

Google Prodcast Season Four Episode 2 | The One With Data Centers and Peter Pellerzi

[JAVI BELTRAN, "TELEBOT"]

STEVE MCGHEE: Hi, everyone. Welcome to season four of The Prodcast, Google's podcast about site reliability engineering and production software. I'm your host, Steve McGhee. This season, our theme is 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 remember, hope is not a strategy.

—

STEVE MCGHEE: Welcome, everyone. We're coming back to season 4 of The Prodcast from Google. Can you believe season 4? It seems like a long time. I don't know, maybe it's just me.

So today, we have a great new guest, Pete Pellerzi, who works with actual physical things, like buildings and machines, and stuff, and not just like, bits. So, Pete, you deal with atoms, I guess. Is that true? I mean, that's wild. You might be the first one on our guest list to do that.

PETER PELLERZI: Well, I would say electrons.

STEVE MCGHEE: Electrons. OK, cool. And co-host today is Matt. Welcome back, Matt.

MATT SIEGLER: Thanks, wonderful to be here again.

STEVE MCGHEE: Awesome. So yeah, we did a little kind of a quick fake intro to Pete. But Pete, why don't you introduce yourself properly? Who are you exactly?

PETER PELLERZI: Sure. So by training, I'm an electrical engineer. Actually, I was a physicist before that. But I have an electrical engineering degree, and that's where I've spent most of my time doing physical things like large substations, transformers. We build those, install them, and we do the physical part of the Google world.

So at Google, I'm a distinguished engineer, which is a title that is relevant here at Google. And I

am assigned to the construction team. So we have different groups that make the physical infrastructure, tons of really talented people. And I work with the construction side, the physical side of it.

And our team is responsible to take the design that's done by our design groups and platforms and so on, make sure that what we do works for the correct generation of server technology. And we bring that into the physical world.

We're the ones who physically pour the concrete, even though that's a small part of our job. But we do the foundations, we do the building, we do the electrical, and the cooling infrastructure. We do all those big campuses that you see in the newsletters. So we turn it into the physical world.

MATT SIEGLER: Give us a sense of the scale of these operations, the people in--

STEVE MCGHEE: Yeah, they can't be that big, right? It's got to be--

PETER PELLERZI: Yeah, they can't be that big.

STEVE MCGHEE: Size of my house, maybe. No?

MATT SIEGLER: Maybe a gymnasium.

PETER PELLERZI: A typical data center.

STEVE MCGHEE: Small shed.

PETER PELLERZI: A small shed. Yeah, no, not quite. It's interesting because there's a lot of startup when you do the first one. So typically, what we'll do is we'll locate a community that meets our criteria. And we work very closely with the communities to land our not just data center, but our presence there for a long time.

So we're not just going to run in and run out. We work with the communities, we work with partnerships with the mayor and the environmental folks. We're very community-oriented. And we started building a campus. We'll buy 200, 300, up to 1,000 acres in a community, and we'll responsibly develop it.

So you'll see our entrance go in, where you'll pull in with your car and you badge in at the guard station. Then you see that the first building goes up. And we always plan for two or three buildings, because you want to make sure you don't do something weird on the first one, like stick it in the middle of the parcel and then you have no place to put the second building.

So we'll start planning the site and we'll build our first building. A typical data center is roughly two football fields on end. So if you took a football field and another football field and you put them together, that's the size of the actual server hall for one of our data centers.

Then you see all the stuff around them, the electrical infrastructure that supplies the power, then the cooling infrastructure that removes the heat out of the building, generator support when the power goes out once in a while, and then all the other structures-- places to eat, offices, entranceway, parking, and all that stuff.

STEVE MCGHEE: So with regards to places to eat, I think that means that you have employees there as well. Like, I'm guessing it's not like a fully autonomous system.

PETER PELLERZI: No.

STEVE MCGHEE: Once it's been built, I'm saying like, of course it's a zoo when you're building it out and it's never probably really done. But when it's operating, when you have services running, you've got like a slice of Gmail and a bit of AI or whatever is running in that data center, are there a ton of ants running around like pushing buttons on machines, or is it totally autonomous, or somewhere in the middle? Or like, what are we talking about here?

