that. It's probably not as strong as the push to using AI. And then the third is it's not just a browser or a desktop-based mechanism. I can spin up agents myself. Now I can click a couple of buttons. Now I've got these agents running. And then it's just programs doing things. And this is not happening with much user interactivity once you've set something up. So I think that's

You know, those are the big differences, I would say.

The Cloudcast (04:35.768)
Yeah, yeah, makes sense, it makes sense. And let me ask you this then too, because maybe let's quantify it a little bit, because it sounds scary, but how big is the shadow AI problem today in your view? I've seen some sometimes alarming stats, but you also don't know, okay, how much of those stats are click-baity headlines or whatever, right? But considering JNI is only a few years old.

I feel like it's probably a considerable portion of like if we go talk about data leaks, which is what we're going to be talking about here in a little while. Like when it comes to data loss prevention,

DLP, right? Did it just skyrocket up the charts when it came to the, you know, how much leakage was potentially happening and how big of a problem this became quickly?

Rohan Sathe (05:22.352)
Yeah, I mean, think just if you just look at the growth rates of even chat GPT back in 2024, right? I mean, they were seeing like consumer level kind of growth rates for like hot companies like when Facebook had launched or even Uber had launched. Right. And so I think because of that, because of all the consumer mind share that's, you know, easily spread to the enterprise. Right. So I do think the problem is very pervasive. And we see that

because as a data loss prevention company, like we see the hard data around usage of AI applications in an enterprise. And so that also has dramatically skyrocketed. And I would say at this point, you probably have 80 plus percent of the employee base using some form of AI. And so the question is, is like, does the enterprise have good control on it?

The Cloudcast (06:14.126)
Yeah, yeah. So let's talk about that for a little while because like we mentioned DLP. So data loss prevention, you know, hardware, and systems. And I will say this, you know, like I've had some history behind it with some like some of the more traditional systems back in the day. And I mean, here's the thing in my experience, those tools were super cumbersome, super clunky and like.

Rohan Sathe (06:38.329)
Yeah.

The Cloudcast (06:39.286)
You had this classic trade off of like, Hey, we're going to lock down the user. Like IT is going to like, is basically was like this whole idea of like, user productivity versus security. And we're just going to lock you down really, really hard. And twofold question, like how has that mindset changed in the era of AI, but also to your company was using AI prior to open AI and gen AI, right? And so like,

Rohan Sathe (07:06.501)
Yeah.

The Cloudcast (07:08.814)
Almost maybe like a twofold question. Like what was your concept back in, I think it was 2018, 2019 when you were thinking about starting the company and then did that mindset change when everything kind of hit in 2022 and 2023?

Rohan Sathe (07:24.228)
Yeah, for sure. And we've had to stay abreast of all of the rapid changes in AI. And we continue to do so. But just some of the backstory of DLP. the problems kind of, or DLP really originates from this idea at Bank of America, I think is like the classic example where you had internal

employees stealing credit card numbers of their customers and doing nefarious things with those.

Back then, the use case was, know, everything is in our own data centers. So like, let's deploy some software on our servers or on our employees workstations monitor if they're taking cardholder information and leaving with it. Right. And there are only a few ways to leave with data. At that point, it was like, I could stick a thumb drive in my workstation. I could use a browser maybe, right. Or I could email it to myself. so a couple of things have changed. So one is.

