[JAVI BELTRAN & JORDAN GREENBERG, "TELEBOT"]

JORDAN GREENBERG: Welcome to season five of *The Prodcast,* Google's podcast about site reliability engineering and production software. This season, we are continuing our theme of friends and trends. It's all about what's coming up in the SRE space, from new technology to modernizing processes. And of course, the most important part is the friends we made along the way. So happy listening. And may all your incidents be novel.

STEVE MCGHEE: Hey, everybody, welcome back to *The Prodcast.* This is Google's podcast on SRE and production software. Hello Jordan. It's been a little while.
JORDAN GREENBERG: It's been a while.
STEVE MCGHEE: Our co-host, Jordan, is back.
JORDAN GREENBERG: Yes, I'm back.
STEVE MCGHEE: Hello. How's it going?
JORDAN GREENBERG: Going well. I'm very excited about our guest today, who usually meets people in times of crisis. I skipped over that piece, meeting our guest at a work conference instead.
STEVE MCGHEE: Yeah, that sounds better.
JORDAN GREENBERG: Yeah, I think so. It was. It was great. Would you please introduce yourself?
HEATHER ADKINS: Yes, hi. My name is Heather Adkins. And, well, I work at Google. No surprise there. And I lead what we call our Office of Cybersecurity Resilience, overseeing all the great people working on security at Google and all of our initiatives.
And I've been at Google for a little over 23 years. So I've gotten to see the company from when we had two products to today, when we have things that fly. And then I guess Sundar has announced we're going to have things in space soon. So that'll change the threat models.
JORDAN GREENBERG: Oh, my goodness.
STEVE MCGHEE: Why not? So your employment is, like, in grad school, basically, is what you're saying. 23, is that right? Something like that?
HEATHER ADKINS: 2002.
JORDAN GREENBERG: Oh, gosh.
STEVE MCGHEE: Yeah, pretty wild. Well, we didn't meet in a time of crisis. I believe it was a Christmas on-call shenanigan of some kind involving reliability and security at the same time, and making sure bad stuff didn't happen, or if it did, it wasn't that bad, or whatever. We don't have to go into the details. But it was a memorable Christmas Day, I can tell you that, at least on my side.
HEATHER ADKINS: All of my Christmases at Google have been memorable. Each and every one of them.
STEVE MCGHEE: Oh, boy. OK.
JORDAN GREENBERG: A present in each of them.
STEVE MCGHEE: That's right. That's right. Well, speaking of which, it's not like I set this up at all, but speaking of SRE and security engineers working together, incidents tend to be that thing that we work together on. But they're kind of different. How have you found working with SRE? And how different are these incidents? And are they tremendously different? If there was a Venn diagram, how close to two circles versus one circle is it? Stuff like that. Can you enlighten us a little bit on the differences there?
HEATHER ADKINS: Yeah, I think the primary difference between a reliability incident and a security incident is that the instigation of the incident is usually a malicious human. And then