PETER PELLERZI: It's somewhere in the middle.

STEVE MCGHEE: OK.

PETER PELLERZI: Obviously, there's always a push to make things as automated as possible. You don't want to be going up and down rows of machines looking for the one that doesn't have the blinky light.

You want to know which one has a bad hard drive. You want to know which one has a bad memory. You want to know which one has a failed power supply, and so on. So those are

processes, break fixes are highly automated. Otherwise, you just can't keep ahead of them.

STEVE MCGHEE: That makes sense.

PETER PELLERZI: But we do have quite a contingent on site. Remember, we always focus on the machines, which is our primary focus. But there are a lot of other services, for example, 24-hour security, people, gardening, landscaping, food services, cleaning, repairs, general maintenance, and so on.

So there's quite a contingent of folks that maintain a so-called automated data center. They're getting there, but they're not science fiction. It's still humans that are fixing things.

MATT SIEGLER: Yeah, fixing things, that's got me thinking about things that go wrong, incidents, concerns, safety. Tell us a bit about how things get managed when they go wrong. What's an incident look like in a data center?

PETER PELLERZI: Well, having been in the data center industry for a long time, you don't give away your age anymore when you get beyond a certain age. So I'll say 25-plus years, that's sort of a safe way to say it. But in my 25-plus years, I've worked here, of course. And before here, I was at IBM, another large data center provider.

Every data center eventually will have some sort of failure. It's just a physical property of the world. Nothing lasts forever. Things break, things go wrong with no one doing anything wrong. It's just that things break. Circuit breaker, trips open because the trip unit failed, and so on. All kinds of things happen.

That's not really what differentiates success in a company, or differentiates a company as being successful. Things will break everywhere. It's how you deal with the break. That's the key. That's the differentiator, at least what I've seen here at Google.

For example, Saturday afternoon, 2 o'clock, perfectly fine weather, beautiful sunny day, something goes wrong. Utility fails. Who knows? A million things can go wrong.

STEVE MCGHEE: Can you think of an example of a particular failure that you can share with us? I know that you have exciting stories.

PETER PELLERZI: Oh, wait, well, I'm not sure about exciting. But, for example, a couple of weeks ago, well-documented, the entire country of Chile lost power.

STEVE MCGHEE: Oh, yeah.

PETER PELLERZI: Oh, yes.

STEVE MCGHEE: That was in the news, yeah.

PETER PELLERZI: Again, we have no control over that. Absolutely zero control. We have no way to do anything about that.

MATT SIEGLER: Who could see this coming?

PETER PELLERZI: It just, they failed. They had a whole bunch of things. You could read about it in the news, why they failed, and so on, but that doesn't matter. We did not experience an outage. So you say, well, why not? This is catastrophic. Well, OK. Hold on.

When something goes wrong on this proverbial Saturday at 2 o'clock or 1:00 AM in the morning, we have very specific protocols at all of our data centers. So, for example, we open up a video chat immediately, a common video chat between anybody who wants to get on. But usually, it's between facility managers, SRE, and so on. But what happens is folks automatically start to dial in. We get the emergency announcement on our email, we'll dial in, we'll help you. So we never strand someone at a data center and say, well, you're on your own, see what you can do. Not at all. And this is the differentiator, we behave like a community.

STEVE MCGHEE: Yeah, that's great.

PETER PELLERZI: And you never leave somebody just hanging. They know they have the support globally. You'll have SRE London, you'll have SRE-- everybody will chime in.

And respectfully, because you don't want a million people talking, but you listen. What can I do for these folks? Can I help? Have I had the same situation? What did I do? Are they forgetting something? This way you're not alone. And this is a super successful strategy, I mean, super successful.

So when we had this situation in Chile, we had multiple data centers down there, and they lost power through the whole country. So what's the strategy when you have a countrywide outage? You don't have that written down anywhere, but you have enough people with enough authority to say, yeah, why don't you get a hold of refueling? Let's start refueling right away, because we don't know how long this is going to last.

