Getting Started with Agentic AI: Discovery & Governance

ProofID Fireside Chat with SailPoint: Transcript

Speaker	Role	Organisation
Tom Eggleston	CEO (host)	**ProofID**
Matt Bruun	VP EMEA, Agentic Technology Specialists	**SailPoint**

About this transcript

This is the full transcript of ProofID's fireside chat 'Getting Started with Agentic AI: Discovery & Governance'. Tom Eggleston, CEO at ProofID, is joined by Matt Bruun of SailPoint to talk through where organisations should start with agentic AI security. The conversation covers continuous discovery of agents and non-human identities, shadow AI, governing agents with the same rigour as human identities, entitlement radius and agent classification, SailPoint's acquisition of Entro, the four teams that need to work together, and where customers should aim to be within 12 months.

Use it to catch up on the conversation instead of watching the recording, or feed it to an AI assistant to create your own summary, explore a specific topic or pull out the insights most relevant to you. Every turn is timestamped and attributed to its speaker, so you can trace any point straight back to the original video.

00:00:04 **Tom Eggleston** · CEO, ProofID

Welcome, and thank you very much for joining us for this really special chat that we've got today, where we're going to have a bit of a spotlight on SailPoint and how their technologies can help you with agentic AI security. I'm really delighted to have Matt with us today from SailPoint. Matt, would you introduce yourself?

00:00:24 **Matt Bruun** · VP EMEA, Agentic Technology Specialists, SailPoint

I'm Matt Bruun, the VP for Agentic Technologies at SailPoint.

00:00:27 **Tom Eggleston**

And what does that mean?

00:00:30 **Matt Bruun**

So SailPoint, around about a year ago, decided that things were moving so fast in the marketplace that we needed to create a specialist team in the organisation that could really take the some of the technologies we already had, some of the technology we were developing and were going to acquire in order to help customers really embrace what's going on with non-human identity and agentic AI.

So very different group of people that are looking at this technology. Different personas, different buyers and so on. So required a different sort of go-to-market. So yeah, we set that up about a year ago. And now I run that on behalf of SailPoint for Europe, Middle East and Africa.

00:01:11 **Tom Eggleston**

Well brilliant. And thank you very much for being with us today and sharing some of your expertise and insights. I think it's really interesting what you just said there about how SailPoint recognised things were moving so quickly and had to kind of change what they were doing. Of course, that's absolutely true for all of our kind of mutual customers as well.

The whole world is moving really, really quickly with agentic AI security. And I think one of the challenges for, you know, for customers and professionals trying to make sense of this is actually, you know, what does it all mean, where to get started. You know, and what they can actually take action on quickly.

And I think that's really what we're hoping to get into today is to really understand, you know, what actions people can take. And then specifically, you know, how can SailPoint help with that. So getting into it, when I think about agentic AI security, there's a model in my head that we like to use.

And you know, I'm sure you probably agree with this, but it always starts with discovery for me. Because, you know, I come back to this idea, if you don't know about something, you can't secure it. Really. And, you know, so that's probably a good intro point for our conversation. So maybe you could speak a little bit about, you know, discovery.

What that means to you within SailPoint. But what I'm actually really interested in. What do customers actually find when they do discovery.

00:02:26 **Matt Bruun**

Yeah. Yeah, I think so. You're absolutely right to have to start with discovery. But not that discovery then stops as you go through this discovery then becomes continuous as well. But when people first start, they have got to start tapping into. But when it comes to non-humans and agents, there are so many more of them than there are humans in an organisation.

You know, it could be 50 to 1, it could be 100 to 1. And they live and lurk in different places that people aren't used to finding them. So you have to have the ability to detect and find all of those wherever they

may be. So there are a number of different ways of doing that. Number one is to start tapping into the sort of the cloud and SaaS apps that people use to generate agents and which are then often used to leverage the non-human identities that are there.

Then looking at into some of the foundries like Bedrock and Copilot and so on, where people are creating apps within their organisation. And then really, what we've done recently is to make sure that we don't miss something which is going on all the time, and that's to tap into information that's already there in the pipeline area where a lot of agents, not just the sort of non-human identities are being created, but that's where they are self-identifying.

