

[THEME MUSIC] Tele telebot

You missed a page from Telebot

JORDAN GREENBERG: Welcome to season seven of *The Prodcast*, Google's podcast about site reliability engineering and production software. This season, we are exploring what it takes to ship software and what that even means to different folks. Join us as we discuss how to ship it as SREs.

And of course, the most important part is the friends we made along the way. Happy listening and may all your incidents be novel.

(SINGING) You missed a page from Telebot

JORDAN GREENBERG: Welcome to *The Prodcast*, Google's podcast on site reliability engineering and production software. A flawless intro this time, though I've flubbed it many a times, many a times.

STEVE MCGHEE: Well done.

JORDAN GREENBERG: My name is Jordan Greenberg. And I have my co-host.

STEVE MCGHEE: Oh, I'm Steve McGhee. That's me. Hi. Hey, Jordan. How's it going?

JORDAN GREENBERG: It's going well. It's going well. And we have also an awesome guest with us today. Will you introduce yourself?

WILLIAM HOLROYD: I am William Holroyd. I'm an SRE for regulated cloud.

JORDAN GREENBERG: Regulated cloud, that sounds familiar. OK, I would love to hear a little bit more about that. Steve, you want to take us into our first set of Q's?

STEVE MCGHEE: Yeah, William, I work in cloud. But what is a regulated cloud? First of all, what is "regulated" for the listeners who have maybe heard that word before, but not in this context? Let's start way back. Let's just define some terms. And then pretty quickly, we can get into what the actual cloudiness of it all is. How should we start?

WILLIAM HOLROYD: Yeah, sure. I think the public name for how this is phrased is "assured workloads" or "data boundary." Regulated cloud is kind of our org name internally.

WILLIAM HOLROYD: In a normal cloud, that's where you get to decide where things are run, who has access, how keys are managed, what you bother to log. But regulated environments are usually decided by a law, such as ITAR, or they're set by some sort of agency standard like CJIS, which is a Criminal Justice Information Systems, or some other standards attestation body. Most people are probably listening to this podcast would be familiar with SOC2 as an example or CSA star.

STEVE MCGHEE: Yeah, and I assume that this goes into other countries which have other acronyms as well? And then it just gets more complicated, I presume?

WILLIAM HOLROYD: Everybody loves their acronyms. GDPR for Europe. California has one. Canada is Canadian Goods Program?

JORDAN GREENBERG: Yeah, the Canadian Controlled Goods Program, where actually, I got to work with you on.

WILLIAM HOLROYD: And then Protected B.

JORDAN GREENBERG: Yes.

STEVE MCGHEE: So then what is it that they're regulating? Is it the location that you run your VM in, or is it lots of stuff? I guess, what are the controls in a broad sense? Well, let's get into the why in a second. But first is the what. What are they controlling? Are you not allowed to use computers that are of a certain color or anything like this? How detailed does it get?

WILLIAM HOLROYD: Usually there's four things that are defined in these. Not necessarily with SOC2 or business level, business-to-business attestations, but when you get into the ones dictated by agencies and law, you have residency, which is where the data is allowed to live. You have access, which defines the people and the roles that can reach it.

And then you have key custody for encryption keys-- who can hold them, who can even access them to be able to decrypt the data? And then evidence, being able to demonstrate that all of that works to an auditor on demand with artifacts.

STEVE MCGHEE: I think I understand what it means to be regulated. And then the why around regulation is probably because laws. That makes sense to me. And different laws in different countries, and you want to make sure you're following the right laws in the right place. That all sounds pretty good.

So as just a normal person on the internet, will I ever encounter one of these? Do I know when I'm using one? Does it feel different? Or is it kind of transparent to the end user?

WILLIAM HOLROYD: End users, it'll be transparent to them. They're not going to know that there's a compliance regime or some other regulated control set behind the scenes. Most people will probably understand that it's there, depending on the industry that they're in. So for example, if you work for a company that's already in an industry, they're already aware of what those boundaries are-- for example, payment data, health records, working with a government agency, financial services, personal data, belonging to people in a specific jurisdiction-- so the EU being the more obvious ones-- if that's your area, you already know that they're there behind the scenes. But from an end user perspective, it should be entirely invisible to you.

JORDAN GREENBERG: What does it mean for a person-- everyday SRE, everyday engineer, product manager, TPM, et cetera?