STEVE MCGHEE: They're able to adapt. It's adaptation--

PETER PELLERZI: We got on the fly. And you've got different people with different pieces, and say, no, you really should look at this and make sure you get that. And you really feel like they have your back.

STEVE MCGHEE: It's like cooperative adaptation. That's pretty cool.

PETER PELLERZI: Exactly. And this is the secret. This is not, let's make sure nothing ever breaks. That's not a strategy. Things break. But how do you deal with it?

MATT SIEGLER: I think some of our listeners would be really keen to hear how-- for their much smaller organization, perhaps with less of an existing roadmap for this distribution of skill, for these many layers of planning over a longer time-- how they could find a roadmap for themselves to build this kind of competency forward?

PETER PELLERZI: Sure.

MATT SIEGLER: What would you suggest to them, starting from where they are right now, how to develop their culture for this forward from small? And what kind of steps to take, like incremental approach to gain this kind of resilience?

PETER PELLERZI: It's interesting because I started at Google when we were very small. When I started, we had just a handful of data centers, and I actually was assigned to finish some of them as one of my first assignments here, because we really had just a few.

So how do you start from a small organization and scale it? Or maybe you're not going to scale it. Maybe this is all you need. You just need a few data centers and you're fine. You need to find trusted partners.

So for example, when we were smaller, we had very low staff. We depended on certain outside architects and engineering consultants, certain trusted vendors, electrical contractors and so on, that had specialty skills. And we used them, and we paid them. We didn't ask for charity. They were on retainer and we paid them for their time. But they had the expertise we didn't have.

So we said, listen, we're going to set up maybe an on-call agreement with you so that we don't have to write you a PO every time. But hey, when I need you on a Saturday, can I get a hold of Andy and Tom, and ask them to dial in or to go to the site and help the site recover? So we used outside vendors as we built our in-house expertise. You don't have to know everything. You don't have to be capable of everything, but you need a plan.

Just something simple like fueling, when you go on a generator, you need fuel. Well, you need to have the discussion with your local fuel vendors way ahead of time before you need them and say, what can you do? Well, I can get you there in four hours. OK, I got it. You sure? Yep, I'll be there in four hours when you need me. OK, now, that becomes part of your business continuity plan.

So start small, and start with the obvious things. Talk to your vendors, put down a few, even just a few words on a business continuity plan. What happens when power goes out? Just five bullets, and that really gets it going.

STEVE MCGHEE: With regards to the Chile incident, I wasn't involved in that at all because it was just recently. But I remember some other kind of power incident I heard about somewhere in the world, and they talked about getting refueling. But they had figured out how long it takes to get the fuel, but also how many trucks needed to be in flight at any given time to keep that up, to keep that--

PETER PELLERZI: That's right.

STEVE MCGHEE: They could keep it going basically, not forever, but for a long time. But they didn't really, they didn't know there was going to be a really big outage. But they thought like, well, we know how to get one truck of fuel, but let's just make a spreadsheet just in case in the future.

And they ended up using it. Like that spreadsheet totally saved the day because they knew to spin up the seven trucks, or whatever it was that can allow, they can get all the way to the fuel repo, and back again before the tank runs out, and blah, blah, blah.

So that kind of preparation I think is, I think that translates really well, not just to the data center world, but to any kind of dealing with response. You don't have to predict the exact outage, but just knowing your capabilities, and what you can do and like, well, this is what we have in our shed of possibilities. So at least you know what's in your shed, which is really important, or your tool belt, maybe.

MATT SIEGLER: Steve, this does sound a little like doing your postmortems after an incident, looking back on lessons learned. Looking at your analysis, how did it match up with what actually went on? And I expect you to do quite a bit of that after an incident like that, to tell us a little bit about that looking backward, after your forecasting, things you learned, things you got right and wrong. I suspect you do quite a bit of that.

