Brandon: [00:00] Large enterprises use EDP as a first step in their discount strategy, and EDP should really be the very last step of that strategy. So you should be using all of your other optimization options first, compute savings plans, CRIs, credits, usage-based discounts, and once you've done everything that you possibly can, only then should you consider EDP.

Announcer: [00:30] This is Telco in 20, a podcast that helps telco execs achieve a competitive advantage with AI in the public cloud. It is hosted by Danielle Rios, also known as DR. Today, we're talking to Brandon Pizzacalla, the CEO of CloudFix.

DR: [00:49] Hey guys, I'm DR. Every telco out there wants to control cloud costs. Those that spend big with AWS can get help through the hyperscalers' Enterprise Discount Program or EDP. The way it works is customers agree to spend a certain amount on AWS services over a set period of time and in return, they get a discount on their overall bill. But here's a fun fact about EDP that a lot of telcos don't know. To calculate your discount, AWS snapshots your usage months before you renew. That means the key to getting the best deal is to reduce your cloud spend way in advance and optimize it continuously. But let's face it, most DevOps teams are overloaded with their day-to-day work. So cost optimization usually falls to the bottom of their list. So how do you stay on top of it? You get some help. I'm not talking about consultants, although they can be great.

[01:42] I'm talking about software that can do the heavy lifting for you like CloudFix. CloudFix is a super cool tool that can help you optimize costs automatically. It identifies easy-to-make changes that add up to big savings without taking your team away from their top priorities. It couldn't be easier to use. So today we're going to talk to CloudFix about the three common mistakes that enterprises make in cloud cost optimization, how you can use CloudFix to save millions of dollars, and how to put yourself in the best negotiating position for your AWS EDP. So let's take 20. Brandon Pizzacalla is the CEO at CloudFix. Hi Brandon. Welcome to Telco in 20.

Brandon: [02:25] Hey DR. Really glad to be here.

DR: [02:27] I'm so psyched to talk to you about cost optimization. We're going to talk about how your company, CloudFix, helps organizations get the most out of their AWS spend by

automating cost savings and finding efficiency improvements. And so I know there are a lot of telcos out there that have signed multi-year big money EDP agreements with AWS in exchange for a discount on the services they use. As the cost optimization expert, what are some of the common mistakes large enterprises make when negotiating their EDP with AWS?

Brandon: [03:00] Yeah, this is a great question and there are a ton of mistakes that large enterprises make. So number one is large enterprises use EDP as a first step in their discount strategy, and EDP should really be the very last step of that strategy. So you should be using all of your other optimization options first, compute savings plans, CRIs, credits, usage-based discounts, and once you've done everything that you possibly can, only then should you consider EDP. Yeah, as an example, I mean, we're working with a customer right now. They spend $250 million a year on AWS, so they're huge. They jumped on the EDP train before optimizing everything else, and now we are learning they're wasting about $8 million a year compared to if they optimize first. So just imagine what you could do with an extra $8 million, whether it is hiring more people or just putting it to the bottom line.

DR: [03:58] Spending on GenAI or-

Brandon: [03:59] Exactly. The second biggest thing I see is large enterprises overcommit to their EDP. The AWS sales team will come in hot. They will try to get you to sign up to this increasing commit every year. They'll make you feel good about how you're growing and your AWS spend should be growing, but that's the opposite of what you want. You will sign a deal that ultimately ends up costing you too much. And we're working with a prospect right now. They do about $60 million a year in AWS spend. We did the math. If they had never negotiated an EDP at all, they would actually be spending less money. So that's the kind of situation you get stuck in.

[04:39] Number three would be not using a full team to do the negotiations. So typically we see a FinOps person managing the negotiations. They know all AWS numbers and the spend and the different instances that are being used. And what ends up happening is they miss these key strategic negotiating points that would be solved if you brought in a team across legal, engineering, finance, and business. And you have to think about

this as you're negotiating a multi-million dollar deal, should you have one person doing it or should you get a team together to do it?

DR: [05:17] Yeah, one thing that was sort of surprising for me as we've been negotiating EDPs for our own usage is they start looking three months, six months in advance at your spend, and obviously the AWS rep is trying to push your numbers up and so your first tip of optimize first and then do your EDP last, we do do that, but you really can't optimize it with 90 days to go before the renewal is due. You really need to be continuously optimized, always have your level of cloud spend at the lowest level. You can't sit there and tell AWS, "Well, we're going to optimize this, and $8 million is going to drop out." They're like, "Nope, your levels are here and your EDP is this number or higher."

Brandon: [06:00] Yeah, that really catches people off guard too. So either they never optimize, they're always running out of time or they start too late. So agreed. Should be always optimizing.

DR: [06:10] Yeah, so in that optimization effort, we see a lot of companies focusing only on the big-ticket items. This happens in budgeting too, just in general. People focus on the big items and try to bring them down and they don't really focus on that long tail. And so why is it important to also look at that long tail of smaller, maybe seemingly insignificant spend on the other AWS products and services that companies use?