WILLIAM HOLROYD: I would say that most people listening, not working on a compliance project, there's a few different ways to think about this. If you're an SRE, the controls that you would need to apply to your resources through organization folder project are actually enforced by the platform. If we're talking about GCP as an example, it's not a process.

So thinking about it from an end user perspective on our platform, when you create an assured workloads or data boundary enabled project for a specific control package, the controls in that package become your guardrails for everything that's within it and also enforced by any organizational policy constraints.

So it's not on a Wiki page somewhere. It means that the resource in that package enforces restrictions, like certain containers don't get created and regions that shouldn't be there. So for certain US regimes, control packages, you're not allowed to build anything in Asia or Europe, as an example.

STEVE MCGHEE: So these are more like hard guardrails instead of suggestions? They're actually embedded in the system.

WILLIAM HOLROYD: Correct. And there is a shared responsibility piece to this as well. So there are certain things that, as a customer, you're not supposed to do as well. So working around the guardrails would not be a great way of adhering to that.

JORDAN GREENBERG: That seems right. That seems right.

WILLIAM HOLROYD: Yeah, I'd say one other impact is, for an SRE, it's probably one of the teams that's more likely to log a bug with Google support if you have an issue. And there are products, as an example, that you would see some of the power of these compliance controls come to life.

So one would be access transparency. Another one is admin access. We can talk about that one later, in depth, if you want. But those will actually enable you to see who actually interacts with your project, your organization or whatever your resources are, from the Google side. And depending on your organization settings or resource settings, you can actually enable them to interact with those resources if it's absolutely required.

STEVE MCGHEE: So you're saying that as a customer who's using regulated cloud-- let's say

I'm a bank or something, or a healthcare system of some kind. I can potentially, for some parts of the cloud or some products, I can see who at Google has access to those systems or who has touched those systems? Is that a feature?

WILLIAM HOLROYD: Yeah, you won't see who in Google has access to the system. It's based on a business requirement, a business justification to be able to interact with the resource. So for example, if you log a bug for a product-- I'm going to pick one here and hope I don't get emails from them.

But Cloud SQL is an example. Somebody opens up a bug with Cloud SQL, ends up getting escalated within the company. An engineer's like, yeah, I don't think we have something that's hit this. We've never hit this issue before. It would be interesting to see what the logs with that database might be, and I need to receive them.

But if they are in one of these control packages, they're most likely going to be locked out of being able to access it. So a FedRAMP environment requires us to have background checks done as part of our work. So in that particular instance, only people with that would be able to interact with your organization. So they could actually submit a request and say, this is actually in relation to this bug you opened with us. And then from there, you can approve it or just say, I don't want you to, and deny it.

JORDAN GREENBERG: And then on the flip side of that, is there ever a reason where a customer would want them to access this workload? Maybe they would give them that specific reason and say, can you please jump into my workload to help me?

WILLIAM HOLROYD: Yeah, so the way we think about things are like data classifications. So I think a lot of people think that if I click this button, they're going to have access to all my data. And that isn't necessarily the case. Just being able to access or interact with a resource requires approval.

The settings you put on an instance, on a resource, are just as sensitive to us as it is with any data you might be actually storing in the database, as an example. So it might just be a configuration on the SQL server, not the data in it that may need to be accessible. So that would probably be the mitigating factor there.

JORDAN GREENBERG: We've heard about how engineers might interact with this. What about if you are a manager of engineers in this space?

WILLIAM HOLROYD: Yeah, so as a manager, I think trying to protect your team from toil is kind of a big thing you'd really want to think of. Anything you can do to automate away toil, anything that kind of makes movements forward in technology upgrades and so on would be of good advantage to you and your team.

From a manager, one thing that we do help enable is as control drift becomes part of a life cycle with some of these control packages. So it's not something where it's like FedRAMP High has adjusted what it means to be FedRAMP High, and now we have additional requirements to adhere to them.

So for us, those control packages are versioned. They'll be improved over time. And we actually will provide the ability for you to actually update your workloads in place and compare it against your currently deployed resources, and then apply the upgrade for the pieces that we own of that control.

So support personnel controls are probably another concern as well in terms of things that you might have to answer to in a regulated environment.