PETER PELLERZI: Absolutely. And I'll give you a, let's stick on that fueling topic for a bit because it's something that you don't do very often. Let's be serious about this. The utility grids are actually fairly reliable. When you need it, it's maybe once in your career or twice, it's a very rare item that you need to refuel because of an outage.

A lot of times, you're refueling because you've used your diesel generators for maintenance purposes. And so then you kind of refuel at your leisure. I would challenge you to do as much real-world testing as possible. And here's something that we learned that is not obvious.

You say, well, how much fuel do I burn? Oh, we burn a truck full of fuel, for X number of hours. And that truck is 7,200 gallons, 7,250, I believe, full tanker truck. So you're like, OK, great, we need one of those every blah, blah, blah. And you do these quick calculations, and you think you're all dialed in, you're golden, but you haven't done it physically in the real world.

So we did. We had that opportunity. We had run the fuel down on one of the sites a little bit and we said, hey, let's do this as though it was a real emergency. So we started refueling. What did we learn?

We learned that the most you can get out of a truck is 300 gallons per minute. That's it. That's all they can pump. So now, we had not figured in the offloading time. And the offloading time turned out to be significant.

STEVE MCGHEE: It wasn't instantaneous?

PETER PELLERZI: It was not instantaneous. You have to drain 7,200 gallons out of that truck at 300 gallons per minute, max.

STEVE MCGHEE: That's a few minutes at least.

PETER PELLERZI: And that doesn't include setup time with the hoses, connection, disconnection, clean up, chocking the tires on the truck, getting the drive-- none of that was included, and we were very surprised at how long it took. And getting the truck in and out of the gate through security, we had not thought it through.

STEVE MCGHEE: That's a great example, because I think often the math, you can do enough math and just totally convince yourself, and then it hits reality, and you're like, oh, wait, there was a whole field of equations I forgot about or I didn't know about.

PETER PELLERZI: Do it. And that's a lot of the real-world testing that we do. As part of maybe not an emergency, but as part of a maintenance iteration, say, hey, what happens if that doesn't start? Or what happens if this doesn't run? Take advantage of every single one of those opportunities because you will learn something new every time.

STEVE MCGHEE: Let me ask you a question. Kind of shifting gears a little bit. When we talk about fuel, and infrastructure, and concrete, and things, I think of that as looking down, down the stack, right down towards the Earth and the fuel or whatever.

And then if you look up, I imagine sort of symbolically, like, after a certain point, you get out of the data center and you get into the distributed system that's built on top of it. And so this is where these big systems of software are running across many data centers.

And then, now, you're interacting with other teams. You're not really dealing with vendors, dealing with trucks. You're dealing with SREs in London, you mentioned, or software developers

in Tokyo, and things like that.

So what are the interactions between your staff, the people that you work with on site, and those folks? Like, is it just, is there a brick wall and you never talk to each other? Are there tickets like, you mentioned in this video conference, the panic room, which is tongue in cheek. Hopefully, it's not actual panic. But are there other methods that you guys-- like, is everything figured out or is there also adaptation going on in that direction as well?

PETER PELLERZI: Absolutely, and we talk all the time.

STEVE MCGHEE: That's awesome.

PETER PELLERZI: Before this call, I was on another call with, probably, 20 SRE folks, talking about the next round of DiRT testing, ISON testing, and so on, that we will be doing this year. It is a constant communication. And the best communication we can have is that our work is totally transparent to the rest of the fleet. So like the Chile situation, the whole country had no power, but we were not impacted.

STEVE MCGHEE: Can you tell us what DiRT and ISON testing are?

PETER PELLERZI: Oh, of course. DiRT is simulated failures. And I forgot what the acronym is. SRE, you all should know better than I do.

STEVE MCGHEE: Yeah, Disaster Recovery Test or Resilience test.

MATT SIEGLER: Resilience.

PETER PELLERZI: There you go. It's simulated. What happens if this goes really bad? How do we recover the customers, the data, and so on? So we coordinate with that all the time. And then ISON testing is where we actually fail things off. We actually shut things off and say, what happens if this shuts off? But basically, it's a *real* shut off. And this is where we walk the talk. We say, well, you should do this, and you should do that. And it's very easy to say that, but we actually do shut things off. And we say, well, what if the utility failed? Well, OK, we could do a simulation, or we could turn it off.

