



CLOUD REALITIES

CR054

FinOps unleashed! with
Kunal Agarwal, Unravel Data



CLOUD REALITIES



LISTEN NOW

Capgemini's Cloud Realities podcast explores the exciting realities of today and tomorrow that can be unleashed by cloud.

CR054

FinOps unleashed! with Kunal Agarwal, Unravel Data

Disclaimer: Please be aware that this transcript from the Cloud Realities podcast has been automatically generated, so errors may occur.



[00:00:00] Passwords right now, just changing all my passwords for my birthdays and my wife's birthday and the kids birthdays to Rob Strong password 123. I think we should all share that one.

Welcome to Cloud Realities, a conversation show exploring the practical and exciting alternate realities that can be unleashed through cloud driven transformation. I'm David Chapman. I'm Sjoukje Zaal, and I'm Rob Kernahan.

And this week we're going to be wrestling with the important subject of FinOps. We're going to explore what FinOps actually is, some of the problem scenarios that you get, and actually some of the best practices that you can apply to your cloud environments.

To drive not only more productivity out of those environments, but do that in a way that the cost is very visible and [00:01:00] controlled. But before we get to that, bit of an interesting one this week, I was walking along and this is about me being confused, Rob, more so than you, but I wonder whether you've got a few on it.

So I was walking along having my morning walk and I'm lucky enough to live in a little village in the, in the UK. And there's quite a lot of trees around. I was walking along listening to a podcast, not, not our one, a different, a different one in my headphones. And my headphones are ones that are like over here and they've got a band across the top of my head.

And I was walking along and a tree swooped in and pulled my headphones off and threw them on the ground. And I was wondering, is this like an ents having a joke? Or is this like, what was going on with that? And anyway, I was confused. What have you been confused about this week, Rob? I just imagine you walking along, bumbling away, just like a thing, and suddenly headphones get ripped off and you're like, hey!

Yeah. Nature's out to get you, [00:02:00] Dave. Everybody's out to get you. You should be paranoid. The amount of confusion I had walking forward. Did that put, did the tree go on your annoyed things list? We were talking about the other day that, you know, you often have. Yeah. I don't know whether I was annoyed about it.

I was more sort of unnerved and weirded out because I couldn't, couldn't tell whether the tree had done it on purpose. Well, I'm glad you're also confused because this way, well, this week, Dave, I'm calling it out a bit about what, why I'm confused, which is the, uh, And I'm not talking about your elderly generation, or your young generation.

I'm talking about, like, your working age generation technology. And I'm confused about why that group of people still have broadly crap security skills sometimes. I'm going to use an example, and I know a lot of people might go, Ooh, yeah, is people who have the same password across multiple systems, websites, or whatever, or people who think they're being clever and just up the number at the end of the password every time they change it, etc.

I'm And in this world, where everything's getting hacked, there's stories everywhere about websites [00:03:00] getting owned and your details going out, people still maintain same password, multiple systems, and it just still, to this day, confuses me. And I think it just comes down to, they can't be arsed. So maybe you could do it, yeah, I think that's, I think you've, I think you have got to the bot, the right, you've got the root cause there, right?

But Now, so let's talk about what Rob's good practices are for password management, shall we? Yes, please share Rob. I think, I think, you know, not, not to say that shall Marcel and I are guilty of this, but let's just assume for the sake of argument that we might be. It was a bit of a look when I said it actually, I'm looking around the group and going, hang on a minute, is



one or two of you among us as well?

What would, just, just say just for the sake of argument that we might be, what would your advice be Rob? Give us some tips, Rob. Yeah, give me some tips. Well, let me get some pen and paper.

I love it. Suddenly, uh, this weekend, [00:04:00] Dave, Marcel and Schauke are changing all their passwords on websites. Maybe, maybe. 20 hours later. Well, there's some basic ones. It's dead simple, right? So it's, um, obviously, um, different password for every system you use. Yeah. Use a password manager, pick a long password.