WILLIAM HOLROYD: But that'll route your cases to personnel meeting those package requirements, whether that's a geographic restriction or even background check requirements. So to give you a more concrete example, FedRAMP High means that the person-level support are located in the US, and they have the enhanced background checks, and that they have

gone through the necessary training for it. So they're fully aware of what the impact of that particular compliance package is and what it means to the customer.

STEVE MCGHEE: I know this is going to be a drastic oversimplification, but I'm new to this. So to me, this sounds like this is like a implementation of like a RACI chart. So it's like, who can do at all times for which pieces of data from which countries?

WILLIAM HOLROYD: As a product manager, the thing that you would probably want to internalize is our supported products list. So product support varies by control package. It's no different from us and any other cloud provider. They have their limitations on which ones will adhere to the different compliance regimes.

That list is somewhat a boundary of what your addressable market is for all of your regulated customers, whether what it is you're trying to ship is FedRAMP High, IL5, ITAR, the EU Data Boundary. Even the Canadian ones like Protected B or the Controlled Goods Program-- I got it right this time-- determine what you can adopt for those customers in trying to deliver your services.

I do want to point out that I said internalize. I would not try to memorize it. That list continues to expand. We do continuously expand the portfolio. I believe we put out blog updates for all of them. I've seen my fair share of the updates on some of the things that I've worked on, which is good to see.

And probably keep in contact with your account manager. Maybe that's more of an engineering manager concern. Really depends on who actually owns the relationship with Google and letting them know what your actual needs are so you guys can be in alignment with those releases when they come up.

STEVE MCGHEE: Makes sense.

JORDAN GREENBERG: That especially, I think it's like sometimes it has to move at the speed of the customer. Sometimes it has to move at the speed of the development team. What can they reasonably implement in x amount of time? And this is very much a conversation that has to be had and moved around, depending on where both of those ends are. I think that that's a really good way to say it.

STEVE MCGHEE: So given all that, it sounds like you've applied a bunch of rules based on a bunch of laws to a bunch of systems. And that sounds like an awful lot of work. Are you guys done? What's in the future? Is it just a matter of spreading this across all of the domains, and all of the laws, and all of the regulations, and then keeping up with them over time? Or what do you think is-- what's next in regulated cloud world for you?

WILLIAM HOLROYD: The short answer is technology, you're going to fundamentally change every couple of years--

STEVE MCGHEE: Sure, yeah.

WILLIAM HOLROYD: --as we come up with new ways of improving things. I don't see regulated clouds, the compliance regimes change from that at all. It's probably going to follow that same model. I'd say I'd probably like to start by separating my opinion from the company for a moment. Hopefully nobody at Google has any concerns with this one.

But for 20 some years, once the cloud came around, we were sold this particular idea that the plan is just one giant data center. You can replicate everywhere. You can failover anywhere. You can have one ops team in the US, wherever, and follow us on rotation, maybe in a couple different locations, backup to whatever region was cheapest.

But this was all based on assumption that the trust between the countries exists. And I don't think that is safe to assume anymore. Watching the news, you can see that governments are kind of reevaluating who they depend on. And every time that happens, there is some new set of lines that get drawn on the map. I actually think we're at the beginning of this, not the middle.

My practical advice to deal with this is that if you think you'll be in a regulated market sometime in the next three years, you would want to design for your boundary now.

It's far cheaper to build inside one than to move into one. You don't necessarily have to apply FedRAMP controls. You can go read the documentation in terms of how the environments change or what comes with that. And you can build it. And then that changes aspects on our side as well. It becomes pretty easy at that point. That's how I see things. So yeah, there's a lot that comes with that.

JORDAN GREENBERG: So what does it mean for an engineer who has to think through these problems? OK, it's probably a good idea for me to start thinking about how I want to implement these boundaries early. What does the plan look like for an engineer in that way? What should they think about?

WILLIAM HOLROYD: Yeah, so in terms of trying to avoid that regulation tax of building later, this means you need to start thinking of that additional dimension in your architecture that most people don't have today, unless you're in one of those industries we were talking about earlier, which is the jurisdiction.

You are already currently reason about zones and regions for availability. Jurisdiction is the level just above that. And it behaves entirely different because it's not a failure domain. It's a legal domain. A region can failover to another region. A jurisdiction has hard walls, unless the other side of that fail is also inside the same legal framework.

