EPISODE 1442

[INTRODUCTION]

[00:00:00] KP: Compute resources continue to trend towards being cheaper, easier to use and faster. Despite these positives, more compute demands more energy, and therefore an increasing carbon footprint. With many companies committing to controlling their net carbon emissions, tools are required for engineers to not only measure their cloud infrastructure, but to make informed choices about how to set up that infrastructure. The operations of Google Cloud are already net zero emissions. The carbon emissions we're going to be talking about are the gross carbon emissions.

Steren Giannini is a Senior Product Manager at Google. He works on Google Cloud Carbon Footprint. In this interview, we discuss how Google Cloud collects and exposes this data to users and how those users are responding in their infrastructure approaches.

[INTERVIEW]

[00:00:56] KP: Steren, welcome to Software Engineering Daily.

[00:00:59] SG: Thanks for having me.

[00:01:01] KP: Can you tell us a little bit about your experience in the world of software?

[00:01:05] SG: Sure. So I'm, today, a product manager at Google. I've been in the software industry for 15 years, I should say. Originally, I'm an engineer by training, then I joined the product management side. But keep a very technical hand, because my products are Google Cloud products. So they tend to be very technical products. So myself, I'm more of a web developer, even though I've done some Android and 3D rendering development in the past.

[00:01:37] KP: Very cool. Well, the product we're going to talk about today is carbon footprint, which I think the users will have a general understanding of what that might be from its name.

I'm curious, is this your first foray into doing anything environmentally related? Or do you have a connection there as well?

[00:01:55] SG: No. That's a good question. I actually was looking at, “Okay, what action can I take using my software engineering skills or product management skills to help fight climate change?” And as I was starting to brainstorm things I could do on the side, I then realized, “Wait, but I'm a PM at Google. Maybe I could do something at Google.” And so that's where we got started with helping Google Cloud customers understanding their gross carbon emissions and helping them make more sustainable choices for their cloud infrastructure.

[00:02:31] KP: Can we give a formal definition? What is the carbon footprint project?

[00:02:36] SG: Yeah. So Google Cloud Carbon Footprint is a product name. But maybe let me step back and maybe define what do we mean when we say carbon? So carbon is a shortcut to say to talk about the greenhouse gas emissions that are responsible for climate change. And they happen to be different gases that are often reported in equivalent grams of carbon dioxide. So this is where the carbon comes from. Basically, when we say carbon, we mean greenhouse gases responsible for climate change.

And now, as part of Google Cloud, we basically have helped customers with two things. The first is helping them understand what is the impact of their choices and help them make better choices, notably regarding which Google Cloud region to deploy their infrastructure into. And the second area is helping them understand and reduce their gross carbon emissions linked to their Google Cloud usage.

So the first one is really about which location to pick when you deploy Google Cloud. As you might guess, Google Cloud, like many other cloud providers, have a very large set of data centers all around the world to pick from, where one can deploy a virtual machine, can store data. And all of these are located on different electricity grids. And those electricity grids, to produce this electricity, more or less, carbon emissions are emitted. And as a consequence, and depending on where you decide to deploy your infrastructure, you might indirectly emit more or less gross carbon emissions.

[00:04:30] KP: So this is a really fascinating idea that actually was first introduced to through the opportunity to talk to you. I've been in the position to choose a region before. I'm in the United States, and pretty much the things I've built have targeted the United States. So even though I've got a couple of regions, I've always just said, “Let me pick mine or see if one's a penny cheaper and go that way.” What are the typical motivations people have for picking their region who aren't yet thinking about the carbon impact?

[00:04:59] SG: That's a very good question. So typically, you have to consider a few factors. The first is an obvious one, is maybe you have some data locality constraints. Maybe in the area you're in, your data must be in a certain location. That's the first point. The second point is latency to your end users, right? If you're offering a web service, you want to put your servers as close as possible to your end users.

Another one is the price, as you mentioned. The price of a virtual machine. And depends for many cloud providers on where you purchase that virtual machine. And lastly, and this is what we've helped customers with, as I said, not every location has the same carbon impact. So picking the right region for your workload depends on those three dimensions. And often, it is something that you have to optimize depending on what you want to put forward. Sometimes it's price. Sometimes it's carbon. Sometimes it's latency. Sometimes it's a mix of everything. And so we've published some data regarding the last part, the carbon impact of Google Cloud regions, so that you can take more informed decisions considering all of those factors.