So there are characteristics about them which mean you can look at, well, how are these non-human identities working with each other. What are the kind of the secrets, the tokens, the credentials that they're using? And how is that being used within the pipeline? Because that's our key way of finding things that perhaps are not as easily discoverable from the first two categories.

That's a big difference from what we were able to do even just a few months ago.

00:04:13 **Tom Eggleston**

So it's moving, as you said, at the top, really discovery from being like a one-time thing to actually part of your kind of, you know, your normal business as usual process, continual discovery.

00:04:23 **Matt Bruun**

Yes.

00:04:23 **Tom Eggleston**

Yeah. And there's another really interesting side of this as well. I think, you know, it's a different aspect of discovery, really, is this idea of shadow AI. Yeah. Because it's not necessarily all about agents. You know, often for some organisations, the starting point is actually what are people using more broadly in terms of AI tools.

00:04:40 **Matt Bruun**

Yeah. So you I mean often that links into something that we've mentioned before, which is things like ephemeral agents or, you know, brief agents and brief NHIs. And so the ability to detect those when they show up. There's really kind of two ways of doing that. We have a technology called non-human identity detection and response.

What that's looking at is, you know, are keys, tokens, credentials that I recognise as being used for other non-human identities, being used with an identity that I don't recognise. So that's a sort of a kind of like a factual, fact-based thing that is that, you see, and then when you see it and detect it, you can then introduce a workflow to take care of that.

So depending on what that agent is trying to do and the data that it's trying to access, you can make a decision. People always talk about things like kill switches.

00:05:35 **Matt Bruun**

Hopefully those aren't required very often. Normally it's just a question of someone has not followed what they're supposed to have done when designing an agent, testing an agent, or releasing it, and it just needs to be brought into compliance and registered properly. And then also the other way of looking at these things is to look at building patterns of behaviour.

So what do we expect from these agents, these types of service accounts, the way they interact with each other. You build up patterns of recognition and then if you start seeing anomalous behaviour, much like we used to talk about with humans years ago as well, when people start logging in from different places or doing different things like that, same with agents and non-human identities.

00:06:14 **Tom Eggleston**

So again, there's this kind of a blend there moving from discovery also into detection. There's kind of both of those aspects to it as well. I thought the example you gave about the use of a credential that is not expected is a really good one, because one of the things that we know about agents is they'll instantly, you know, have access and use whatever they can find.

Yes, within the bounds of what they have access to, which humans won't necessarily do. So the ability to detect that, I think is really quite powerful. Yeah. So you've obviously worked with a number of customers who've been through a discovery programme. If you've got any good war stories about what they've found.

00:06:47 **Matt Bruun**

Yeah, there's quite a few. I mean, I think now I think customers are at a stage where they expect to be surprised. They know they're going to see things or be presented with things that they didn't want to see or they hoped not to see. So now they're kind of expecting that. So we're sort of over that initial surprise.

I think one of the things that really alerts customers is when they've still got agents and credentials out there from employees who've left the organisation, that's a big one, because or contractors who've been in and given someone the keys to do these things and perhaps even vaulted things as well.

But those credentials are being shared across other different agents and different identities. So that's usually a big red flag. Another one that's always a big surprise is quite how much access agents, but also third parties can have through the agent. So people being able to access really important proprietary data or PII information when they're using an agent because it's crossing boundaries internally and externally.

So again, to your point, the agent is trying to help. That's what it's been set up to do. But it doesn't necessarily consider the things that we might do as humans. And I guess that's really where things like compliance and audit then start to come in, because you have to set standards. And standards have been set by organisations like the EU Data Act.

They're still present in GDPR and NIST and DORA and so on. And agents just make that even more challenging to abide by. So having the ability to discover everything then bring everything into compliance and then look for where those anomalies are, find them, stop them and then bring them back into compliance.

That's what most customers are beginning to realise they need to do. And just to build a process around that, it's not a project, it's a programme. And it needs to keep going.