Back then, the detection of what was sensitive was very limited to rules. So like I would write a rule that says a credit card number is like 16 to 19 digits. It always has these certain prefixes that passes some algorithm when you run a checksum on it. And that's how you identified what a credit card number was, right? Unfortunately, the problem with these rules-based kind of methods is that they're super noisy because it turns out

If you're just scanning anything that's 16 digits and has some prefixes, like a lot of stuff actually just passes that and is not a credit card number, right? And so DLP had gone a very bad rap of being super, super noisy because it was just rules-based. The big transformation that happened and no pun intended was just the Google kind of transformer paper that talked about sort of new techniques towards solving natural language problems, right? And so that's kind of the idea that we leveraged to say, hey, you can

get a lot better about classifying data using NLP. And the reason why my experience is relevant there was even back in 2015, like Uber was at the cutting edge of leveraging machine learning and doing that at scale. And so this was a great opportunity for me to leverage that expertise to solve a different problem in a different domain. You asked a question about how have things changed with LLMs. The only difference has been the technology has gotten even better. So like we had kind of transformer-based architectures,

Rohan Sathe (09:44.748)
that were still very good, but not as good as large language models. And so that's the only thing that is adapted is like, OK, now we can leverage something even more sophisticated in the way we do detection. So that's kind of been the biggest thing.

The Cloudcast (09:58.486)
Yeah, yeah. And I know I've talked about this a number of times recently on the show, but I kind of see a lot of these things we've been talking about. I see two sides, right? It becomes this cat and mouse game of like, okay, there's the on the front side, there's the users. then like, yeah, like you were saying, an employee goes and signs up for chat GPT, uploads a bunch of documents and then does it, you know, what they think it's like, hey, I'm just going to use this to, you know, make my job easier, make it faster, make this, you know, document better.

Rohan Sathe (10:27.192)

Yeah.

The Cloudcast (10:28.406)
by the way, that's really sensitive data that they shouldn't have been uploading. But then on the back end, you're using AI to kind of, you know, this large amount of detecting signals from noise and the bigger the environment, the more noise. And so you kind of have this like one versus the other and one side leading to potential losses, the other side trying to prevent it. And is it as simple as it's a cat and mouse game or?

And also like what happens when like AI powered attacks meet AI powered defense, right? Like, like at what point does it just kind of get like AI everywhere and it's just the AI is fighting each other kind of thing. Like tell me a little bit about like the thought approach, if you will, to all of this, right?

Rohan Sathe (11:02.585)
Hehe, yeah.

Rohan Sathe (11:12.527)
Yeah, yeah, for sure. And yeah, I mean, I think in general security, even outside of data loss prevention, like even just like antivirus or something is very much a cat and mouse type of thing. That's just a general kind of security challenge. So you have to stay kind of up to date with newer threats, And it adapts software to accommodate for preventing those kinds of things. And so, yeah, I think that, you know, luckily with data loss prevention,

Most of the incidents are actually just poor hygiene. So it's not like the user who's using chat GPT is like deliberately trying to share sensitive data to get past a control. They just don't realize the implication of sharing your financials with open AI or whoever. Right. And so it's more of an education thing ultimately for that use case where it's just like poor hygiene.

And it's like, how do we just make sure employees operate with the right hygiene and educate them to the right kind of places to use AI? But you're right. Like there's this kind of meta idea of like now, Nightfall is an AI product. So we're kind of fighting people from using AI and we're also going to apply that to AI agents. So then we're really an AI kind of fighting an AI, which is, which is very meta. But yes, I think that's, that's the world we're going to.

The Cloudcast (12:34.254)
That's right. That's right. So let's dig into kind of the more practical aspects of all of this then. So let's kind of do a maybe a bit of architecture on all of this. Like how does this work across say like SAS points and endpoints and browsers and AI tools? like I know like again, going back to my old old DLP days, like, hey, there'd be like.

Rohan Sathe (12:52.515)
Yeah.

The Cloudcast (12:58.99)
a client that ran on your desktop and like if you tried to plug something in it blocked it or if you tried to email something it like flagged it or like how are you tracking things increasingly when you know a lot of things are sass based a lot of things are browser based you know it's not like you have you know my laptop running necessarily a lot of apps anymore I'm not I'm doing just about everything in a browser and so like how do you

Rohan Sathe (13:01.667)
Yeah. Yeah.

Rohan Sathe (13:14.83)
Yeah.

Rohan Sathe (13:23.736)
Right.