So 13 plus 14 characters, mix them up. Don't use obvious phrases, maybe mix two completely unrelated words together. Put a mix of characters and symbols in and don't do the all my zeros. Are my O's really and E's or 3's because that's really easy to guess and then it's down to if it's a really important account like where your email goes or you store all your personal data, two factor authentication, it's really easy app goes on your phone and then if somebody does own your password, they can't get in anyway, right?

And it's dead simple. And that's why I'm confused because people continue to not have what I would consider to be a pretty basic level of security on the consumer advice by [00:05:00] Robert Kerner. I still get, I still see it and they go, Oh, don't do that or short passwords. Do you use, do you use all of that? Yeah.

Do you live in Breda? In fact, my whole family is drilled in the ways of good password security. But how do you remember all these passwords, Rob? Ah, see, password manager, which you put a very, very strong password on the front of. You know what Rob's password is? It's, it's Rob's password manager, all one word.

One, two, three. Well, there's that website with all the, um, with all the, there was that where you say people still use password one, two, three, admins use I am God, things like this. But there was a study done and reporters were stopping people in the street and say, will you tell me your password if I give you a chocolate bar?

And an alarming amount of. People gave up their password for a bit of chocolate, and my head is exploding. You should have saved that for a cyber episode [00:06:00] on that note. Well, I think not only, not only have we got the root cause of why people do that in that confusion, but also some consumer advice that might be one of your best ones yet, Rob.

And, uh, if any security experts think I've just given some bad advice, please let me know quickly because I need to change all my passwords. On that note, let's welcome our guest today to help us with the thorny subject of FinOps. I am delighted to say we've got Kunal Agarwal, who is the chief executive officer and co founder of Unravel Data.

Kunal, good to see you this morning. Do you just want to say hello and introduce yourself and Unravel? Yes, thank you so much for having me here. I'm Kunal Agarwal, co founder, CEO of Unravel Data. The company is the leader in data observability, um, and FinOps for the data analytics stack. That's really what we're designed for.

And we do it in a very automated fashion to really simplify these activities using AI. [00:07:00] Before AI became cool, which was when October 2022 company was founded in Duke University about nine years ago now. So very excited to be working with large scale data teams. That's the best part about the job, getting to work with some really amazing talent across data scientists, data engineers.

In variety of different industries. So it's usually folks that are running large scale data systems that have a need for something I can gravel. So why don't we kick off the



conversation by just talking about the requirement for FinOps like lots of organizations, I think. Go to the cloud, get onto the cloud, they may or may not have built the capability to manage their cloud, and often it can be the case that organizations costs start to go in a way that they hadn't modeled or hadn't business cased.

So let's use that as the jump in Kunal, is that a scenario you recognize? Yeah, absolutely Dave, and then even more [00:08:00] so on the data analytics side. just by the sheer volume of data that people have to go and process. And then all the use cases that are getting blown up on data analytics, machine learning, AI.

So you've got this massive amount of data and these influx of use cases coming in, and that starts to balloon even faster. So just your overall cloud build for your web apps

But then when you look at the data analytics, Zach, like that is its own beast. And that scenario of that kind of, let's call it cloud sprawl for want of a better term, or you know, there might be a better term you can introduce us to for what that looks like. Would you say that that's a relatively common issue in organizations because one of the things that really actually annoys me and I would, I would say that being the many things annoy you.

Let's be honest. This is just one of the long list. Long list. That's true. Wrong. You are right. I'm on it. I think if I'm brutally honest, [00:09:00] you have days when you're on it, when you're not on it, just depends on how the days go, doesn't it? One of the things, is organizations that are, you know, making a fuss about decamping from the cloud because they haven't made the cloud economically viable for their organization now in a very minuscule amount of cases, I can perhaps buy some aspect of that, particularly if that organization is like a very, very highly skilled tech organization and they think they can bring it back and run that tech themselves, maybe, but I wouldn't have said for the majority.

that that's the case, and often cost is cited at the heart of that. So from a FinOps perspective, then, scenarios you recognize and what do you advise your customers in those sort of situations? Yeah, um, look, cloud has a lot of advantages, especially when, again, you look at trying to run these data analytics workloads on premises and getting [00:10:00] back to your data centers is good when you have predictable, constantly running workloads that don't fluctuate a whole lot.