00:08:47 **Tom Eggleston**

Yeah, some really interesting stuff in there. I think the one about, you know, agents having so much more access than you'd expect is that really resonates. And there's something I like to say to customers, you know, as an exercise to do carefully, you know, use Claude or Copilot or whatever it is, if it's connected up to your SharePoint or OneDrive and just ask it, show me everything I've got access to, and then ask yourself the question, should I really have that access?

Because, you know, as people, we're quite lazy and we don't necessarily go looking for things, but an agent's just going to find it. And kind of to your point, you know, what agentic AI does is it really shines a light on some of the human processes that maybe we've not completely nailed down as well as we need to, because humans are not working at machine speed.

Yeah, I think all of which kind of brings us into a point around, well, in my model it starts with discovery. The next thing is governance, right? So, you know, once you've discovered all these agents, we need to have some governance around them. We need to know what they're bounded to do. Here's a question for you.

If you've already got SailPoint's IGA tools, why can't we just use that? Why do we need something different for agentic? What's the difference really, between sort of, you know, human governance and this kind of non-human or agentic?

00:09:59 **Matt Bruun**

Well, the theory of the case is the same. You still need to have things like, you know, role-based access, intent-based access and so on. You just need to govern it. The difference is two things. First of all, you've got to be able to do it in real time. So first of all, I mean, we talked about this in the last session that people have tried to get to zero standing privileges.

And that's even more important now with agents and NHIs. Because without that, you're going to be just constantly playing whack-a-mole. And so bringing them into your governance models is absolutely the right place to start. Yeah. And so for us that the tools aren't really separate, they are absolutely linked.

So it's very when we do the governance and we are bringing everything. Sorry. When we do the discovery and we bring everything into a unified registry. One of the key things that you'll see is probably like the second column is who is the owner, who is the individual that is accountable for that agent or that non-human identity.

Because linking human identities to non-human identities really helps us understand, well, what should we be able to do and what shouldn't we be able to do? Some of the organisations we've spoken to, particularly some of the ones who are more like software organisations themselves start with human identity, say based on my role as a software developer in this team or this level of software developer, I can create this level of agent with this level of autonomy boundary, you know, and privilege and persistence.

But if I'm a in a different person, in different part of the business, I can't do the same things. So and then, of course, if something goes wrong, when you're discovering something or something goes wrong

in real time and you need to find the owner, being able to very quickly tie it back to that and introduce that back into your human workflow process for this is critical.

So for us we see them as absolutely interlinked. However, there are lots of customers that at the moment aren't SailPoint IAM customers, but they do have a really critical need for getting their NHIs under control now. So you can still use our SailPoint Agentic Fabric for all non-human and agentic identities and still use that independently of our core IAM IGA tool.

00:12:26 **Tom Eggleston**

I really like the idea of agents, you know, having a human as the owner and looking at the relationships between all the different humans and using it as a way of determining what access you've done, that, of course, that's something that SailPoint's really strong at with identity graph and all that side of it.

Another point that was interesting there. You talked about this sort of the registry and the classification of agents. I like to think about this kind of concept of almost like entitlement radius, kind of it matches with the term blast radius, which kind of happens after the effect. Yes.

But when we think of entitlement radius, we bring that right to the beginning of the conversation and think about what is the total, you know, radius of entitlement that we're going to give to this agent. And that's partly about what you inherit from the human. But also it's about the classification of each individual agent.

So, you know, talk me through that a little bit.

00:13:16 **Matt Bruun**

So we look at classification of agents in two ways. The first is the boundary. When we talked about the level of access that it's got internal external. But then things like its autonomy then the level of privilege that it's got, but also the persistence we talked about, users no longer being in the company or former employees.

That's a key link between the two things there. The other key element of classification is the data. What data does it have access to? And you know, in our experience, about 80% of enterprise and corporate data is unstructured. And so being able to balance what the classification of the agent is by its type, then adding in all of the contexts around the data that it can access, where that access sits, what kind of data it is.

That gives you a far better picture of the level of risk for an agent. So in your entitlement radius, you're absolutely right. There's the entitlements. And then yes, you're right. Blast radius sounds very much like what's happened after the event, but it helps you model the potential blast radius and say, okay, this is just too risky for us to do.