STEVE MCGHEE: And then how does this affect the upstream systems? Like, are they actually capable of surviving such a thing?

PETER PELLERZI: Well, it shouldn't. It shouldn't. But again, like offloading the fuel truck, it's something we don't do very often because we never want to run on diesel. For all kinds of reasons — environmental, and so on. So we try never to run on diesel generators, but so even more important to actually test it because it's something you don't do very often, which means, oh, it'll be fine. Really, it won't be fine.

STEVE MCGHEE: Yeah, out in the world they call, sometimes, they refer to this as chaos testing. That's mostly in the software failure modes, but you can imagine a similar kind of thing. It is like, induce a little bit of planned chaos into the system and see what happens. Like, throw the monkeys in the room and see what they do, kind of thing.

PETER PELLERZI: Like, we tried to refuel because we had done some maintenance, and we said, well, let's actually try it under emergency conditions. And boy, did we learn a lot. So that's part of this-- again, it's, if your strategy is to hope that nothing ever fails, that you're going to be very disappointed.

STEVE MCGHEE: Indeed.

MATT SIEGLER: Pete, tell us a little bit where we're going next generation tech, either through power technologies or something new going on in data centers that you can actually share with us or people listening that we wouldn't expect. What's coming down the road for us?

PETER PELLERZI: Density.

MATT SIEGLER: Say more.

PETER PELLERZI: Yeah, so let me take you a little bit through that. 10, 15 years ago, most of the data centers, well, all of the data centers were some Intel AMD chip, and it was very similar to what's in your PC at home. Of course, I'm being very simplistic here, But they were air-cooled chips, meaning they had a heat sink stuck on them, and a little muffin fan, and it blew the air across the heat sink and took the heat off the chip, and all was good. And they were fairly low power consumption, 100 watts, something like that, maybe 150 watts.

Then the industry, of course, was hungry for more capacity. So you stick more things on the die, you make that chip a little bigger, you stick more chips in there. You build a bunch of chips into one chip, Application Specific Integrated Circuit, an ASIC, which is a conglomeration of things that work together on one sort of package. And that march has been going on forever.

STEVE MCGHEE: Yeah, Moore's law is part of this. Like, it's just getting bigger and bigger, and tighter and tighter, and yeah.

PETER PELLERZI: More, and more, and more. I want to increase the clock speed. I want the thing to go faster. Because they're expensive, I want a higher return on how much processing power it can do for the thousands of dollars that I've invested. And you've seen the Intels of this world and now NVIDIA. I can do better. It's four times faster, it's 10 times, and so on.

So all that is wonderful. And it's not my area of expertise. It's people who have all kinds of PhDs in that, that can talk much better. My world is power and cooling. So we observe these chips, and they require more power every year, more power, more density.

But they give a lot of performance. So the business gets a lot of return on their investment. Power turns into heat, simple as that. I supply more power into the chip, or the data center, or whatever, it does work. That work turns into heat.

STEVE MCGHEE: Yeah, at this company, we obey the laws of thermodynamics. Is that right?

PETER PELLERZI: Power goes in, does work, work turns into heat. That's the way it goes.

So the power that we're providing now is probably several X more than, and I don't remember the numbers, but it's several X more than it was 10 years ago. And it pushes you to certain conclusions. And the whole industry is coming to the same conclusion that we did six, seven, eight years ago. So we were ahead of the curve.

We said, look, if we extrapolate this up, there will come a point, where you cannot put a heat sink strapped to a chip and blow air across it. You're going to reach a thermal limit. You simply cannot get the heat out of that integrated circuit using the traditional aluminum heat sink and a fan. You're going to have to use a higher heat capacity media.

STEVE MCGHEE: I think I know where we're going, yeah.