Brandon: [06:35] Yeah, absolutely. And I think we understand why people spend human time on big-ticket items. They're not cheap resources. So when you're looking at what you might be able to optimize, of course, these people go after the $100,000 a year, $200,000 a year savings. That totally makes sense. But the long tail is where millions of dollars, literally millions of dollars a year sit and no one spends the time doing it because they all look so small. So when I say long tail, it's these $100 a year, potentially $1,000 a year services that are just sitting there and no one ends up going to fix them.

[07:12] That is literally why CloudFix was built to automate all the finding and fixing of these long tail optimizations. So, I think the biggest long tail optimization that I've seen is something we've done internally. We've also helped our customers with it,

is optimizing Amazon QuickSight dashboards. It sounds super unsexy. No one wants to do this, but this is a $24-a-month dashboard user cost, and every time you share a dashboard, then a new user gets added forever. And so long story short, we automate this optimization and we've saved 250k just internally in the last few years by cutting those $24-a-month users.

DR: [07:55] That makes me feel bad because I'm added to a bunch of QuickSight dashboards and I'm never logging in, and so that's expensive. I didn't know that. Why do you think this happens? If there's so much money on the table, why aren't companies figuring out how to optimize the long tail themselves?

Brandon: [08:10] To do this you either need to pay an expensive team of experts just too much money so that ROI doesn't make sense or build this automation in-house on your own. And generally what we've seen is no one has the time or no one wants to put the investment in to spend a year or so building this automation and then maintaining it and improving it.

DR: [08:31] Yeah, super unsexy work, going to all the couches in the world and getting all the pennies and pulling it together. So we're in this age of AI, how is CloudFix using AI to optimize cloud spend and what improvements have you seen compared to the methods that you used before AI came out?

Brandon: [08:47] AI has been extremely valuable. So right now we use it in two ways. One is what we would call it, the UI Copilot. It actually uses AI to understand what's going on, on your screen and walk you through the steps you need to take in your AWS dashboard to implement optimizations and savings. The other way that we're using AI is to provide extremely deep insights and recommendations. And this is something that there wasn't value to do it if you're putting a human on this task, but with AI, we can ingest thousands of lines of code from Lambda functions you're running, from S3 buckets, from your ticketing software, from usage statistics, and all of that can be analyzed and combined to give the GenAI tool an understanding of why you are using an instance for and what you are doing with it, and then provide really insightful recommendations on how you should adjust your parameters in order to save money while still being able to accomplish what you're trying to do with all of that code.

DR: [09:51] And I think that's what really wigs people out, kind of scares them off is, yes, you've identified an optimization, something for me to reduce or eliminate, but I'm scared of the impact. And I think that's what really causes people to pause on optimization strategies because they're like, "Well, it's running and if I delete this thing or stop using it, what are the downstream effects? And I don't really want to mess up our operations," and so they don't change them.

Brandon: [10:17] Absolutely. You nailed it. And part of what the CloudFix AI tool does is provide FinOps and DevOps teams with all of that information so they can understand exactly what's going on and make the decision if they want to make certain changes to optimization parameters. The most recent example is we had this Lambda function that was running and it had been running for years, and what was hard to tell is it looked like it was doing something and no one really spent the time, which maybe be a day or two of work to dive into what was going on, but as soon as we turned on the AI automation, it immediately identified that this is essentially just looping nonstop doing nothing. And that wouldn't have been obvious without AI being able to analyze everything.

DR: [10:59] That's so awesome. And so I know you have a bunch of different innovative strategies that you guys are using at CloudFix besides obviously applying AI. And so what are some of the other innovative strategies CloudFix uses to maximize AWS savings?

Brandon: [11:12] This is a great question because some of the things you can do to optimize AWS savings actually have nothing to do with the technical side. So what I really like to chat about with our customers is a few key points. Number one being establishing unit economics. So this sounds kind of boring, but I think it's interesting. You'd go and identify your key cloud cost drivers, so on a per-user or per-transaction basis, and then align these metrics with the broader business goals that you have and track them and you incentivize people to head towards your unit economic targets. You can implement that, by the way, as a finance leader. You don't need to be technical.

DR: [11:51] Correct.

Brandon: [11:51] Number two, we call this implementing chargebacks. And simply a chargeback is just a financial penalty that you apply to teams or BUs for inefficient cloud usage. So you can say it's best practice to use a certain instance for this type of work, and if a BU decides to use a different instance, they are actually hit with a penalty. So you align people's incentives financially with what you want their behavior to be on the cloud side.

DR: [12:19] Yeah. Like an internal tax.

Brandon: [12:21] Yes.

DR: [12:21] You guys are lazy and not optimizing, and so we're going to charge you more. And hopefully, that aligns the behavior to be like, "Okay, well we want to spend our budget on other things. We have this dumb tax, let's try to figure out how to eliminate the tax."

Brandon: [12:34] And I'm sure almost everyone's experienced this before, the engineers want to be driving a Ferrari all the time, but you don't need a Ferrari to take a quick trip to the corner store. So that type of optimization can be done.

DR: [12:46] Exactly.

Brandon: [12:47] Number three, which I think is very important, a lot of people are missing actually, is optimizing your discounts. So I think I could do an entire podcast on discount optimization.