We need to bring down this level of access. We need to bring down this level of autonomy or the amount of time it's going to run, or limit the access to the data that it's got.

00:14:33 **Tom Eggleston**

And just staying with governance for a little while. Obviously one of the things we do with humans in SailPoint obviously has had a real strength in this is this idea of certification and not just allowing, you know, we've got automation that's great. And we actually need to go and check that access is being done correctly.

What does that mean in the agentic world. You know. Does that what does it look different.

00:14:57 **Matt Bruun**

Yes it does. And I think that's you know, with the recent acquisition of Entro, it's closed the gap that we had around particularly around credentials. So the tokens, the secrets and so on that non-human identities are reliant in order to interact with each other. And so, you know, the kind of things that we're looking there around, certification is either forcing recertification, making sure if we see things like, you know, expired or shared or overshared credentials, making sure that we highlight that.

And then also things like unvaulted credentials, you know, if people are using secrets that are in plain text and available to anybody. You are creating problems either immediately or for further down the line. Now, what we don't do is try to solve all of that. What we do is enable customers. We highlight the problem.

We enable them to put it into their workflow, and then using whatever systems and tools they've got to solve that. There are really good, you know, PAM vendors out there that are, you know, vaulting is their world. And, you know, we don't want to get into that world. They do that really well, but we want to help customers work out how to leverage that better and reduce the risk that's associated with, you know, improperly governed NHIs and agents.

00:16:12 **Tom Eggleston**

I think it's really interesting there as well. Again, you're talking a little bit about bridging from this agentic world into the human world, right. So we can detect something on the agentic side. And then we're using kind of human workflows to actually take some action and do something about it.

00:16:26 **Matt Bruun**

So yeah, that ability to take action. Most of the time you want to automate it as much as you can. It could just be as simple as a simple workflow to automate the credential being put into a vault, for example, like that. Other times it might be, you know, we've highlighted this particular user as a persistent offender. They're always using the same credential for whatever agent they're building.

They're sharing that credential. Or their build quality is poor. And it's never passing our kind of initial tests on these things. So there's an education kind of workflow or education opportunity there to help them improve as well as reducing the risk.

00:17:10 **Tom Eggleston**

I really like this idea around education. And you mentioned patterns earlier. I think there's always a danger we all do this in the identity space. You know, we think about security first, you know, and there's a danger that you sort of become almost like the Department of Business Prevention. Yeah.

So you can't do that. It's too risky. You're not allowed to do that. And that's obviously a really important part of all of this. You know there's a lot of risk around NHI. Mhm. And of course we need to manage that risk and stop bad things happening. But I think we have a bit of a responsibility as identity professionals as well to help our customers, to help, you know, CISOs within organisations also to become an enabler because for organisations there's not really a choice.

You know you have to push forwards. Certainly in a lot of industries. You have to push forward with an agentic agenda, because if you're not doing it, your competitor is doing it and it's not a nice position to be in to just be saying no. Yeah. So I'm really interested in this idea of how we can develop patterns and sort of give those to people and say, actually, here's some patterns.

They're going to help you go faster.

00:18:14 **Matt Bruun**

Yeah. I think you make a really good point, though, about not wanting to just be people that either say no or business prevention, but actually what you really want are your users to participate in this because they want to protect the company, the assets that belong to the company. The information.

One of the going back to where we started with discovery. But one of the tools that customers use is something called shadow AI remediation. Now, that's the ability to discover what are what SaaS apps and agent building platforms are in use within my organisation. That's a browser plug-in.

But then what it does do then is it enables you to sort of tighten that down and say, well, here are our sanctioned apps and here are our unsanctioned apps. So that's an education point. But then if I then let's say that my organisation has decided and I'll use this as a real example because this is what we do.

My organisation decided that our corporate tool is going to be Gemini. But we're not going to use Claude, for example. So. Okay, so if I tried to go to Claude, what it will do in the first couple of times I go there, it will educate me that you're trying to access an unsanctioned app for corporate information.

Please, please. And it will redirect me to our version of the Gemini Enterprise tool. And then after a while, it might not. It just might redirect me there straight away. And then after that, I might just. It might have twigged that I probably need to just keep going straight to Gemini in that instance.