PETER PELLERZI: Liquid. Now, we say water cooled, but it's not water. They're different fluids, but basically it's water with anti-corrosive inhibitors and things like that. But essentially, the manufacturers, in particular NVIDIA, said, we can't get there without a cold plate strapped right to our chip.

So we're going to take copper or some other proprietary material. We're going to stick it on the chip, and you're going to pump cold water in one side and out the other, and remove that heat using water, which has 3,000 times the heat capacity as air does per volume.

STEVE MCGHEE: Yeah, computers notoriously love interacting with water, I think.

PETER PELLERZI: Well, yes. So we feel, or at least I feel, that Google was extremely forward-looking six, seven, eight years ago, where Google said, look, we need to learn how to do water-cooled processors. It was not simple. It was not trivial. A lot of folks are trying to do that now, at the 11th hour. And yes, you bring water directly to the chip. That is a leap of faith.

STEVE MCGHEE: Yeah, I remember hearing about this a long time ago. And they were like, the first plan was we bring these pipes through the ceiling and bring water down. And then we realized, that's a bad plan because drips happen, little things like that. Man, that's tough.

PETER PELLERZI: Well, but we've spent the last, and I can't remember if it's seven or eight years, let's call it seven. We've spent the last seven years, a lot of really good, sharp people spending a lot of sweat perfecting this. You just don't run a couple of pipes and yeah, here I've got liquid cooled chips.

There's a lot to it, monitoring, flows, temperatures, leaks. What do you do if there's a leak? How do you not have this? So we've learned a lot. And that is our secret sauce now.

And I love when people say, well, what's the secret sauce? Well, let me think. It's seven years of really hard work by a lot of really dedicated people. And then, magically, a secret sauce appears.

STEVE MCGHEE: That's right.

PETER PELLERZI: It's hard work, period.

MATT SIEGLER: Speaking of expertise and new tech, somebody told me that artificial intelligence is optimizing our data centers for power. Is that a thing?

PETER PELLERZI: Yeah. Well, again, it's funny. I guess it used to be called machine learning. And again, I'm not the subject matter expert on this, but it is very useful to us. And we've been using it for several years. I think there was a public announcement on our use of--

MATT SIEGLER: Yeah, a while back.

PETER PELLERZI: I think it was called machine learning, ML, but now, it's AI, on how to run our cooling plants, and we got some very large savings. Again, I hate to say numbers because then someone always says, oh, no, but I think it was between 15% and 40%, if I recall.

STEVE MCGHEE: That's awesome. This is a PUE rating. I think that dropped overall.

PETER PELLERZI: Yeah, exactly. And here's the secret. As human beings, if you sit in front of a control panel and have to adjust the knobs every two minutes, you will lose your mind. I mean, you can't do it. Also, you don't have predictive capabilities.

So we had, again, some super sharp folks here. They worked with the London folks, DeepMind. And we came up with this little-- little-- it's always little when you don't have to do the work. It was actually a lot of work.

STEVE MCGHEE: From a distance, from London to here, yeah.

PETER PELLERZI: Yeah, they came up with this ML approach to say, look, we can not only take all the inputs from all these devices and find an optimal running point for each one, but we can also take predictive weather data.

STEVE MCGHEE: Oh, cool.

PETER PELLERZI: So I can make adjustments now based on what the weather will be two hours from now.

STEVE MCGHEE: That's awesome.

MATT SIEGLER: That's amazing.

PETER PELLERZI: So it's real time weather data with real time operating data. And you watch it. It's fascinating to watch because it will start to turn down the mechanical cooling, a chiller, which is our refrigerant-based air conditioning. It will start to reduce that. And you're like, what is it doing?

And then you watch it for a little bit longer. And it makes a decision that if I increase the fan speed and the pump speed a little bit, I can get the same cooling for less power. So it starts to turn down the refrigeration because it anticipates the weather two hours from now. Fascinating to watch.

STEVE MCGHEE: So at Google, we're happy about dealing with large numbers, like we like it. And I like to always say that a small change in a large number is still a large number. So this happens a lot inside of Google whether. It's a good change or a bad change.