there's a Venn diagram overlap of when the malicious human causes a reliability problem.
But for the most part, what we try to add into the incident management function during a security incident is the element of, you have to anticipate the behavior of someone else. And there's this kind of philosophy in the field of, when you find the hacker, when do you turn them off?
And in a reliability incident, usually when we find the root cause, we want to fix it as fast as possible and move on. In a security incident, you have to step back and play a little bit of a game of chess with the other side. And so that is kind of a different sort of externality to security incidents that we often get to have a real candid conversation with.
This is going to change completely with AI, because we're not going to know whether that's an agentic AI hacker or a human. And that's going to make things very interesting for us. But I think otherwise, there's a lot in common.
I think when you look at the methodologies for analysis, getting to root cause after the fact, doing a blameless postmortem, there's actually a lot in common with how you do reliability and how you approach security incidents. And in fact, at Google, our SRE teams and security teams use the same incident management frameworks for that reason.
STEVE MCGHEE: Yeah, I was going to say, the tactics, and the tooling, there's a fair amount of overlap. And I was going to actually just say the exact same thing. I would suggest trying, companies out in the world that aren't Google, building bespoke copies of each other. I've seen that happen in some places.
Have you experienced that? Do you ever work with cloud customers, for example, that quote an SRE who works somewhere else, not at Google, where they're trying to have these tools, and they're like, oh, should we use the same one as secure-- they have whatever. Should we use that? Blah, blah, blah. You know what I mean?
HEATHER ADKINS: I'll be honest, security people are usually pretty picky about their tools. And even at Google, we have a-- the incident tracking for reliability is a tool. Everybody uses it. And it's really great, because if there's something down, you can go, and you can check it out and what's going on. I have to consider whether there's a bad guy reading my incident management log. And so-- no. And also--
STEVE MCGHEE: Yeah, that's tough.
HEATHER ADKINS: Yeah, it's tough. So putting an advertisement in our internal outage management tool that we are doing a security incident would be probably a misstep on the OPSEC side for us. But the other thing that I found is security people, increasingly, we've been trying to integrate as much security data context into that tooling.
And so you end up with a situation of, do you want to try to put all of that integration into your incident management tracking tool, or do you want to try to create a custom tool? So there's a little bit of fragmentation still.
But I think, in general, what's helpful is just to make sure you've got how and the mechanisms right, and that two teams could work in each other's tools without it feeling too alien, depending on what kind of issues you're working on. But again, I think each environment is going to see unique elements of that.
STEVE MCGHEE: Yeah, I was just watching the-- last night, actually, I was watching the movie about Alan Turing that came out a few years ago or whatever. And that near the end of the movie, spoilers, the war happened and stuff.
But there was a point where they're like, we know this thing is going to happen because we broke the code. And they're like, we can't tell anyone because that would spook the other side into knowing that we broke their code. And then that would just change the whole equation. And we can't actually do the thing we want to do, which is like, shut it down right now. And it's like, oh, boy, that's a tough lesson.

HEATHER ADKINS: Right. It's your instinct, yeah.
STEVE MCGHEE: I mean, it's arguably tougher when there's direct bombings on vessels and stuff. But that's hard. I mean, how do you approach that? How do you know when that tipping point is? Maybe that's a whole different podcast.
HEATHER ADKINS: But yeah, there's a whole discipline of this in information warfare. And there's some great manuals the DOD-- the US DOD-- has published on this. But there's no easy calculation. And especially if you don't have perfect visibility into how your threat actor is thinking, or in the case of kinetic warfare, how your enemy is thinking. But there's thousands and thousands of years of military study on this, everybody always trying to answer that.
For me, though, I think the most important thing is to focus on limiting damage and to focus on having a complete picture. And so we're always trying to figure out, how do we put together the flow diagram of how the attack happened, where it started, what are all the things that a threat actor might have touched, all the data, all the machines, put that picture together as quickly as possible and really focus your pre-event automation on building that.
Because the faster you get the complete picture, the faster you can turn stuff off. And the idea is that we cut everything out, and you make that kind of cutoff happen, and then you move forward.
STEVE MCGHEE: Got it.
JORDAN GREENBERG: So there's a creativity piece of it, where if you are a malicious actor and you're looking to compromise a system with the tools that you have, or if you're upleveled with AI tooling, you get to be very creative, because you get to see the full picture of what someone has done.
And when you are creating something to be used by humans, security is built in. We do secure by design. We work on involving security right away to fix something so that we can catch these cases before they happen.
So this is something that we've mentioned in previous security and SRE talks, where we've had Jessica Theiddat with us, talking about how people do make these-- or compromise systems. I would love to hear a little bit about how maybe some of the principles that we have in secure by design that help us build reliable systems, some pointers or tips that maybe you would educate SREs with in a book, maybe.
STEVE MCGHEE: There's a book that has those words on the cover. You may have seen it.
JORDAN GREENBERG: Oh, my gosh.
HEATHER ADKINS: A real labor of love.
STEVE MCGHEE: Look at that. Look at that.
HEATHER ADKINS: About 100 of us wrote a book. And if you've ever had to write a book with 100 people--
JORDAN GREENBERG: No, I'm sure it's great.
HEATHER ADKINS: But no, it's really amazing. You got 100--
STEVE MCGHEE: It's a good book. It's a good book.
HEATHER ADKINS: Thank you. We released that in April of 2020, the first month of the pandemic. And so--
STEVE MCGHEE: Perfect.
HEATHER ADKINS: What a way to go into a pandemic. It's been really great.
JORDAN GREENBERG: At least you had something to do. You could read the book.
HEATHER ADKINS: Not so good for advertising it. But we had a really great collaboration. And the approach we took with this was, in every chapter, we wanted a thesis for why this is a reliability topic, and why this is a security topic, and how to blend the two.
And it was through doing that that I sort of realized that a lot of the work we'd done on insider