And if you've got them people who are persistent offenders, there's a different education opportunity for them perhaps.

00:19:55 **Tom Eggleston**

So yeah. Education opportunity. I like that phrasing.

00:20:00 **Tom Eggleston**

And I suppose that sort of approach of education and giving these patterns also applies to people who are building agents, you know, and giving them the tools to do that in a way that actually is going to be secure and in line with policy.

00:20:12 **Matt Bruun**

Yeah. Yeah, absolutely. I mean, I mentioned this earlier, when you've got this sort of you've got the key platforms that people are building it with and then the foundries as well, and then sort of their own agents as well. The way they're built and more and more people now are designing and building their own agents through low-code and no-code solutions as well.

So setting the standards for what you want those agents to look like, how they're built, how they're tested, how they're deployed, how they're announced, and then how they're run. All of that is that's all

about, again, back to basics. With non-humans, it's all about compliance and then enforcement and then catching the outliers and bringing them back in.

Because I'm back to your point around being a business enabler. That's actually what we're trying to do here. There's we're trying to help customers get more agentic technology out there more quickly, but in a way that is safe, secure, and that they can have confidence around pushing out and rolling out and really adopting to get all of the business benefits that we know are the real reasons why people are doing this.

00:21:15 **Tom Eggleston**

And just picking up on that point of enabling, again, you go out and speak to customers all the time.

00:21:21 **Tom Eggleston**

Where are customers at the moment? You know which organisations are enabling, which customers are doing nothing because they're too scared and which customers are kind of, you know, pushing ahead, maybe in an unsafe way. You know, where do you see people at the moment?

00:21:34 **Matt Bruun**

I'm going to use your phrase about patterns. So there are patterns of success that we see in customers. And often I'll use that language when I'm talking to them. But the ones who are probably, I would say doing the best because they are furthest along in terms of their adoption and their use, but also the way that they are remembering to embed identity security and protection right from the very start.

And that's because they are taking the mandate that's coming from the board or the CEO to we, you know, we can reduce costs, we can improve our top line, we can improve our operational efficiency. They're taking all of those things. And then they're realising that there are probably four groups within most organisations that are responsible for managing all of that and for making it work.

The first group is the traditional identity and access management teams because they're still ultimately the people that the board and the CTO, the CFO and the CEO look to make sure that the identities are secure and being managed. There are increasingly now the sort of the AI and agentic teams, the people who are actually building the applications and the agents that are pushing out.

But there are two additional teams. One of them is the AppSec team. And again, that going back to what I was talking about around all the credentials and tokens and secrets and things like that, the AppSec teams are the ones who really care about the security at a very, very detailed technical level. They've got to know and trust that, as I said earlier, not just the build quality, but then all of the ways that these agents and non-human identities are going to be interacting and how they are identifying to each other, that's a security thing rather than an identity thing.

Traditionally. And then the final group of the is the people in the SOC. They are the ones that are having to respond in real time to, you know, they will get the workflow for various different alerts based on the classification and the risk levels that we talked about earlier. They're the ones that have to make decisions.

You know, what can I automate? And therefore where might I need to do something more? I mean, ultimately that might end up in the kill switch. But as I said before, the kill switch always feels like it's too

late. We'd far rather be able to manage everything, bring it into compliance, see everything, and then as new things emerge, be able to detect them.

And as we were talking about Hugging Face, that's a really good example of one of the phrases that you hear a lot at the moment, which is agentic drift, that's when agents just exceed the boundary and the intent that they were originally designed for. It's the natural state to go and find more things and to do different things and find ways around things.

That's again, that's how Mythos got to that stage as well, where, you know, the person came in the next day and it presented that user with the choice, you know, should I do this? So that is what happens when you don't control what an agent can do. Now, you can't control an agent's nature, which is to do things. So but what you can do is use its identity to make sure you stop it from doing something when it does start to drift, because it will pick up based on its anomalous behaviour, its anomalous intent, and different and other signals like trying to access different data sources or more privileged data sources where you need to start taking action.