So in your world, what is the number that you used to tell people about your job and how it's interesting? And what is something that you can share about the scale of your world? Is it the square footage of cement, or is it the density of something? Like, what's your favorite number?

MATT SIEGLER: In joules or watts?

PETER PELLERZI: Yes.

STEVE MCGHEE: Yeah, choose your units wisely.

PETER PELLERZI: I choose my units wisely. There's all kinds of capacities and whatever, but this is what impacts me. I'm an electrical engineer by trade, so I'm an electron guy.

So I'm always fascinated by the other side of the house, which is the cooling side. Because, to me, it's a little mystifying. You make cold water, it goes in, it gets hot, it comes back out, you make it cold again, fascinating. Because you're always fascinated by what you don't know that well.

So I went out to one of our sites and I looked at the pipe, and I said, what's that? And they said, oh, that's the chilled water supply to the data center. I said, oh, so what's, OK. And what's that? Well, that's the return. The hot water comes back. OK. So how big is that? It's 42 inches in diameter.

STEVE MCGHEE: That's pretty big.

PETER PELLERZI: So that's the scale. And essentially, a 42 inch diameter pipe, if you creep down a little bit, you can walk through it. And I was like, that's what it takes to cool one of these things? And they were like, well, yeah. They're very nonchalant, like, oh, yeah, oh, yeah.

STEVE MCGHEE: I'm pretty sure there was a scene about this in one of the Star Trek movies, where someone gets transported into one of these pipes and they flow through it.

PETER PELLERZI: Exactly. This particular site was one of our larger ones, has a 42 inch chilled water supply and return.

MATT SIEGLER: That's a big number.

PETER PELLERZI: And to me, that just-- it's like, really?

STEVE MCGHEE: And I bet it's quickly flowing through that diameter as well. It's not just--

PETER PELLERZI: Yeah, exactly.

STEVE MCGHEE: --easily going, yeah.

PETER PELLERZI: Several feet per minute.

STEVE MCGHEE: That's awesome.

PETER PELLERZI: And that just sort of for me, put it in scale, like, well, who else has this? And that's just the hyperscalers.

STEVE MCGHEE: Yeah, would you say in your opinion-- I don't think you have a full understanding of everything on Earth-- but would you say that the Google data centers are different from other data centers at just a--

PETER PELLERZI: I would because they have different purposes, and it's not good or bad. Data centers are kind of a personal thing to companies. If you're this kind of company, you optimize for an enterprise type data center. If you're this kind of company, you-- so Google is personalized to what we do. They look the same. All data centers have the same purpose.

STEVE MCGHEE: Big box, yeah.

PETER PELLERZI: Big box, you put electrons in one side, it gets hot. You take the heat out the other side, and you get rid of the heat. But it's all the stuff inside. How do you talk to the servers? How do you interface with the networking?

All that is really very unique to whatever company you're in. So it's not a good or a bad, but I think we have a really good mix, that's good across enterprise and hyperscale type ML. So we have a really good mix.

And it's not by accident that we've had really good forward-looking leadership, like seven years ago, folks in platforms and delivery embraced the water cooling. And that was a really good idea.

STEVE MCGHEE: Yeah, that's great. Yeah, it's really easy to say no to stuff like that. But the fact that they said yes is great.

PETER PELLERZI: It is, because you're like, why do we need this? But you need to look four or five years ahead of time, if you can, and say, yeah, we should learn how to do that.

STEVE MCGHEE: When I talk to companies-- so part of my job is talking to enterprises about how they use cloud and data centers. And they're coming from their data centers or some other data center or whatever, and they're moving on to cloud, and blah, blah, blah.

PETER PELLERZI: Sure.

STEVE MCGHEE: So one thing that comes up a lot is they like to cite MTTR and MTBF. And I'm curious how you feel about these statistics in your world. Because I know that in my world, in the distributed systems game, we like to say that those numbers actually aren't meaningful to us because it's not a large, homogeneous set of components that all have similar failure

domains or failure modes that we actually have novel failures. So like treating this as it's not a normal distribution. So why are you taking the mean kind of thing?