So like GDPR is multiple jurisdictions in this super jurisdiction, as an example. So the first design question becomes something like what is your unit of isolation in terms of where you can self contain your entire platform within a jurisdiction zone-- its own data replicas, its own backups, its own DR targets? And if you have any sort of cross jurisdictional replication that's not a safety net, even if you encrypt it, that is still a liability. So you have to think about how to self-contain in those jurisdictions.

JORDAN GREENBERG: If we consider who the audience is for our engineering teams, a lot of them are going to be SREs and SRE adjacent. So I think what would be best is focusing on what does an SRE, what does an engineer want to know? What do they need to implement? What should they focus on? And what should be something that one of the other roles that we were talking about earlier-- engineering manager, or product manager, or account manager even-- what should they be focusing on when they work on this architecture early?

WILLIAM HOLROYD: I would think it would be your data model. This kind of applies to existing solutions as well. So it's like if you can't move the data, you kind of need to know what the data is. You need to know which records belong to which jurisdiction, which fields are sensitive, where a piece of data is even allowed to be processed, not just stored, as an example.

Date of classification in terms of what the various fields or types of data that you have really needs to be an attribute in your system somehow, and when your system comes across this attribute, how it's supposed to handle that processing differently.

In terms of the engineering team-- I've been talking about engineering things that I have to put into the technical stack, but the operations model itself kind of changes too. So everything I said about data applies to the people just as much. So a single, global on-call rotation with route in all of the different regions that you have, all the different jurisdictions, that's the human equivalent of global replication.

If the data can't cross the line, then at some point, the hands on it can't either. So it's like planning for regional operation cells or various teams in different regions in terms of who can deal with the data in that jurisdiction is important. And to make that easier, it's building the tooling so that operating one of these cells, if you will-- that you can't fully see or fully operate is normal and not an exception.

JORDAN GREENBERG: Gotcha, OK. So earlier you were talking about some of those little acronyms, those funny numbers and letter combinations-- SOC2, GDPR, Assured Workloads, AW, CCG-- do you have any favorites, if you had to pick a favorite compliance regime?

WILLIAM HOLROYD: By what measure.

JORDAN GREENBERG: Well, let's think of it from a business context. Where do most people implement and then stay there for a while?

WILLIAM HOLROYD: So from a compliance regime, it would be very deliberate to have to move into that. I think most people, if they get into SOC2, I think that's probably a good starting point if you're not there already. I think that's the gold standard, just to be able to sell business to business. I think that's essentially where you would need to be to enhance your business.

There are other things that you can do that you can achieve along the way. I think CSA Star's requirements are fairly in aligned with SOC2. So if you set that distinct goal of SOC2 and work towards it, there are other attestation frameworks and security compliance regimes you'll pick up along the way. But moving beyond that into the acronyms that we just mentioned, which are government level, that is deliberate. That's where your business is signing a contract somewhere.

JORDAN GREENBERG: Gotcha.

WILLIAM HOLROYD: But not necessarily required.

JORDAN GREENBERG: And so when we were talking about this earlier, we were thinking about if you expect to interact with government agencies or regulated agencies in any way, you should start building for those regimes now. So how can an engineering team think through what they need so that they can plan effectively? Because a lot of these compliance regimes sometimes get very complicated to implement. So where should they start to plan how they can get there?

WILLIAM HOLROYD: That's a great question. There's a lot of different pieces that come into play. Like I mentioned, there's the operational cost, depending on where you are going to live in the world and the various jurisdictions where you need to hire. Even what their citizenship status or work authorization is needs to be taken into account, which is probably the one thing that most people forget about, depending on which compliance they're coming into.

It affects everything across the board, not just the cloud environments like I was talking about.

There's the software piece. FedRAMP actually has you go back and look at any place you could use any sort of hashing or cryptography, and you will have to change it. So there's a software impact there that might be actually part of your data model, depending on what it is you're actually storing and how you're storing it.

You might have to re-encrypt all of your data if you're bolting on later. There are things along the way, I think, in SOC2 as well, that even dictates how things get to production. So devs not having access to production, automating most things, there's a release management aspect of getting things into production need to be taken into account.

How those services talk to each other, not necessarily intra-data center or client connectivity to the services, but how they're even interacted from, I guess you'd say, an internal boundary, so even from an employee just trying to manage the platform.