So you've got something that you just have to run, and it runs every day or every week. And it's predictable that it needs so much horsepower to run, and it's good for doing that. When it comes to data analytics, people are still experimenting. And there's more people jumping on the data analytics stack than just the hardcore IT folks inside the company.

Marketing is jumping onto it, and they're trying to understand customer behavior, for example. Product is getting into it because they're trying to understand how to improve product. Legal is getting into it, right? Every department in the company is trying to do it. So if you want to democratize that and have people experiment with the data that the company holds.

to say, Hey, what can you glean on the insides and understand how we can improve the customer experience in our product, then there needs to be that [00:11:00] experimentation DNA inside the company and the cloud is much better for that. Secondly, is even if you have workloads that need to run every day, like, say, for example, a recommendation engine.

Or a fraud prevention program, the amount of data that you're processing on Monday may be different than what's happening on the weekend, for example, or the amount of data you're processing on October may be very different than what's happening in December



because of the amount of, you know, people that are active on your products.

And cloud is great for that. So it gives you the elasticity. It gives you a way to have infinite, uh, resources and compute, if you may, to go and scale these operations up. So all those things are positive. However, with having infinite resources, comes an infinite bill. So having those guardrails around, what is healthy?

What is good? What's our budget? What, what are [00:12:00] the boundaries in which we need to operate, those need to be defined, and those need to be monitored, and only then can you have an efficient usage, but that's, that's really the, um, you know, the two forces, right? On one side, you want to go and experiment, you want to allow everybody to jump on the cloud, and the other side, you're trying to curtail it and curb it and organize it in a way that, uh, you know, you try to, uh, grow this efficiently.

And I think that's a, that's a healthy, uh, fight between those two sides, especially in innovative periods for any industry and company. How many organizations you see right from the outset have finance involved in this process? So like the way that I T organizations, which is typically where main cloud is getting managed today, they often have from the last generation of I T a certain set of processes in place for how they work with finance.

And I think the cloud and finops demand new relationships, doesn't it? So before we get into you. Some of the mechanisms by which you can manage your cloud [00:13:00] more effectively. Maybe talk to me about some good examples and maybe some not so good examples of how you've seen finance as a department with organizations interacting with the cloud.

It's a mixed bag. If you've got workloads that are being migrated to the cloud, then finance does get involved pretty early because they're trying to scope out and understand, Hey, we're moving our data centers to the cloud. What's that going to cost me? Hmm. Where it gets challenging is when you have new workloads coming on the cloud and there's no idea of what it should be, what it should cost, and what the guard drills could be.

And that becomes a rather, uh, combative exercise between the two departments inside the company where they go, oh, I thought it's gonna be X dollars, but, or X pounds. But it actually turned out to be 2x and 3x or 4x of that. I mean, not even through half the year. So unfortunately, we don't see finance being involved in the planning stage.

That's really where they need to start. Instead of the policing stage, which is at the end of the month. If you're there in [00:14:00] planning, and you're able to understand from the business what the requirements are, and then involve the technology side of the house, in our case, data engineers, data scientists, Along with the people who are trying to create these new innovative data products to understand, hey, look, there may be an experimentation phase, which is, you know, we're going to run these four or five different experiments for creating a new recommendation engine, for example, and then there's going to be a stable or a more stable production phase.

And these are the kind of budgets that we need to start working towards. But that happens during, um, the design and the whiteboarding phase itself. Um, and then being able to give these folks access to various systems that they need. So that plays into a conversation that we have a few times, which is the difference between efficacy and efficiency.

And when you talk about new workloads, people who are shooting for New capability don't always have efficiency in their minds. I and we have this debate about if you combine [00:15:00] efficiency right up front, how do you do it? How do you do it effectively? So as you



go live, it's efficient. It's effective. It actually computes, you know, with reasonable amounts of cost control.