DR: [12:57] Maybe we should.

Brandon: [12:58] Maybe we should. Yeah. The three main points I would make to keep it high level and something that everyone could take away from this conversation is number one, stop using the super inflexible reserved instances. Switch to using compute savings plans to cover about 30% of your consistent compute workload, and then for the remaining compute workload use convertible reserved instances and those can be switched to different regions and different instance types. So that's what the convertible aspect means, and we saw so much value in optimizing discounts and how painful it was to actually do it.

[13:38] That's why RightSpend was built. CloudFix RightSpend is a product we've had around for about a year and a half now,

and it literally goes and automatically optimizes your entire discount portfolio to save you a ton of money. And the latest customer we have on RightSpend had done all their optimization previously. They negotiated an EDP really well, so things were looking great for them, and they actually told us, "I doubt that you could help us." We went through our RightSpend analysis and now they're a customer of ours and we're saving them an extra $5 million a year on top of what they thought was a fully optimized account.

DR: [14:16] No, that's awesome. I love RightSpend. It's such a great idea and it's kind of interesting because CloudFix, it's a startup like Totogi. You're constantly trying to figure out how to get that brand awareness and market reach. And so CloudFix originally was a SaaS product software that assumed customers would just connect their AWS account and just let it go work on its own. But I think Brandon, and you've learned that customers haven't been exactly comfortable working this way, and so you're changing your approach. How have you pivoted CloudFix and how are you onboarding customers now so that they're successful with your product?

Brandon: [14:53] Yeah, it's a great question and I think if you step back, it seems very obvious that people won't just let certain things automatically run in their AWS prod instance, but I kind of had to learn that the hard way. So what we've been doing is just much more human interaction and human touch offering professional services and consulting suite where we are able to work directly with senior executives or with technical teams or with FinOps and DevOps teams in order to help them understand which optimizations they should be focused on, why they should be implementing them, which our AI tool helps with as well. And then actually helping them implement them through the change process that they might have internally. So solving all of these other points of friction that DevOps and FinOps teams face internally, we're helping do that for them so they can actually realize these savings and transfer that to the bottom line.

DR: [15:50] Yeah, that's awesome. I think even though we're in this age of AI and everything is being autonomously done, people still need help. And again, I think what you guys are doing really well is starting to get people on that flywheel of optimization, just getting them onto the on-ramp and on the freeway, and then you have a wider touch later when they start to realize this

	is safe and it does work, and we're saving tons of money and it's great business value. So that's awesome. So I always like to end my podcast with something fun. And when I met you, your last name is super interesting, and so I call you Brandon Pizza, not Pizzacalla. Who doesn't love pizza? I mean, pizza is like the best. So I love pizza.
Brandon:	[16:34] I love pizza.
DR:	[16:36] I think a lot of people like deep dish or thick crust. Don't hate me. I'm a thin crust girl. And so my question to Brandon Pizzacalla, what's your favorite pizza crust?
Brandon:	[16:47] I think I could probably only get this in parts of the US so maybe I do this once a decade, but I'd go for the cheese stuffed.
DR:	[16:53] That's like a heart attack.
Brandon:	16:54] Yeah, I would feel terrible after. But my go-to though is also thin crust, to be honest with you. And the biggest tip I can give to anyone that likes pizza is just ask for it to be well done when you're ordering it. And I promise you it's going to be a 10X better pizza.
DR:	[17:09] That's awesome. Well, I will do that because I really do love thin crust and I like it when it's crispy. And so Brandon, I think our listeners have learned so much about your tips on cloud optimization. I think this is totally actionable, not just to use CloudFix, but just in general negotiating that EDP and really getting your teams aligned, saving real money. And so Brandon, thank you so much for coming onto the podcast.
Brandon:	[17:32] Thank you so much for having me, DR.
DR:	[17:37] Stick around. We end each podcast with a Telco in 20 takeaway. I've got two minutes to tell you something you need to know. Brandon and I just talked about why telcos aren't optimizing the long tail of AWS spend for easy savings. My opinion, DevOps teams have cost optimization on their list, but they blow it off for higher priority, more urgent tasks, and as a result, it never gets done. If I was the CFO of a telco, I'd be pissed my team wasn't doing everything they could to optimize the crap out of our AWS spend. I'd want to be sure we were minimizing any wasteful spending on AWS and perhaps

introduce a check and balance on the process to make sure the work got done. That's where CloudFix comes in. It can help you save up to 20% on your total AWS bill and as much as 50% on your EC2 spend.

[18:30] That's usually the biggest line item on an AWS bill. Just imagine what you could do with all that extra money. And guess what? I can help you save even more. I negotiated a special CloudFix discount exclusively for my buddies in telco because it's something every operator using AWS needs. If you're spending more than $20 million a year with AWS, give me a shout and we'll get you ready for your next EDP negotiation. And after you save beaucoup bucks with CloudFix, let's connect on LinkedIn or X at TelcoDR. Tune into more amazing Telco in 20 episodes, like and follow, and leave us a five-star review. And don't forget to swing by our killer YouTube channel for more telco insights. And sign up for my super awesome email newsletter on telcodr.com. Later nerds.