I've actually built a very small tool that has like three sliders, but you can say, “Okay, I care about latency and carbon, but not about price. Or I care about carbon and a little bit about price, but not about latency.” And the tool gives you the regions that we suggest you to use.

[00:06:31] KP: Very neat. I'd love to delve more into that tool. Can you describe maybe with those three variables, is there a way we can visualize the underlying distribution? How do the interactions take place across this choice?

[00:06:43] SG: It's a tool that uses open data. And so the data that it uses is, for the latency, it considers – It simplifies by considering the distance between where your users are located that you can define in that tool and where the region is located. Regarding price, it uses the public

prices of, I think, a compute engine VM. And regarding carbon, it uses the data that I mentioned that we have recently published and opened to everyone. So this data has two numbers. One of them is called a carbon-free energy percentage. And the other one is the green carbon intensity.

So maybe I can dig a little bit deeper into what those numbers mean. I think, notably, the first one need to be explained. So what do we mean when we say carbon-free energy percentage? So overall, Google invests in enough renewable energy and carbon offsets to neutralize its operational carbon emissions. So you should think about that as Google Cloud is net zero for its operational emissions, operational carbon emissions. That's net.

What about gross carbon emissions? So basically, gross carbon emissions are that, in a certain location, at a certain hour of the day, the energy might come from only carbon-free energy sources, or might not. So when it comes in a given hour, when it comes from carbon-free energy sources, we consider that, “Okay, this is a carbon-free hour.”

So basically, you can think of this carbon-free energy percentage as what percentage of the electricity comes from carbon-free sources on an hourly and original basis. So you see how the net zero is really at a global scale. Like globally, at the end of the year, Google operational carbon emissions are net zero. But at a local scale, in a given region, in a given hour, then the gross carbon emissions might not be zero.

And so with this number, this carbon-free energy percentage number, you can basically understand what percentage of the time the region is running on carbon-free energy sources. So if I take an example, a very concrete example, let's say you are running a compute engine VM in Las Vegas. Las Vegas has a carbon-free energy percentage of 19%. If you were to move this VM to Oregon, Oregon, as we've published it, has a carbon-free energy percentage of 90%. So, 90%. Then you can clearly see that you would go from roughly 20% of the time running on carbon-free energy sources to 90% of the time running on carbon-free energy sources. So basically, you would really reduce your gross carbon emissions just by moving from Las Vegas to Oregon. So you stay on the West Coast. You stay in the United States, and you massively reduce your gross emissions. And by the way, you also reduce your bill. Because it happens that Oregon is cheaper than Las Vegas. So this is what this tool can help you with.

Basically, sometimes, our customers don't really have insights into what's the best choice regarding price or sustainability. And this tool can help you with that. But of course, the tool is a dedicated page, a dedicated tool that we have created. But we've also worked to add very small labels next to the region selectors of the Google Cloud Console. Or even on the Google Cloud documentation, when you have a locations list, you can see a very small leaf and a small low CO2 indicator next to the regions that have the highest carbon free energy percentage or the lowest grid carbon intensity.

And so, basically, what we did with that small label is, “Oh, you don't need to understand those numbers. You don't need to understand what carbon-free energy percentage means. You don't need to understand carbon intensity, which is the gram of equivalent carbon dioxide per kilowatt hour consumed.” No, no. You don't need to understand all that. You can trust that if you see this small leaf and this small low CO2 label anywhere on the ducts or the user interface, then basically that means that this region will have low gross carbon emissions. And you can – Hopefully, by picking this region, you will have lower impact.

And by the way, we've tested this change. So before rolling out the display of those small, low CO2 labels in the ducts or the user interface, we perform an A-B testing experiment, right? Because this is what we often do when we do a changes like that.

And what we've observed is that after running the experiment for enough time, was that 19% of users were more likely to pick a low carbon region when we were displaying this small label. And actually, if we only looked at new Google Cloud users, we noticed that 50% of them were more likely to pick a low carbon region when this symbol was displayed on the user interface. So basically, we were very satisfied with those results. And that made the point of rolling out this change all over the Google Cloud Console and the Google Cloud documentation pages.

[00:12:36] KP: Yeah, I wonder if we could pause on that for a moment. I mean, effectively, you're saying a large percentage of users, given this information, use that to inform their decision. I don't have the luxury and a lot of products I interact with of getting an ecofriendly version? I think I'd take it if it were there. Do you see the opportunity for this to become a broader standard in other industries?