But what many do is as they're doing something new, efficiency is not really in the mind's eye. So there's always this phase that has to kick in quite quickly. quickly behind it. And that's the kind of debate about how do you mix the two? How do you make it effective in your view? Do you see that as a people shoot for efficacy and then go?

Whoops, need efficiency. Or have you seen anybody successfully combine the two together? Because it's a it continues to be a debate that we have about how you do new cloud native things and try and combine the two together. Yeah, I'll give an example. Um, there was this blog by Shopify. If you Google Shopify BigQuery, Costs, this block will come up, they were trying to launch a new data pipeline and the engineering groups inside of Shopify figured that they needed to understand what kind of data scans would happen for every run of the data pipeline and they did that in test.

They [00:16:00] saw, I think it was about 75 gigs of scan. So BigQuery, which is this data analytics platform on Google charges based on the amount of data scanned, and it was 75 gigs of data for each run. And they figured this pipeline, which is a marketing tool actually, would be called 60 times a minute. So every second is about 2.

5 million times a month. And if you multiply those numbers on BigQuery's cost, they calculated this is going to be a 950, 000 data pipeline. Wow. Now, the engineering prowess of Shopify is such that they figured, look, can we use some techniques? And in this case, they use this technique called clustering.

Which is, think about it as a way to organize your data in a table to make scan lighter. And they were able to get that scan down to, I believe, 0. 1 or 0. 2 gigs from 75 gigs. So that same 2. 5 million runs would now cost them 1, 300 as [00:17:00] compared to almost a million dollars. That's the dramatic difference you can see if somebody is proactively understanding and looking for the different optimization or tuning techniques that can make these data applications actually run faster and cheaper.

Unfortunately, not everybody has that kind of engineering prowess as the Shopify data engineering team has. There's very few companies out there in the world. And then secondly, it's hard. There's a lot of steps to go through to even understand this. I'm, I'm sure I, I'm not sure how many listeners on this podcast even know what, uh, clustering is or what are the other techniques that you can go to to go and, you know, bring these kind of costs down.

And that's why people look at it in hindsight. But if there was a way to go and understand these problems easier. Oh, there's a way to educate these audience about all these different techniques and things to watch out for, and they were [00:18:00] incentivized to go and do these things in the get go, then you would see more of these things happening proactively.

I think the point you make there on engineering prowess and organizations that have got Very, very strong engineering can really be on top of this stuff and really understand the nuances and the detail in all the different layers. You have to get right to have an optimized position. Of course, no disrespect to anybody, but a lot of organizations that are transforming into this sort of world don't bring that prowess with them.

So this must have been something at the point You guys were thinking about unravel and when you thinking about finding an organization that you could see as a problem so what do you want to take a step back from the specifics of an optimal come back to it in a second just to talk about your founding journey what were you perceiving is the problem when you were



thinking about unravel and how did you go around so setting a purpose.

It's exactly that. We've, we figured that all these data systems are gonna be used by a broader range of people with varying [00:19:00] skillsets. Um, you're gonna have a person that knows marketing really well. Or a person that knows product really well, that's going to jump into these data analytics stacks, rather than having the IT teams only jump on these stacks.

And that means mistakes are going to be made. That means things need to be simpler, for it to be a level playing field that everybody can truly extract the value out of this. So at Unravel, we believe if we were able to automate a lot of the things that a data engineer or a data scientist would do or need to do to get both performance and cost efficiency out, then we would, number one, help them reduce the amount of toil and time that they're spending doing this stuff.

And secondly. More people could get on the data stack and ensure that their data pipelines and applications are running well. And thirdly, they don't have to have a [00:20:00] PhD in how distributed computing systems and things like that work. Uh, honestly, and when my co-founder and I, we met actually at Duke University in in America, he's a professor of computer science at Duke.