The Cloudcast (13:26.146)
begin to detect things, how do you remediate things? Explain to everyone how the pieces all fit together.

Rohan Sathe (13:32.688)
Sure, yeah. So at the core of it, there's this idea of exfiltration vectors. So these are the places where data can leave the environment, right? So the things that traditional DLP had kind of focused on were things like USBs or somebody printing sensitive data. I who really does that anymore? But those would be kind of the ways of the past.

Now, as you mentioned, people are primarily using SAS applications. They're using their browsers, they're emailing stuff through email clients. So the idea is we need to make sure we have coverage in the most common places. Folks may be exfiltrating sensitive data. And to do that, we need to insert ourselves both at the SAS application layer so we have API-based integrations to common SAS applications that

an organization might use, like Office 365, Slack, Google Drive, stuff like that. But we also need to sit on an employee's workstation. So we've got software running on the workstation. And that's what's monitoring for any data exfiltration. So the places we have controls on are SaaS applications and endpoint devices. And that gives us full visibility into what's going into what's happening in the SaaS application, what's happening on somebody's workstation.

The Cloudcast (14:51.566)
And this may be an antiquated question, but from a architecture standpoint, then where does the AI and AI processing piece fit in? it in real time analyzing everything and it's flagging things? is it more like, I go back to my old days of like, is it signature based and you got to update the signatures on the end points? like, how does it know?

Rohan Sathe (15:08.676)
Yeah.

Rohan Sathe (15:18.168)
No, yeah.

The Cloudcast (15:21.39)
Like, and because it's not, or at least I don't know this, I'll ask the question. Like, is it a, like a small LM running locally or like, how do you handle where things run and get processed?

Rohan Sathe (15:30.713)
Yeah.

Rohan Sathe (15:34.371)
Yeah. So most of our processing, have pre-processing and then after pre-processing, just the rest of the processing, right? So most of the general processing happens in our cloud. So we're a cloud-based SaaS service, right? So we are hosting models in our own cloud infrastructure. And that's where we're doing inference on, hey, we saw this metadata, like this looks like it's sensitive or this looks like it is exfiltration and not just a standard business workflow.

For example, most of that is happening in our cloud. We do do some preprocessing. So some of that preprocessing is in our cloud, as well as on the endpoints themselves. So we may say, hey, we preprocess some data. And we actually don't need to send it up to our cloud, because we know it's either just a business workflow or something that is not sensitive.

The Cloudcast (16:24.994)
That's fantastic. Now, let's kind of flip over and kind of maybe talk a little business and culture for a second. I know that you've mentioned in the past, like shadow AI is just a technical problem. It's a behavioral and organizational challenge as well. And I really like that. So like, what's your advice for...

CISOs, CIOs, security leaders, CDOs, there's so many C-levels acronyms out there, it's an alphabet soup these days. when folks are trying to manage this but also strike that fine balance between user productivity and security.

Rohan Sathe (17:11.362)
Yeah, mean, the first step is I think you can really be in kind of a monitoring only mode, right? So firstly, you just need to establish a program. Like are there a set of AI applications that are authorized by the organization that we must direct people to? Or do we have the stance that people can use whatever as long as we have some visibility and are able to have some sort of agreements with those companies that they can't share, that they can't train on our data, right?

So that's kind of step one is like, what is the program? What are the AI tools in use? And to validate that, you you can use our software and shameless plug, or you can use other tools. But

I think the idea is you want to get some sense of like, what is the landscape of AI applications that are in use today, right? And based on the policy that you have internally and what the data is showing, then you can start to put in controls, both in terms of preventing people from using certain applications or at least

monitoring when they're sharing sensitive data to those applications and then coaching them on, you know, hey, this is redirecting them to the right places to use AI or, like you can share data, just don't share this type of data and here's like our policy on that.