00:25:08 **Tom Eggleston**

I think as identity specialists, that really brings home that really important point, that identity can be the backstop. You know, the other detections, you know, they might be fallible. That's why we need identity there just to really draw a line around what an agent can do.

00:25:22 **Matt Bruun**

Absolutely. And I think it's why, when it comes down to it, if you think of all the different cloud service providers that there are from productivity to ERP, you know, Workday, communications, all of those different things, they're everyone is talking about providing, you know, agentic security within those.

But ultimately the one consistent theme across all of those things that makes everything work together is identity. And that's why we're still seeing customers want to go to dedicated identity providers to make sure that they are the link between human and non-human identity to make all of this work together, properly and securely.

An example of that. Most pertinent at the moment is Hugging Face. And I was checking into it. Could SailPoint have stopped what happened there? We couldn't stop the initial vulnerability because that exists there. But the point at which those stolen secrets were being used to access other areas, and it was linking to other non-human identities and agents from going around things.

That's the point at which that we're looking at that's anomalous behaviour and the use of shared or stolen credentials and secrets to do things. So for me, that was a good test about whether, you know, when you do need to do something, could we find it, could we stop it? And do we know how to stop that as well as raising the flag?

You've got to be able to tell people this is what you actually need to do in order to bring that into control, like remove the shared secret and now use a new secret and vault it.

00:27:00 **Tom Eggleston**

Hugging Face was really, really interesting. I mean, we were prepping this yesterday and he was saying, well, we're going to talk about Hugging Face in three weeks. It'll probably be a different one. But at the moment that's the one that everyone's aware of. And I think it really brought home for me. You know,

actually identity is almost like a backstop that you need to have in place because some of the other detection technologies or detection tools are out there.

They're not 100% reliable. And actually having that identity layer that's actually able to pick this stuff up. Limit the radius, determine when credentials are being misused or whatever it may be. Really is fundamental in actually, you know, being able to limit the damage that can potentially be done.

00:27:35 **Matt Bruun**

Yeah. And just to take it one step further, I was with a partner in France a couple of weeks ago, and this partner quite smartly said, you know, people know that people have accepted they're going to have some kind of breach for a long time now, pre-agentic. Now most customers are saying we know we're going to have some kind of breach, but it will run at agentic and non-human identity speed.

And what this partner was saying is you've got to be able to prove that you did everything. So going back to the identities. And so when you have cryptographic evidence about the flow that sits around an agent, its behaviour. What you tried to do to make sure that everything was discovered, put in a unified registry, brought into governance.

And then when something goes wrong, if you can prove that you have done everything you can, the level of your fine is going to be far lower than had you been fast and loose and just focused on the business outcomes. And, you know, that was a partner's perspective on that. And I think they're right, actually. You know, people do know it's going to be costly if they get exposed. I mean, the reputation of the company. And we were using Hugging Face just now. And before that it was Mythos. You know, some of these things are scary and they can have significant impact on companies.

But it's one thing having some reputational damage. It's another thing being wiped out by the level of fine. So being able to prove that you did everything you could, and consistently is absolutely critical.

00:29:15 **Tom Eggleston**

Yeah, absolutely. I think that definitely resonates this idea of, you know, observability. But even beyond observability, auditability of everything that happened all across the way is absolutely vital. And it makes me think again. I'm actually going back to your previous point about the pattern for success and all the different groups, you know, that sort of that security SOC team, that's their world a little bit more is absolutely vital part of this.

And the takeaway from that really is to try and boil down that pattern of success. I think what you were saying was it kind of takes a village, right? You know, it's not you're not going to be successful as an identity team trying to get control of your agentic AI security initiative. The real pattern here is collaboration across the different parts of the organisation that have a stake in the outcome.

00:29:58 **Matt Bruun**

Yeah, absolutely. And going back to the I talked about shadow AI remediation and that tool. That's a way of you know, one thing I'd say to all customers is it's critical you just get started and just do something. And this is a really good way in. Of gathering information that's just about your organisation, because there's plenty of information out there about other people and other organisations and the Hugging Face example.

And but when people understand that they're looking at information and insights, which is about their business. So the number of SaaS apps that I've got running in my organisation, the number of agent platforms that are there and the number of unsanctioned apps or that are still being used and so on. So that gives people an understanding of, okay, now this is about me.