Uh, and he was researching database optimization techniques. And I had seen in 2013 2014 time frame when quote unquote big data was just coming out being used by some of these companies, some of these large corporations. That was a light bulb moment where we saw this is such a powerful technology, but it's so rough around the edges.

That's not a complete product that everybody can use. And if we were able to solve the security problem or the governance problem, and in our case we took on the observability, the cost efficiency, and the performance management problem, then it could be used really broadly by many more teams and many more organizations.

We have now gotten support for a wider variety [00:21:00] of systems than there used to be in 2013 2014. There's multiple clouds, multiple systems that are running on clouds, because our promise really is it doesn't matter where you're running these data analytics systems. If you're running them, then there's this thing on the side called Unravel that's making sure that everything is running as a well oiled machine.

Your applications are finishing on time every time and your efficiency is something that you can automate and don't have to worry about that. And have you got a particularly good example of an organization that you can name or even not name that you went into initially that was struggling with this stuff and what was done then to unpick it?

Because I know that lots of organizations are at that point in their journey at the moment, I think. Yeah, um, every industry is jumping on the ML, AI, data analytics bandwagon. I'll give you an example of a company that's not a high tech company. So it's a logistics company, one of the biggest ones in the world.

[00:22:00] And they saw their logistics business just boom in the pandemic era. And they realized that they could use the data analytics that was being generated from all the things that were being shipped, all the planes that they had to really get customer service done better. And they needed to do this on a large scale data system because of the data points that they had.

Millions and millions of data points coming up at them every hour. But what they realized is the systems that they had and the people that they had didn't know how to go and operate



on these large scale distributed systems. So they had about 3,000 engineers that came from, let's say, the Oracle world.

onto this new Databricks, Snowflake, BigQuery, you know, this modern data stack world. And they needed to get their outcomes done immediately. So they couldn't wait for three years to train these guys and then get them onto the stack because they needed that analytics. So they [00:23:00] jumped right into it. And a lot of things went sideways.

Number one, analytics project was delayed. People didn't know what to do. The analytics that they had wasn't reliable. Things were failing. Things were broken. And the costs were actually 3x, I want to say 3 or 3.5x of what they originally budgeted for. So it was very high. And what they, what the root cause of all of that was people just not knowing they were making mistakes and how to go and resolve those mistakes.

So they try to do it in a manual way, they try to train people, but that takes its own time. That's when they partnered with us, Cogs Inside, and what we did for them is we scanned their entire data environment, all the applications running on top of it. And then the algorithm automatically started to surface all the different issues.

around reliability, around cost, around performance to the [00:24:00] people that were running these applications and then told them, this is a problem that's happening, this is why it's happening, and this is how you can go and fix it. In plain English, so that they could go and resolve those things. We saw about 78 to 80 percent reduction in cost in nine months.

Um, so it's pretty dramatic, pretty quick, but more importantly, the folks were upskilled overnight because they were seeing their problems and why they were causing those problems in real time. And we started to see a decrease in the repetition of those problems go down. Where people now understand, ah, this is what I keep doing all the time, I should stop doing that.

And now we see these guys, you know, having a greater throughput in how they're creating new applications and bringing them online. And everybody just be more efficient on this environment. I think that goes back to the point you're making earlier about engineering excellence. I've no doubt that the people who transition to the new platform were excellent engineers on the old platform, but they didn't understand the new.

There was a rush for that [00:25:00] functionality. And guess what? A three X increase over what you expected. And then, oh, dear, we need to do something. But that's also a very dramatic story about how you can get the cost in line much, much faster. I'm very sure there was a rapid learning curve for many on that platform.

Those for those few weeks. You're absolutely right, Rob. And unfortunately, fortunately, people have had 15, 20 years to learn Oracle and there's playbooks and best practices. And, you know, you have amazingly smart system admins to take care of all these different types of problems. But the speed at which innovations coming to us on the machine learning AI stack and the demands of the market to go and produce measurable, impactful outcomes quickly and the pressure of that, those things put together is kind of making efficiency and cost sit in the back seat.