threat from a cybersecurity perspective matched almost perfectly with how SRE was thinking about reliability. And if you think about, how do I keep an employee from doing something malicious, or how do I keep an employee from making a mistake, the solutions are often very similar.
A peer review, for example. If you typoed an integer setting in a configuration file, that could either take the website down because you typoed it or because you really wanted to. So we really sought to look at this Venn diagram phenomenon that we were starting to see and pull out examples, pull out capabilities, approaches to how to think about the problem, but also how we solved the problem, and give people some examples.
And as such, it's a really dense book. And my recommendation, if you do tackle it, is there's an open source version online. We felt very, very passionate that while you can get it in a book form, and it's been translated into many languages, we also wanted it to be something that people could come back to, and refer to, and to share.
But it's really good in a book club. Take one chapter at a time. It's really good if you've got a particular set of problems and you cherry pick one chapter here or one chapter there. It's not exactly a beach page-turner. So a little dense and deep, I've been told. But we did try to share as much wisdom in there as we could.
STEVE MCGHEE: Awesome. We appreciate it. I know I've shared chapters at a time from the online one you're referring to with customers and be like, it's this one, read this one. Please get the book if you want, but just read this one for now.
HEATHER ADKINS: And you'll see, there's different authors on every chapter, but there's always more than one author. And that's because we pair them up, somebody who's an expert on reliability, someone who's a security expert. And we really built that together, the narratives.
STEVE MCGHEE: That's awesome.
JORDAN GREENBERG: And things are usually better when they're built in that way, when people are working hand in hand, and they're able to be really good at what they're good at. It's usually excellent. Shifting gears just a little, a recent white paper was published on agentic AI hackers. Is this something you're worried about right now?
STEVE MCGHEE: Is this real?
JORDAN GREENBERG: What else are you worried about right now? Is it real? What is that?
STEVE MCGHEE: I'm afraid it's real.
HEATHER ADKINS: We're worried about everything. We are starting to see the adoption of large language models by threat actors. And it's about what you're seeing in business at the moment. You have productivity enhancements. In the business, we're seeing summarization of complex data. Here's a 4,000-page paper, give me a summary, kind of stuff.
We are also seeing these productivity enhancements by threat actors. You could imagine, write me a phishing email, translate that into another language, make it more realistic. We're also seeing the large language models get used for tasks, so sort of stepping up from just productivity enhancements to full-on tasks, including some early versions of polymorphic malware. So they will use a large language model to generate a new version of the malware for each use. And this is really--
JORDAN GREENBERG: Wow.
HEATHER ADKINS: Yeah, it's scary. We've known about polymorphic malware for probably, like, 2 and 1/2 decades. It's certainly not the first example. But this helps it scale. And it's so easy to do.
But imagine your antivirus software, or endpoint detection response software, as we now call it. It's probably not going to catch it with its kind of very basic capabilities. You'll need very advanced capabilities to catch these things. And I think we're going to see more of this.