The Cloudcast (18:27.768)
Sure, sure. Well, and let me kind of a random bring it back around to something you said earlier in the show that you said something about like, okay, know, shadow AI can potentially get pretty bad in the agentic realm as well because now it's not just you doing things, but agents doing things on your behalf. so as we kind of, the industry kind of, I won't say shifts, but...

Adopts in addition to gen AI this whole idea of a gentic AI like how does that impact? The thinking and how does that impact the the consumption patterns? as you go from kind of one technology to the other because a gentic AI almost seems like it takes the problem and makes it exponentially worse at times or could has the potential to you, right? So how do you think about that Ron?

Rohan Sathe (19:14.157)
Nah.

Rohan Sathe (19:17.878)
Yeah, no, that's exactly right. think the ease of spinning up agents too, and just the distribution capability has made the problem even more pervasive. So an example of this would just be like, you know, we had a customer, for example, that was in healthcare, they're recording all of their customer support calls. Those are getting transcribed. So they had an agent that was taking the transcription and then running that against some sort of LLM to do some analysis. Now all of a sudden you've taken

A very sensitive customer interaction that contains sensitive PII or PHI and fed it to a model. And the employee that set this up was just trying to do something good. Like they weren't doing something malicious, right? They just wanted to set up a workflow that they were probably running manually and just improve that through automation. And so the fact that somebody can easily do this where they don't really have to even know how to code because they're just clicking buttons just means that data can.

move from place to place so that you even realizing. it's not as simple as just like me normally having some human interaction with my browser or something like that. It's just like I click some buttons and then magic is happening kind of in somebody else's cloud, right? So I think that's kind of the new surface area that just needs focus and attention and coverage. And Nightfall can help with that, but I'm sure others can as well.

The Cloudcast (20:45.346)
Yeah, and help me out a little bit too with it because like, what is the most common, agentic platforms that you're seeing the most adoption with right now? And, and what are some of the kind of the use cases that you're seeing that you're like, Hey, this is acceptable or this is bad. Like, tell me a little bit more about how folks are, using it and then you're potentially allowing it or preventing it.

Rohan Sathe (20:57.196)
Yeah.

Rohan Sathe (21:10.103)
Am I going to get some sponsorship money for the plugs here? No, I'm just kidding. But I would say I think the ones we're seeing is we're seeing N8N for sure. OpenAI has this bot killer, bot builder agent kit. So those are the ones that I primarily see. mean, at this point, we can probably assume that any foundational model company is going to have some form or flavor of just creating agents super easily. So I would expect.

The Cloudcast (21:12.75)
There you go, there you go.

Rohan Sathe (21:38.967)
the classic names, anthropic or whatever, to have something that folks should be privy to.

The Cloudcast (21:47.502)
Makes sense, makes sense. No, I appreciate it. final question to kind of close this out. If anyone out there is interested, what's the best way they could get started? Maybe they want to reach out to you or if you're going to be speaking anywhere or tell me a little bit about what you got coming up or if anyone wants to get hold of you or dig into Nightfall more.

Rohan Sathe (22:02.381)
Yeah.

Rohan Sathe (22:08.109)
There's a lot of stuff we have on the product roadmap that's going to be coming out here shortly, kind of increasing some of the coverage and capabilities that we have with Shadow AI. But folks can reach out to me at just firstnamerohanatnightfall.ai or just go on our website, ask for a demo, and our team's on that pretty quickly. So there's a couple of ways to get to us.

The Cloudcast (22:30.158)
Fantastic. Very much appreciate your time this week, Rohan. And for everyone out there, thank you for listening. We certainly appreciate you listening as well. And if you enjoy the show, please leave us a review if you can wherever you get your podcasts. And please tell a friend as well. as always, we're looking for feedback. Show at thecloudcast.net. On behalf of Brian, who wasn't

able to make it this week, and myself, everyone out there, thank you very much for listening. And we will talk to everyone next week.

Rohan Sathe (22:57.838)
Thanks.