[00:13:00] SG: Yeah. I think, definitely, within Google products, you have seen this small leaf popping up a bit everywhere. And it is actually the exact same symbol that we use all over the products surfaces. So let me give examples. If you use Google Maps and do a route search, then you will have multiple routes suggested to you. And the one that is the most sustainable will have this small leaf associated to it. If you use Google Flight, when you see the list of flights for a given destination, you will also see the ones that are most sustainable. And again, the label is the same. It's the small leaf. And in case of Google Flight, it gives you basically which flight has the lowest emissions.

So you can see how within Google products, this pattern appears where it makes sense. And then I hope we will see that more and more in more products, because we know that those micro optimizations can have large impact. And maybe let me go back to cloud and picking a region. I'm convinced that after moving to the cloud, picking the region with the lowest carbon impact is probably the most impactful change compared to the effort required that customers can do.

And if we show the information in the right place at the right time, then this can have a lot of impact, right? If we make somebody aware that they should deploy to Oregon, instead of deploying to Las Vegas, because the gross carbon emissions will be lower in Oregon, then that's all we need to direct these users towards the greenest option.

And so, yeah, picking the right region is probably the most impactful change that customers can do, of course, have to move into the cloud. Because I can talk about these two if you want. But there is also a big difference between the gross carbon emissions of running on Google Cloud and the gross carbon emissions of running on-premises, for example.

[00:15:09] KP: Yeah, I'd love to get into both. But I've got one follow up for you about the example of moving our VM from Las Vegas to Oregon. You already read over the numbers, we went from something in the teams to 90% on the carbon, which is huge improvement. The price also, as you'd mentioned, goes down. What about that third variable of latency for people who live in Las Vegas and have enjoyed a close server? What's the new experience going to be in terms of millisecond? Sir, on average?

[00:15:36] SG: I mean, that's a good question. Like, if your goal is to offer some kind of real-time video against training platform to Las Vegas residents, then, yeah, you probably want to optimize to be as close as possible to them. If your goal is to offer a website, then Oregon, Las Vegas, at the speed of light, I don't think there is that much difference. Also, your users are probably not only coming from Las Vegas. They are coming from all over the West Coast. S

And so, this is why it won't have much impact. Of course, we are not talking here about crossing the entire globe. We are talking about a small original change within the same area of the same country. Now, there are also some workloads, by the way, that do not care about latency to end users. Like, everything that is about data processing of a data set, the data set is located in a certain region that you can also use our numbers to decide where to store your data. And then processing that data is not latency-sensitive. It doesn't need to happen in order to serve a user request. Therefore, this can be done not only in the right place, but also at the right time. And this is something that we do not really communicate about. But basically, to come back to what I said earlier, the location and time of the day determines how much gross carbon emissions are emitted.

And so if you, for example, run your batch jobs when the sun is shining, the carbon intensity of the electricity will be lower at this time of the day. Because when the sun is shining in California, for example, then the solar panels are producing a lot of green electricity, or putting a lot of green electricity on the grid. Therefore, it's probably like you will have lower gross carbon emissions if you run your batch jobs when the sun is shining.

[00:17:43] KP: I'm sad to report this is sort of the opposite of what I've done a lot in my career. I've set jobs to run one minute after midnight.

[00:17:50] SG: We all do, right? But if you think about it, if the price is the same, you can look up the carbon intensity of the grid where your VM or your batch job runs. And you can see that, “Well, maybe I should run it one minute after noon.”

[00:18:06] KP: Yeah, or if it's a 30-minute job, I should run it at 11:45. So that it really covers the optimal points. There's sort of an optimization to be done there. Is that something that you see a lot of DevOps people considering?

[00:18:22] SG: No, because I think that's probably the self-step of what you want to do when you want to reduce your gross carbon emissions. If we think in terms of impact, like, step number one is move to the cloud. Google Cloud, notably, has a very low power usage effectiveness. So maybe for our audience, power usage effectiveness is basically how much energy you need to run the facility, whether the servers are running over, or how much energy is used for IT. And so the lower this ratio is, the more optimized, the more effective your data center is.

And Google has been committed to communicating about this for a while that, for many years, it had one or the lowest power usage effectiveness of the industry. And every year, we keep lowering it. So I think right now we are at 1.1 for Google Cloud Data Centers. So much lower than a typical on-premise power usage effectiveness. So basically, if you think about it, the first thing you should think of is using moving to the cloud. And by the way, as I said earlier, once you move to the cloud, the thing – Like when the operational emissions at the end of the day are net zero, because Google is already offsetting all of that for you. So okay, step number one, move to the cloud. Step number two, probably picking the region with the lowest carbon impact. And then step number three, like optimizing the time that things can be run at. But there are also things like trying to avoid wasting resources.