I think a lot of people, thinking about what we've talked about, it's more of a customer comes to us, and they want us to make a promise on how to actually store that data. And a lot of it is just trying to meet the law. But you also have to think about it from an insider risk perspective as well. So what controls do you have internally that prevent people from doing things that they might not otherwise have access to? So as you build up these controls, your surface for attack is going to change. So you have to think about that aspect of your implementation as well.

STEVE MCGHEE: This is a lot. This is a complicated set of things. It sounds like there's a lot of

security wrapped up in this as well. So have to consider your security posture, how it relates to your data format, I guess, and just your data model, and then how that all relates to the country and rules and industry. It's a lot. I guess I'm just saying it's a lot.

So how do we help? How does Google help? What are the things that we offer? Clearly, it's not just, here's a firewall, good luck. There's probably more to it than that, I'm guessing. What are the different products that we-- I mean, we don't want to just enumerate all the products. That's no fun. But what are your favorite things that Google offers that you think are genuinely cool, or impressive, or helpful, or things like that. Here's your chance to shine. What do you think is going on here?

WILLIAM HOLROYD: I'll go back to what does Google do well.

STEVE MCGHEE: OK, great.

WILLIAM HOLROYD: I think that was pretty easy. So going back to the comment about the four columns, I guess you'd say for this-- the residency, access, key custody, the evidence-- we've tried to make that as easy as possible for our customers. I've gone through FedRAMP before on another cloud provider. And it was its own development pipeline, its own project.

And so for us, we've developed things like Data Boundary to exist so that answering those four questions is partially answered by control package. You just apply it to the organization, the project, the resource, the folder, depending on the regime. And you just get that for free. So it doesn't have to be manually built over your long projects.

When it comes to access and the evidence pieces of those four, we have a couple of products. One is called Access Transparency, the other one being Admin Access. Access Transparency is what enables customers to see when Googlers interact with their customer accounts or resources.

But like I mentioned earlier, it has nothing really to do with data most of the time. It's just who is interacting with your resource based on how we classify access to an account. For example, just looking at the account sends an org level alert of who is trying to access what resource, even if it was a failure, as an example.

WILLIAM HOLROYD: And then that just shows you who was trying to access. So that's let's say it's a valid reason for somebody to access your account. Our other feature would be access approvals, which enables customers to block or approve the engineering specialist to help resolve the issues if they need to go down that route, if they actually need to interact with whatever the resource is.

And I think this is the area which I think is the most fun for me, at least in the org that I live in. We really have a solution for every type of sovereign level deployment that someone is interested in.

We have our typical cloud regions that everybody knows and love. We have dedicated clouds, like the one operated by SANS in France. And then we have distributed clouds that can live at the edge or even in air-gapped environments. So we really cover the ground for anything that you might need from a sovereign deployment perspective.

STEVE MCGHEE: That's awesome.

JORDAN GREENBERG: Multiple flavors, regulated, dedicated, air gapped, lots of different types of clouds, just like real clouds, actually. I never really thought about it.

JORDAN GREENBERG: So thank you so much, William. Where can we find you on the internet?

WILLIAM HOLROYD: LinkedIn is pretty much the only real public place I have. I have a Twitter account, but I don't really use it.

JORDAN GREENBERG: That's easy enough.

WILLIAM HOLROYD: It's William Holroyd on LinkedIn.

JORDAN GREENBERG: Perfect. We'll make sure to include that link in our podcast documentation. So for now, thank you so much for listening. And thank you, William, for joining us. I really appreciate learning all about the different letters and numbers and types of clouds that we get to work in now, so appreciate that.

WILLIAM HOLROYD: Absolutely.

STEVE MCGHEE: And as always, let the regulated pagers be silent, I guess, right?

JORDAN GREENBERG: Yes, every pager.

STEVE MCGHEE: All the pagers. But also those ones.

JORDAN GREENBERG: All right, well, thank you, listeners.

WILLIAM HOLROYD: Cool. Thanks, William.

JORDAN GREENBERG: Have a great day.

WILLIAM HOLROYD: Thank you for inviting me.

JORDAN GREENBERG: You've been listening to *The Prodcast*, Google's podcast on site reliability engineering and production software. 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 hosts Jordan Greenberg, Steve McGhee, and Matt Siegler.

The Prodcast is produced by Paul Guglielmino and Salim Virji, with contributions from many SREs behind the scenes. The podcast theme is "Telebot" by Javi Beltran and Jordan Greenberg.

[THEME MUSIC] You missed a page from telebot