And then we might also begin to see the full, end-to-end identification of hacking, what we call the cyber kill chain, trademarked by Lockheed Martin. But the fear is somebody will be able to type into their tool, "please go hack company," whoever you're targeting.
And it will do all the research. It'll find out what kind of software you use, find a vulnerability in it, conduct a phishing attack, exploit the vulnerability, laterally move through the network, find the data, steal the data. Game over.
STEVE MCGHEE: It's not great.
HEATHER ADKINS: It's not great.
STEVE MCGHEE: I don't love it.
HEATHER ADKINS: It's not great. But that sounds like a really scary idea. But we're going to see those kinds of tools in the hands of what we call the red teams as well. And red teams are the friendly hackers that we-- most big companies like Google have them. And we pay them to hack Google and tell us what we need to know so we can fix it.
They're going to have these tools as well. And so we'll get to simulate this against the environments. And we'll get to practice. And we'll get to see what it's like. And we'll be able to build similar kinds of defenses into IT, production, et cetera environments. It'll change the face of, I think, networks and how we set up networks pretty significantly.
STEVE MCGHEE: So yeah, terrifying. OK, got it.
JORDAN GREENBERG: We're all scared.
HEATHER ADKINS: Mission accomplished.
STEVE MCGHEE: You did it. Well done. We can end here. No, you were speaking a little bit in the future tense, but I have a feeling this is also present tense. Without divulging things, is this kind of working? Like, should we just give up? Is it over? Did the bad guys win with this? Or is it doable? Are we seeing defenses arise as well? Like, when we practice with our red teams, are we like, ah, I got an idea, let's do the thing? Is that true?
HEATHER ADKINS: I think we're on a glide path. I think we can see the horizon. We're starting to see the tools show up. This is where the white paper is helpful, to highlight what we're actually seeing. So it's real. It's coming. We should not be surprised by it. We shouldn't be scared by it. We should be thinking proactively.
The good news is that a lot of the good hygiene that we should be practicing will be helpful here-- making sure we're doing software upgrades when there are vulnerabilities, making sure we're putting in strong two-factor authentication for users. At Google, we use hardware-backed security keys.
And in the infrastructure, make sure we're configuring it properly, that we're not leaving vulnerable configurations open to the internet. The hardened perimeter still has value. And we've been talking about, you don't want to have a bonbon security model, where it's only hard on the outside and soft on the inside.
JORDAN GREENBERG: Hmm. That's a good one.
HEATHER ADKINS: Yeah. Having controls on the inside are also very important. Castle analogies are also very popular in our field. A castle is not just a building with a wall around it. You have a moat. You sometimes have multiple layers of walls. Those walls are usually pretty thick. And they usually extend down into the ground so you can't tunnel underneath them, or sapping, I think, as they call it in the official term.
And then you've got the keep. And the keep is where you keep the crown jewels. And the crown jewels are on the top floor. And there's usually no stairs. And so if you think about it, you've got these layers of defense in a castle.
The more layers of defense you have in your infrastructure in a traditional security model, the harder it will be for even agentic AI to go hacking. And the really fun thing is that if you've ever

played around with LLMs for any kind of information research, very happy to send you off into a very unproductive part of your research.
Which means that agentic AI hackers are not going to be perfect. And they're going to spend a lot of time trying things that don't work, just like humans do. But their intuition for it will be a little less predictable. So the more barriers we can put up while we are rethinking how to defend modern IT enterprises is going to be helpful.
And then I think there's also key technologies that get put in place. And here, I'm taking some inspiration from the biological sciences, where pharmaceutical research is looking at nodal biology, which is, what's a problem that happens in a human that leads to 30 diseases? Because if I can solve the one problem, I solve 30 diseases. That's nodal biology.
And it's very similar for how we build networks. What's the one node that if I harden that node, makes everything hard? Access management, for example, might be one of those things. And so if we can really get hyper-focused on agentic AI at the nodes, that's an interesting thread of research that we can pull on.
STEVE MCGHEE: Well, the thing that I think this reminds me of, relating it to reliability, as well, because it's just the idea of your surface area when it comes to vulnerabilities, not necessarily for attacks, necessarily, or infrastructure failure, but just churn, just change.
One thing that I worry about is, we're going to start changing Google faster. And we already are. Just, more code is getting written by machine-assisted humans who are perfectly attentive to what they're doing and stuff. But just the pure volume seems like a problem.
It's just going to be harder and harder and harder as it gets faster and faster and bigger and bigger. Same plan? Different plan? What should we advise people on here? Because I think it's probably not just Google that's facing this challenge.
HEATHER ADKINS: The first thing I would say is, however you think we are going to use AI is wrong.
STEVE MCGHEE: Darn, OK.
HEATHER ADKINS: And this will change over time. But the analogy I love to use is, when we harnessed electricity for the first time, the military planners thought, this is amazing. We're going to build these weapons, and we just zap people on the battlefield. It's going to change military warfare forever.
But that's not what happened. What happened is they put electricity on the railroad lines. And it increased supply chains, so it could increase the depth of your campaigns and the depth of occupation. And so when I think about AI and how it's going to get used, we have all of these things people are trying right now.
We have vibe coding, of course. But we also have people trying AI to write patches. There's a really great debate going on in our field, actually, as we're recording this, about the value of AI-generated bug fixes. Should you accept them blindly? Should you still do a human review? And I would say, every week, these debates are changing. And where they evolve to in two years, we don't know. So what I would say to organizations right now is, figure out how to play safely within boundaries, doing lots of pilots, lots of experimentation. Just know the first 600, 700 experiments may not be actually what you, as a business, or an open source project, or a project team, settle on.
You have this amazing tool to experiment with. Experiment, see where it is useful, abandon it where it's not. Don't be afraid of saying no just because it didn't work out. And just know that in two years, if you were to rewind to November 11, 2025, when we are recording this, probably look really different today than it did two years from now.
JORDAN GREENBERG: And hopefully in a good way. Hopefully, it empowers us, like the internet or electricity. And we're able to better serve users by having things fixed faster, by