And here, maybe I can mention something that we've worked on. At Google Cloud, we have worked on analyzing what we call unattended projects. So there are some resources in Google Cloud that our customers have created and might have forgotten. And those resources, if it's a virtual machine, then this machine is located to you. And even if it's idle, it's still consuming some energy.

So what we've created using some heuristics and some machine learning is a system that tells you, “Hey, it looks like you've forgotten about this project. Would you like to delete it?” And when you do so, of course, you save money. But of course, you save carbon, because we are able to tell you this project has that amount of idle resources. And if you were to delete that project, you would stop emitting that much gross carbon emissions. So that's something that you can now see in the Cloud console. If you open the Google Cloud console, you can see those recommendations. And they tell you how much gram of equivalent carbon you would save by

following the recommendation. So avoiding waste is another action people can take. And then there are also a lot of like overall the architecture. How much energy consuming it is?

I think reading online, people claiming that moving to one programming language is more carbon friendly than another programming language. But honestly, that's the last thing you want to do, right? Like, moving your business from one language to another language is probably the hardest change you can implement.

[00:21:42] KP: And riskiest.

[00:21:43] SG: And it's risky. You aren't going to rewrite your entire business. Like, yes, it can have impact. Don't get me wrong. Like some languages are more low-level. They're more optimized. They might burn less CPU. And therefore, cost less money and emit less gross carbon emissions. But I think that that's not where you should start, right? You should start by those first steps, like move to the cloud, pick the greenest region, consider removing ways, and having a more optimized infrastructure.

[00:22:16] KP: Well, having an easy access in the interface to make these choices, it seems like a no-brainer. If I'm going to go do my next developer project on the weekend, that I'm going to make the ecofriendly choices. For an enterprise to do adoption, I'm sure there are just a few more hoops to jump through. What state of maturity do you find most customers are at at really taking these things into account as they make their decisions?

[00:22:41] SG: Yeah. So to recap, to come back to what I've worked on in Google Cloud, it's informing about the impact of the region selection and providing this Google Cloud Carbon Footprint product, which is a way to measure, report and ultimately help you reduce your gross carbon emissions linked to your Google Cloud usage.

So what we've seen is that we've seen a lot of demand for this second part, the carbon reporting tool that we have launched. So basically, for a very long time, we had a lot of customers asking us, “Hey, we need help to understand the carbon emissions that are coming from our Google Cloud usage.” Because those customers, they have requirements to report on their indirect

emissions. And when you use a cloud provider, you are using a service that itself emits emissions directly or indirectly.

And so our customers need that data because they have some reporting obligation. They want to know. They care about this topic. Or maybe they have made commitments, like Google. I've talked about carbon-free energy percentage. Well, Google has made a commitment that, in 2030, every location at every hour will run at 100% carbon-free energy resources. So that's a pretty strong commitment.

But all customers do commitments too. They commit to be net zero at a certain date. They commit to reduce by a certain percentage of their emissions. So because of those commitments, they have an obligation to basically measure and report on their emissions. So that's why after hearing this demand for a while, we were thinking, “Oh, maybe we could compute all of that out of the box for every Google Cloud customer.” And so that's what we've delivered with Google Cloud Carbon Footprint. Right inside the cloud console, you can click a new carbon footprints item, and you have access to a dashboard that tells you your yearly gross carbon emissions, that tells you you're broken down by month. So you can see how you're trending. It's broken down by region. Because as we discussed, the regions have large impact on your emissions. It's also broken down by product that you might be using so that you can see which Google Cloud product makes more sense to optimize. And lastly, it’s broken down by project. So a project is a construct where you group resources together under the same project. So that gives you a sense of maybe your project where you run your CI, your continuous deployment and integration system, is actually emitting a lot of emissions. And this one could be optimized, for example.

So to step back, we've heard demand from big, big customers to better understand and report on their gross carbon emissions. So we've provided this out of the box dashboard. But also, of course, the ability to export this data automatically into a table so that customers can then incorporate it into their own reporting and calculation over time. So all of this is automated, available to every Google Cloud user out of the box. And so we've seen great reception from early testers. So I mentioned, companies that we worked with, they have helped us test this product before we announced it at the last Cloud Next last year. So customers like Etsy, or Twitter, or L'Oreal, or even partners like SAP helped us validate the data before we opened the

product to public preview. And now we are working hard to bring that product to what we call general availability, which would be a stage where it would be covered by Google Cloud terms of services, and customers could rely on them to incorporate the data into their carbon reporting.