Where people are saying, I don't want to be focused on that, the engineers are, but then the business is going, how do I go and get ROI for these hundreds of millions of dollars [00:26:00] of investments that I'm putting on top of these things? So it has to be hand in hand, but the timing is such that people think that they can only pick one.

But there's definitely ways in which you can educate the market and get them to do both.



Very good. So if we return then to FinOps more generally and talk about just some good practices. So when you when you look across the industry, what would you say are the three or four strong good practices that people should get in place to ensure that their workloads are being managed effectively?

So we think about FinOps in a couple of stages. So number one is Inform, and then uh, Govern, and then the third one being Optimize. Um, Inform is exactly that. Which is, let's at least understand what's happening on the environment, and understand who's doing what on the environment. You would think that sounds like a simple enough problem to solve, but it's not.

We, we see incomplete Inform phase in a lot of companies that we start working [00:27:00] with. They may have just a high level bill information. How much did I spend on my cloud? Okay, a million bucks, but how much did Dave spend? How much did Rob spend? How much did the marketing department spend? That kind of information is opaque or incomplete in a lot of cases, and if you don't have that information, then you cannot do the governance and the optimization after that, because you cannot govern what you can't see.

How am I going to put guardrails on Rob's workloads, for example, if I don't know if I can measure his workloads not to pick on Rob, but that's just an example of that. And Rob's not going to be a rogue user. And then you start to think about optimize, which is hey, if then I can understand. Rob's usage out of that million dollars a month was 100, 000.

Now I need to put guardrails on 100, 000, for example. How do I know that those workloads could not run for 50, 000? Um, that's where the optimize phase comes in. Is there ways to improve and help Rob complete the tasks that he wanted to do [00:28:00] without costing as much money? So those are usually the three phases that we see complete that FinOps journey.

So the best practice is always To measure everything from the get go, measure what applications are running, measure what systems they're touching, measure what data they're using, map out the org chart, and understand who's doing what, what departments they belong to. What projects they belong to. People use techniques like tagging or some sort of a identifier to go and understand one workload.

So once you have things at the most granular level, then you can start to group them and bucket them and start to do analytics on who's doing what. That itself would take companies a long way in their FinOps journey because then you have data to back your questions. It's not a gut feel. You should be [00:29:00] able to go inside and do a query of, what does this particular user or what does this particular application do and how much did it cost and what efficiency score does it have, especially for customers now who are running multi system environments.

multi cloud environments even now, this becomes an even bigger challenge. So having the telemetry data, being able to connect the dots between multiple systems, multiple environments is the only way that you're going to get the complete picture. That's where I would start and then start to do more advanced techniques of optimization and looking for inefficiencies and hunting them down.

Hello just for the avoidance of doubt they don't let me near any production environments anymore i get rugby tackle before i get to the keyboard so uh that's more of the governance they put around me these days on something uh rob there was a very good point about you can only govern what you can measure and.



See, and that is often, unfortunately, the [00:30:00] whacking great bill that comes in that finance. So we're gonna have a call about this on Monday morning type thing, and I think it's that early. It goes to the point about engineering excellence, the front early sight of what the compute requirement of the payloads likely to be.

And then the sort of it went live. Is it doing what we expected to do? Rather, what most people do is publish it to live lovely automation. And then let's all go down the pub, come back on Monday morning and go. I did not expect that bill. You know, we've learned from the companies that we work with is yes.

So we can all sit down here as finance people and business people and say, this is the bill, this is what it needs to be. We need to cut this down. But at the end of the day, the engineers are the ones who will make the changes and can actually impact change. And that's usually the part, which is the toughest part, is to get people to take an action, is to get people to actually focus on that efficiency piece.

And what we've learned in organizations, the reason why that is, is twofold. Number one, they don't know what to go and do. [00:31:00] And number two, they're not incentivized to go and do what they need to do. So that first part of knowing what to do comes back to, do they have the skill set? Do they have easy way to understand what they need to go and do?