But I'm curious if you have a data center full of eleventy million similar parts, maybe those numbers do make sense to you at a certain layer, but not at another layer. Is there a point, where the normal distribution disappears in your world, and do you take advantage of that when you're trying to measure your success? Is that a thing?

PETER PELLERZI: It is. So what you're describing falls under our operations team, to monitor and to measure. Certainly, again, I have to emphasize no one is left alone. We do not throw things over the transom and say, well, good luck. No.

STEVE MCGHEE: Your problem now.

PETER PELLERZI: That is not the community we have here.

STEVE MCGHEE: Great.

PETER PELLERZI: Just we don't have that. So what's of concern, the metric that we look at is availability, 99.999, whatever, five nines of availability. That's our target.

And the way you get to that target is, minimize failures. So you don't have to fix things. And then fix things as fast as you can, the MTTR, so your downtime is minimized. And that's how you get to that overall availability.

How do we get there? We have some unique features or unique opportunities. The manufacturers don't have many customers that have a bigger installed base of their equipment than us, so we actually have better data than the manufacturers of the equipment themselves because they don't have a concentrated pool of equipment that we have. So we see all their latent defects. We see all their strange applications and so on. So we work very closely with the manufacturers.

The other thing we do is we build fault-tolerant designs. Again, if you're going to expect something never to fail, that's not so good. We have a very, very robust spare parts program so that if something fails, A, we can recover very quickly using a fault tolerant design, and then B,

we have the spare part in hand.

So we can see what fails across a large fleet, even though the manufacturer may tell you we've never had a failure on that. And they're not lying. They've never seen it because they sell two to this person, five to that person. But we buy 600. And we say, well, if you have 600 of anything, and you have a 0.001 failure rate, we're going to lose one every two years. So--

Spare parts strategy, repair strategy, fault tolerant design, and working closely with your vendors. And we do really, really well with availability. I would say we're probably the best in the industry.

STEVE MCGHEE: Yeah, we like to joke that at Google scale, a million to one odds happens all the time.

PETER PELLERZI: Exactly. And it's important not to assume ill intent in your vendors. Because they, you say, well, it just broke. You should have told me. They don't know because they don't have a large enough statistical base. So work with your vendors. Again, this is a community sort of argument. Work as a community, not as a bunch of individuals.

STEVE MCGHEE: Totally. Thanks, Pete. This has been great. It's been cool to touch grass with the actual data center. Touch cement, maybe, I don't know, touch steel.

MATT SIEGLER: Get our hands, hold a shovel.`center provider.`

PETER PELLERZI: No, no, cement is what you mix with the gravel and water--

STEVE MCGHEE: Oh, sorry.

PETER PELLERZI: --to make concrete.

STEVE MCGHEE: Concrete. I always get that wrong. Well, I appreciate it. We learned a lot, especially that last bit, just there. Any final words for our friends about how people can keep up with this kind of stuff? I know we have a podcast that Stephanie Wong did a few years ago about how the data centers work in communities like you were alluding to. Any other resources

you want to point people at who are listening to the show?

PETER PELLERZI: There's a lot of good content out there. It's a fascinating industry right now. It's really, really exciting. The scale is really exciting. There's a lot here.

STEVE MCGHEE: Still growing man. It's crazy.

PETER PELLERZI: There's a lot here.

STEVE MCGHEE: All right, thanks, Pete. Have a great day, everyone.

MATT SIEGLER: Yes, thank you so much.

PETER PELLERZI: Take care.

—

[JAVI BELTRAN, "TELEBOT"]

JORDAN GREENBERG: You've been listening to Prodcast, Google's podcast on site reliability engineering. Visit us on the web at SRE dot Google, where you can find papers, workshops, videos, and more about SRE.

This season's host is Steve McGhee, with contributions from Jordan Greenberg and Florian Rathgeber. The podcast is produced by Paul Guglielmino, Sunny Hsiao, and Salim Virji. The Prodcast theme is Telebot, by Javi Beltran. Special thanks to MP English and Jenn Petoff.