This is about my environment. This is about the things that I didn't know and about behaviour that I probably don't think I like. The fact that, you know, execs are using a tool like Gamma to turn really nice, important corporate information into beautiful looking PowerPoint presentations. Well, that's great for the way they look.

But as soon as you're using a tool like that that's out there and it's searchable. It's findable for other people. So that's a really good place to start. And that often is. And quite often the identity teams are using to say, okay, I want to or the AI teams as well, they're working together to say, now we know what's going on here.

Now we need to take a deeper look, and then we'll go in and look at actually what are the non-human identities and the agents that are actually out working within my environment. So it's a nice way of starting and then justifying a deeper look. And then, you know, we had a customer a couple of weeks ago who we ran our discovery tool there for about four days.

We found 6,500 agents and non-human identities that they didn't know they had on. Of the 250 or so agents that they were using, I think there were about 30 different agent building platforms that had been used. And then again, almost all of them were using shared secrets or unvaulted secrets. So that is usually enough to make most organisations, whether they're in the SOC, the identity team, the AppSec team or the AI team and the agentic team take notice and say, right, we've got to get this under control.

00:32:27 **Tom Eggleston**

I love that because it's that's a real wake-up moment. But that wasn't a big investment or a big piece of work to get to that level of insight. You know, that's a pretty simple browser tool. I think it was a.

00:32:38 **Matt Bruun**

Browser tool, first of all. And then the actual true discovery tool, which is that you need to get into the platform to be able to use. But yes. So they went from looking at their general environment through the shadow AI remediation tool. Then they use that to identify. Well, this is it seems like we've got a lot more going on than we thought.

Now let's take a deeper look. And that's when we discovered all of that data.

00:33:04 **Tom Eggleston**

So that's quite a useful point, I think, to actually, you know, talk a little bit about Entro. I think one of the things that is difficult for customers at the moment is it's all moving so fast, it's kind of hard to keep up with, you know, what each of the vendors are doing and what each part of the product does.

So obviously SailPoint recently made an acquisition of Entro. You mentioned it a little bit earlier, and I think that's the tool you were just referring to. So I mean, yeah, maybe it might be useful just to explain a little bit about what Entro does and what it, you know, how it adds to the existing SailPoint platform to uplevel the capability a little bit.

00:33:42 **Matt Bruun**

So sure. Very happy to. So yeah, Entro has been part of us now for a few weeks. It is now integrated into what was Identity Security Cloud. So we can already see the data that comes from the discovery part of Entro surfacing in Identity Security Cloud, which enables us to start putting together the human identity piece as well, so customers are already able to see that.

And then in terms of how it works, I talked earlier. So the name Entro comes from using entropy strings to go off and find this data. So that's where it's tapping into the. And the difference I think, between what we could do before and what we do now is really where it's tapping into the developer pipelines.

And so looking for these credentials, secrets, tokens and so on. And it's basically looking at, you know, finding and exposing and surfacing the data that already exists, which identifies all the different agent types. Then it classifies them. And then you bring that together with the, like, the data access tool that we've got.

And then you get much richer data, much richer classification and prioritisation on what you need to address. The other thing that Entro brought was almost at the other end of this, which is runtime, in real time. So their ability to discover, you know, agents, ephemeral agents that surface really quickly, or long-term agents that are going to be in production for a long time to immediately be able to find those based on their activity.

So this is the non-human identity detection and response piece that's able to then either again one extreme kill switch. Other extreme. Fine. It's been built by someone that we know, someone that we trust. It's been built using the tools to the standards that we expect. Actually, we're going to let that agent run because there isn't actually an abnormal risk that it's just an agent that we haven't seen before.

So but because everything else about it is correct, we'll just be able to watch it. And if it starts doing anything anomalous, like trying to access data that it shouldn't, or trying to do something with that data that it shouldn't, then we can step in and again, back to workflow. We can do something to insert ourselves into the process.