That's the kind of part that companies like Unravel are trying to solve. And the other part about incentivization comes in two folds. One is, Is the culture of the company such that we are a cost conscious, efficiency conscious culture? Or is this the wild, wild west? And that's something that leadership really has to drive.

And we've seen this done in multiple ways. I'll give you a couple examples. We've had companies where they have leaderboards now. So you put up leaderboards in the old offices, it would be this on a hundred inch monitor in the pit where all the engineers are sitting. And those leaderboards could be either way.

It could be the wall of fame or the wall of shame. You do not want to be on the wall of shame. Where the top ten most inefficient, expensive users are popping [00:32:00] up. And you want to get off that and you're going down feverishly trying to optimize things. And the wall of fame is where good behavior is rewarded and you're a good citizen on the platform.

We've even seen companies take it, you know, to the other extreme dramatically where they actually reward people and sometimes monetarily. To go and improve things on the environment, we've seen companies actually share a percentage of the cost savings that they got on the cloud back to the engineers that impacted that change, and that's driven people to go and drive costs down really, really quickly as well.

What you've been looking at this week, Sjoukje. So each week I do some research on related ideas and transformation and tech. And this week, I thought we should take a look at how will AI influence FinOps models in the future. So AI is going to play a big role in future FinOps models. These [00:33:00] models may leverage AI to automate cost optimization, streamline the financial processes, and provide real time insights for decision making.

So with the evolution of FinOps, diverse pricing models are also emerging. Pay as you go, pay for savings only and other gain share models are becoming more attractive for enterprises with complex cloud estates and high cloud consumption. So some features of AI powered FinOps in the futures are advanced predictive analytics, AI powered decision support, accurate invoice analysis, and AI driven risk management.

So a question, is this still the future, or is this already here? AI powered FinOps or AI powered



cloud cost governance is here. The adoption for technologies like this is in this nascent stage. And the way this technology gets adopted faster is exactly what's [00:34:00] happening right now, which is the cloud costs.

Along with the usage are not growing linearly, but exponentially when you look at the use cases and the kind of outcomes that companies are trying to drive, take an example of LLMs and AI models. Um, there is a huge, huge need for, uh, checks and balances on how much computing horsepower will these models and AI engines actually need, require, and burn down.

So, when you also combine that with the things that a engineer has to do every day, on one end, you're thinking about, let's go and create some great data outcomes, and on the other side, let's make sure that these costs are efficient and Performance is reliable. It's too many things for a person to do, and it's too complex of a problem to go and solve because there's so many things that can go wrong.

It's actually better [00:35:00] a machine solve the problem than a human solve that problem. Yeah, I think, I think for me, there's a lot of talk about AI going into developer productivity, and I think that's about writing code and cutting code. I think there's an angle here as well about the AI saying. That algorithm will work, however, is deeply inefficient.

Have you considered or as the developers writing it, the AI is watching it and saying, this is going to loop funny. You're using this data item a lot. Are you sure? And it's like that prompting. So I'm a big believer in good decisions as far up the stack as you can to avoid the bigger problems and the power wave down the stack.

And I think there's a, I think it may creep into code automatic code optimization as the lines are being cut. So you get this, I mean, if you ask an AI to do it, it probably already comes pre optimized. But if the human does it, the AI, prompting back, saying, You sure, Dave? Because that looks deeply inefficient at the bottom.

And you're at the bottom of the wall of shame this week, so maybe you don't want to push that out. Actually, AI generated code [00:36:00] is great for getting the code out, but that code is not optimized. There's actually more unoptimized code that AI is generating than humans are generating. Because the code that that AI has been trained on is static code.

And that static code is not checking for. Is this efficient? It's just checking for, will it get the answers out? Yeah. Yeah. So as more and more code is being generated through ai, we're actually seeing more inefficiencies creep up in the system. So it's good for speed and getting things out, but it's not good for efficiency and growth.

But Rob your point around. Can we shift this as left as possible so it's not being caught in production? Can we do this in the planning phase? Can we do this in the testing phase? Can we do this in the CI/CD phase of moving these applications from dev to prod? That's exactly where you should be doing this.