[00:26:40] KP: And what are the challenges in that last step of getting a GA?

[00:26:45] SG: That's a good question. And that's a lot of my current focus. As a product manager, my goal is to make sure the product meets the customer needs. So we knew it met customer needs before opening it up to the public. But since then, we've heard some feedback that we had heard a little bit, but not enough, to block the launch. But we want to add one or two small features to the scope in order to really make it useful to a larger set of users.

So for example, you might know that in cloud there is something called identity and access management. So this is like which permission do you need to access this carbon data? Today, we've taken the shortcut to say that if you have access to the billing data, then you have access to the carbon data. Well, before going to general availability, we want to introduce a dedicated IAM permission, a dedicated role, so that a company could grant this carbon viewer role to a certain person in their organization without granting them access to the billing data.

So for example, if you have a director of sustainability at your company, you could grant them this IAM role and so that they could access the report without accessing the billing data. So that's just one example among a few other features that we want to deliver. But the other big aspect is, of course, productionization, which we always want to add more monitoring before going GA. We always want to make sure that things are reliable and available. But also, in our case, there is finalizing the methodology and really polishing the methodology that we use to compute those per customer gross carbon emissions. And maybe I can talk a little bit about this methodology so that you understand where we are today. And what we want to tweak before the product is actually generally available.

[00:28:50] KP: Definitely, yeah. Let's get into the methodology.

[00:28:53] SG: Sounds good. So as a customer, when you open this dashboard, you see your gross carbon emissions broken down by region, month, and project, and product. But how do we compute that? Right? How do we go from measuring Google Cloud's data center energy consumption into providing our customers with a dedicated dashboard tailored to their use? This is quite interesting and quite unique. We believe we've worked on something quite unique here. Because basically, for a very long time, Google has been using some internal accounting system to understand which product are using how much resources.

And so basically, we are leveraging this system to understand, "Okay, which product is using which resource? Is it all those resources shared between products or dedicated to this product? Then Google Cloud servers are instrumented to measure the energy consumption, be it idle or dynamic. So idle is when the server is not used. And dynamic is when it's used by one or more service. So basically, we measure the energy consumed. We measure who is responsible for the usage of that resource. And then we cross that data with data coming from one of our partners, which is called Electricity Map. So Electricity Map is telling us, for every Google Cloud Data Center, the hourly carbon intensity of the electricity.

So basically, when we run one server, that for every kilowatt hour consumed, what is the equivalent gross carbon emission? So we use hourly data from a partner to achieve that. So that basically gives us the internal gross carbon emissions. But now, how do we translate that into a per customer carbon data where we basically need to allocate this data to Google Cloud customers? So this is where our methodology is basically computing for every Google Cloud SKU. So like, everything you can purchase on Google Cloud, we have internally associated it with emissions per unit of usage, right?

So to do that, we leverage our billing data at the global scale to a portion to allocate the gross carbon emissions on every Google Cloud SKU. And then we simply use the data of what our customers had been consuming. So when you look at your bill, your Google Cloud bill, you can see exactly what you've been consuming. So we basically take this data and use the carbon per unit of usage data to compute your carbon footprint reports that you can break down by region, product, and customer, obviously.

We do a monthly aggregation, because we don't want customers to rely on some daily noise that might come up from the methodology. So we really want to report to avoid people getting a sense of false precision. So that's why we do a monthly aggregation. And that is what you get on the user interface and the data exported.

And so, this is where we are today. And before we go GA, what we want to add is – So today, this is only reporting on the gross carbon emissions that come from electricity consumption. So maybe we should have talked about that. But basically, depending on how the electricity is produced, it emits more or less carbon, right? If it's coming from a coal or gas power plant, it will emit a lot of carbon. If it's coming from renewable energies, like wind, solar, or maybe even from nuclear power plants. And it will be what we consider carbon-free, or it has a very low carbon intensity.

So energy consumption has different carbon impact. But that's not – Or there are other activities that emit carbon. For example, when there is a power outage, then Google Cloud data centers can start up diesel generators. Those will be directly emitting carbon emissions. So this would be planned to the report. And also, there is some upstream and downstream activities. Like when Google purchase a server, or builds a data center, then the manufacturing of those things will have some carbon impact. So we want to incorporate that also into the methodology that we don't report yet. So today, it's only the emissions coming up from the electricity usage, be it electricity production, or indirect carbon emissions that are coming from the manufacturing of those power plants.