having systems increase their reliability because something was found with an agentic helper analyzing things with you. That's the future we want.
HEATHER ADKINS: I remember the first security incident we had at Google was, we used to have these company-wide ski trips every winter. And a few of us would stay behind to keep the website on. And it was me, an SRE, who's still at Google also, and just a handful of other people. And we experienced our first major denial of service attack.
And we were just sitting on piles and piles and piles of packet data trying to figure out what was going on and handcrafting Python scripts to process the data to see where it was coming from. And actually, all of the toil and struggle with that incident was just data crunching, just to make sense of what was happening so the humans could actually make some decisions.
And if you think about the role of large language models, we could just probably put all that in there. And two seconds later, it would say, here's a cluster of activity, do this and do that. Amazing. Now, the humans know exactly what they need to do. And we can fulfill our value. We've taken the toil off of their plate and allowed them to actually go put in a tactical fix.
But then also, you've got all this time to think more strategically, what's a better long-term fix? How would we deal with this as a more systemic class of issues? Which we did eventually go on to build. But I just remember, a couple of days just doing nothing but Python and reading logs. And it wasn't very fun.
STEVE MCGHEE: Yeah, I know that some of my experiences with AI recently has been looking at logs and just being like, Gemini, what--
HEATHER ADKINS: Yeah, exactly.
STEVE MCGHEE: --what is this saying? And it's pretty good at that. If you could just have a little angel on your shoulder that's like, what that means is the following in English, that's actually a pretty generic good use of compressing logs into actionable insight, which, I don't know, sounds very business school, but it's useful. It's good.
HEATHER ADKINS: Also just presenting you with the questions you should be asking. Sometimes, you're not asking the right questions. So that's another thing. What questions should I be asking? Gemini is pretty good at that too.
STEVE MCGHEE: Totally.
JORDAN GREENBERG: Yeah, so with all of that said, where can we find you on the internet, protecting the internet?
HEATHER ADKINS: I do keep a relatively low profile, but I am on X and on Bluesky, ARGVEE. And I'm also on LinkedIn. And I tend to post about cybersecurity. And I'm very close on geopolitics. These are very tightly coupled these days. And I'm very passionate about medieval European history as well. So you'll see some tweets.
STEVE MCGHEE: Awesome.
JORDAN GREENBERG: Oh, so the castle information makes sense.
STEVE MCGHEE: I like it when it crosses over. Yeah, that totally works. I love it.
JORDAN GREENBERG: Amazing.
STEVE MCGHEE: Well, thanks so much for your time. This was really cool. It's nice to see you again.
JORDAN GREENBERG: Thank you.
HEATHER ADKINS: Thank you. Good luck to everybody.
JORDAN GREENBERG: Awesome.
You've been listening to *The Prodcast,* Google's podcast on site reliability engineering. Visit us on the web at sre.google, where you can find books, papers, workshops, videos, and more about SRE. This season is brought to you by our hosts Jordan Greenberg, Steve McGhee, Florian Rathgeber, and Matt Siegler, with contributions from many SREs behind the scenes. *The*

Prodcast is produced by Paul Guglielmino and Salim Virji. The podcast theme is "Telebot" by Javi Beltran and Jordan Greenberg.