In fact, an analogy is you should not Get a fire extinguisher to do when your house is burning down, you should have a fire extinguisher in your house beforehand itself, [00:37:00] or even try to avoid that fire. That's exactly the way that company should be thinking about these pieces. And now the tech exists to go and make those pieces happen.

So the technology that that we have through Unravel, you can automatically check for inefficiencies and do code reviews with efficiency, performance and reliability in mind before that code hits production. Great. Kunal, thank you so much for taking the time this morning to talk to us. That was a great conversation.



Thank you. It's fun being here with you guys. Really good to talk to you. And we end every episode of this podcast by asking our guests what they're excited about doing next. And that could be something in your, in your personal life, or it could be something in your professional life. So Kunal, what are you excited about doing next?

Oh, um, personal life is always about cars. That's what I look forward to every weekend. Starting to do this new autocross series, uh, which is once a month, and it does have a little bit of, uh, uh, [00:38:00] rivalry there, although it's not professional race car drivers. It's more gentleman drivers, if you may. Um, but I don't know if you guys have done autocross.

It's, it's a rather fun thing to do. Even before we get onto autocross, I like the idea of being a gentleman driver. Speaks to your personality, Dave, quite well, actually, that. Have you got the vision of the turtle back gloves, the tin of sweets, and the cap as you're driving along? Yes, Goodwood, Goodwood revival, you know, that kind of stuff.

Yeah, yeah, uh, the Goodwood are a blast. I'm there almost every year. Oh, you really? I live really close to Goodwood. Oh, really? Okay, to Chester. Yeah, come on there. Yeah, if you're over in the UK in the summer, when Goodwood's on, do give us a shout. We could go for a beer or something. Oh, for sure. We actually sponsor one of the things in Goodwood.

So we bring in our partners and customers there as well. How exciting. It's a blast. It's the best car show that I've ever been to. But yeah, autocross, you know, we leave. rent out a [00:39:00] runway or two, or a big parking lot, if you may. And then we make our own technical track with cones. So you're not really hitting top speed.

You're doing 60, 70 miles an hour on a, on a good stretch at most. But it's really technical, so you've got slaloms, and u turns, and fast corners, and things like that. It's about who does the best lap time around it, so it's a blast. Plus, you can't really kill yourself, so it's pretty safe. That, that sounds incredibly good, like, something extremely risky, but it's like playing, playing, uh, 10 pin bowling, but with those gutter, you know, the gutter bouncer things on it, the gutter inflatables, a bit like that.

Exactly, exactly. Because it's one car on the track at a time. So yeah, accidents rarely happen. Worst thing you can do is get a couple of cones stuck to your tires. Well, that sounds brilliant. So a huge thanks to our guests this week. Kunal, thank you so much for being on the show.

Thanks to our our passive aggressive producer Marcel, our sound and editing [00:40:00] wizards, Ben and Louis, and of course, to all of our listeners.

We're on LinkedIn and X, Dave Chapman, Rob Kernahan, and Sjoukje Zaal. Feel free to follow or connect with us and please get in touch if you have any comments or ideas for the show. And of course, if you haven't already done that, rate and subscribe to our podcast.

See you in another reality next week

About Capgemini

Capgemini is a global leader in partnering with companies to transform and manage their business by harnessing the power of technology. The Group is guided every day by its purpose of unleashing human energy through technology for an inclusive and sustainable future. It is a responsible and diverse organization of nearly 350,000 team members in more than 50 countries. With its strong 55-year heritage and deep industry expertise, Capgemini is trusted by its clients to address the entire breadth of their business needs, from strategy and design to operations, fueled by the fast evolving and innovative world of cloud, data, AI, connectivity, software, digital engineering, and platforms. The Group reported in 2022 global revenues of €22 billion.

Get The Future You Want | www.capgemini.com



This presentation contains information that may be privileged or confidential and is the property of the Capgemini Group. Copyright © 2024 Capgemini. All rights reserved.