[00:33:49] KP: So I'm envisioning now the Director of Sustainability you'd mentioned. And let's think of a case where they're really empowered by the company to go and make positive change with carbon emission in mind. They use these tools to do some sort of optimization, move some things between regions, and the optimization is done and completed. Are they set for life? Or should we check up on this annually? What might be a healthy process for reviewing the way the infrastructure is set up?

[00:34:18] SG: So, usually what we see is that our customers want to report a yearly on their carbon emissions. And so, their emissions have indirect – Also coming from indirect emissions, which is their usage of Google Cloud. So what we see is that once you have started to measure

and report, you want to at least do that once a year. You want to, of course, confirm that you are trending down. But here, there are two actions. One of them is passive, which is just because of the fact that you're using Google Cloud, like Google by itself has made some commitment and is constantly making progress to having a higher carbon-free energy percentage.

So, I think, depending on your region, it's very likely that more and more Google regions will have higher carbon-free energy percentage. And so, just by doing nothing, you should see your gross carbon emissions going down, right. But of course, there is the active monitoring, which is maybe your Google cloud usage has increased. So if your usage increase, of course, your gross carbon emissions could increase, too, if you don't pay attention to also trying to optimize the emission per unit of usage, if you want.

[00:35:33] KP: Make sense. Well, as you think maybe a year or so out when the product is gone GA and people are engaged with it, are there things on the roadmap or areas you want to explore in the future?

[00:35:45] SG: Yeah. So we want to provide even more recommendations. I mentioned that we already provide one. But we would like to also help to give those advice like, “Hey, it looks like you're running a batch job in this location. Well, did you consider moving it to this region, which is very next to yours, but actually has a much lower carbon impact?” That's the kind of thing we want to get into at one point.

And then we also want to help customers estimate. So I told you how we are helping them select a region. We are helping them report on their usage. But what about trying to anticipate the impact of their choice? So because of the data that we have, we have the ability to predict, “Oh, if you were to pick that infrastructure in that region with that amount of usage, then we would predict you would emit that much gross carbon emissions.”

So basically, by doing some estimations, we could help, beforehand, customers select one architecture over another one. Of course, these would be estimates. But they would give a sense of the carbon impact of their choices the same way we give them a sense of the carbon impact of their region choice.

[00:37:03] KP: Interesting. Yeah, I like that idea. Well, you mentioned it's in preview now. Can average users get access to it? Or do people need to wait for it to go GA?

[00:37:14] SG: Yeah. No. It's open to everyone. In November last year, we launched it on stage at Cloud Next. And cloud.google.com/carbon-footprint is where you can get started. But as I said, if you're already a Google Cloud customer and have been using Google Cloud for more than a month and a half, then in the cloud console, you can look up on the left nav, on the left navigation, you can look up for carbon footprints. And if you click that, you should already see some data appearing, because it takes us a month and a half to produce after you first started using Google Cloud to produce some data. But once it's down, you should see it being updated every month for your projects.

[00:37:56] KP: So if I've been active on GCP for a couple months, I'll have my dashboard right away? All I've got to do is go to the console?

[00:38:03] SG: Exactly. No API to enable. Nothing to set up. Like the UI is already here in place for you. It’s just waiting for you to click on the icon of the left nav, which you can find under the hamburger menu, you scroll down, click carbon footprints. And yeah, if you've been using Google Cloud with a billing account, which most customers have, then you should see data.

[00:38:27] KP: Well, very cool. I hope we get everybody to go check that out and make some smart choices with these ideas in mind in their future work. Steren, thank you so much for coming on Software Engineering Daily.

[00:38:37] SG: Thanks for having me. Yeah, I'm looking forward for the audience to check out those products. As I said, my advice is you want to reduce your emissions? First step, as we said, move to the cloud. That's really proven. And second step, probably pick the region which has the lowest carbon impact. That's what I would recommend. And what I want the audience to remember is, as I described, Google Cloud gives you the tools to optimize your region selection and to report on your gross carbon emissions. Thanks to Google Cloud Carbon Footprint.

So it was a pleasure talking to you about this journey. It is a topic that I was very glad to be able to walk on at Google. And I'm always glad to join Software Engineering Daily. Thanks for having me.

[00:39:23] KP: Oh, the pleasure is all mine.

[END]