00:36:03 **Tom Eggleston**

I think it's great to see SailPoint now having, you know, moving into the runtime space as well, or the real time detection as well. I think that's a really interesting move. And, you know, I think Entro clearly is adding a lot to the platform and those capabilities. So we talked a little about the kind of average state of customers at the moment.

And they're kind of, you know, at different stages. Fast forward sort of 12 months. Where would you hope that the average customer has got to with agentic security?

00:36:32 **Matt Bruun**

Yeah. You could probably see me shifting in my chair there. So that is a good question. Actually I haven't been asked that before. Where do I think? I'll go back to what I said earlier about different customers are on different stages of that journey based on where they are. I think every customer needs to at least understand their posture.

What does my posture look like? I have to accept that in almost as a customer, you have to accept that you are going to have things in that environment, in your various environments, that you don't know

what they are and you don't know what they're doing with it. So number one has to be all customers, at least have a sense of what they've got and be on their way to having a plan to do something with it.

Depending on the type of organisation and the industry and the data they have, that will all vary, of course, but for, you know, the people who've got the most to lose and the most at risk, they should be well down the road to having full time governance that is auditable and is compliant with the different standards, but also that they're able to detect and respond in real time.

So that's the if you like, the financial services types of organisations and so on, because they have the budgets and they have the teams, and it should be fully integrated into their workflow, right from a design and build of the agent all the way through to the end of life of an agent and the retirement of one based on what's going on?

Again with the human identity there. For other organisations, having discovery is the absolute minimum, having a plan to work out how to deal with that. And at least having sort of implemented and tested things to know, you know, does my plan work? So I think that's the sort of it's a I've sort of I don't know whether I've evaded your question by saying it's going to be different for everybody, but I do think it will be.

But I think back to where we started. And you started with your five pillars. Everything comes from discovery. They've just got to get started.

00:38:31 **Tom Eggleston**

Discovery. That makes a lot of sense to me. I mean, there's two things. Just get started. I think that's a really great phrase for everyone to take away from this. And really what you said there, you know, in 12 months everybody should have done discovery. Everybody, whatever kind of organisation they are, they should you should know where you are and you should know what you've got.

You know, how you govern that and what steps you take are going to vary based on the kind of organisation. But I think that's a really good takeaway. You know, that's something that any organisation can really be getting on with now and then.

00:38:59 **Matt Bruun**

You can have a plan, then you can say if you've done your discovery, then you can go and say to those four groups that I mentioned earlier, the AppSec teams, the SOC teams, the identity teams, and the agentic and AI teams. We need to get together and come up with a plan that our CEO and our CFO will sign off on, because what it's going to do is help us grow at the speed and level that we want to, but in a way that is safe and sustainable. You know, it's all about for me. I always have sort of the same two words.

It's readiness and resilience. You've got to be ready to do this, but you've got to be making sure you're resilient as well.

00:39:37 **Tom Eggleston**

Thanks so much, Matt. This has been a really, really interesting and useful conversation. I hope for all of our customers that are listening in today, in helping them to make sense of what can be such a fast-moving and complex space. Have you got any closing words that you'd like to share before we finish?

00:39:53 **Matt Bruun**

Yeah. I mean, we talked earlier about just getting started, and I wanted to make sure people understood that, you know, as a trusted partner of SailPoint, ProofID is ready to help those customers, particularly with that first step of just getting going. So you're ready to help those customers use our shadow AI remediation tool to discover what AI tools and platforms are in use within their organisation, and then to start that process of doing a little bit of education on how do you tighten those things up to think about sanctioned and unsanctioned usage, and then educate their users so that their users become conditioned to wanting to be part of helping the company grow, but in a safe and secure way.

And, you know, I'm really pleased that you and the team are ready to go ahead and do that and help customers right now.

00:40:44 **Tom Eggleston**

Matt, thank you very much indeed for joining us today. And thank you to everyone listening for joining us as well.

Continue exploring agentic AI security

Want to explore agentic AI security further? Visit our Agentic AI hub for more practical guidance and resources: https://proofid.com/solution/agentic-ai

If you're thinking about how to put this into practice within your own organisation, our Agentic AI Enablement session can help you understand where you are today and build a practical path forward: https://proofid.com/agentic-ai-enablement