EPISODE 1496

[INTRODUCTION]

[00:00:01] AD: Couchbase is a distributed NoSQL cloud database. Since its creation, Couchbase has expanded into edge computing, application services, and most recently, a database as a service called Capella. Couchbase started as an in-memory cache and needed to be re architected to be a persistent storage system. In this episode, we interviewed Ravi Mayuram, SVP Products and Engineering at Couchbase. To learn more about Couchbase, check out couchbase.com/sedaily.

This episode is hosted by Alex DeBrie. Alex is the author of the DynamoDB book *The Comprehensive Guide to Data Modeling with DynamoDB*, as well as *The DynamoDB Guide: A Free Guided Introduction to DynamoDB*. He runs a consulting company where he assists clients with DynamoDB data modeling, serverless architectures, and general AWS usage. You can find more of his work at alexdebrie.com.

[INTERVIEW]

[00:01:07] AD: Ravi Mayuram, welcome to Software Engineering Daily.

[00:01:10] RM: Glad to be here, Alex. Thanks for having me.

[00:01:12] AD: Absolutely. And you've been on recently talking about different things going on at Couchbase. We're going to talk about this cool report that you came out with at Couchbase, about enterprise adoption, the cloud, things like that. But before we get into that, maybe just start us off with a little bit about your background, what's happening at Couchbase, things like that?

[00:01:29] RM: Well, my background, I've been here at Couchbase, CTO, running engineering, other functions, for almost nine odd years now. It's been a fabulous journey of building a modern database. Truly advancing where databases have been. There has been much innovation here for the last almost four – The last round of innovation was in the early 2000s, late 90s. And so,

it's been a fun time to be in this data-oriented space. As you know, it's all about digital transformation, which is all data. And glad to be part of this journey.

I started my career writing some very low-level stuff for HP, doing operating systems work. And so, from there, the career has sort of brought me over here. And whatever I learned along the way, brought most of the knowledge to what we can – How to modernize the database stack.

So Couchbase, we are on a very exciting journey. We have been a successful company that's been around for almost 10 years now. And the journey from a very small startup of a dozen odd people to now a public company has been a really exciting one for all of us here.

[00:02:28] AD: Absolutely. I totally agree with you. Like, the database space, especially what's going on with cloud and databases is really, really interesting one right now. And it's cool to see some of that stuff. On that sort of topic, we're here to talk about the cloud. And y'all released this report, the Couchbase cloud evolution report in 2022.

I have a few specific questions, but maybe just give me what's the overall feeling that you're getting? You talk to a lot of enterprises, I'm sure. They're moving to the cloud. Using Couchbase in the cloud. Like, what's the feeling you're getting about cloud adoption?

[00:02:55] RM: I think basic thing is that Cloud is – As you know, many of us have done a lot of stuff in the cloud. It's nothing new. It's been around for more than 20 odd years. But now it has come to a point of maturity, where sort of the utility computing, we are in the sort of the first phase of it so to say. And that's mainly driven by the cloud. Everything being available as a service, as opposed to the good old days of sort of procuring your own hardware, procuring your own software, having the full-on it experience, so to say. And that is now gone away to being things being available on-demand, on-tap.

So, it has gone away from the product, install, deploy, configure. That mindset to a service that meaning things are on the tap. And evolution is, first, people started to go to the cloud for convenience. Now, it is more mature where people are going there for basically a metered consumption. Otherwise, you have always had to have this huge OpEx, CapEx sort of discussions. Those have gone away to, “Well, I'll pay as I go,” sort of mindset. That's the

maturity of the cloud. Being able to offer you that ability to pay for only what you use. And it's got many different names. But I think something like serverless onwards, there are many ways to describe that movement. But that's where the industry is. And it's now gotten into the phase of adoption from being like early adopters, and early majority to where it is coming to use the classic convention of an evolution of any technology, it has come to a point of early majority to late majority, and maybe even laggards adopting it at early stages of that.

And so, then it becomes like a full-on movement where pretty much it feels like the conversation is everybody is like, “I got have some of that cloud?” How do I do it is the only question. Why should I do it is no longer the questions. So, that's the flip that you have. Actually, and the pandemic has accelerated that in a major way.

And so, it's fantastic to see. This is what many of us have built in the past, this whole model of elastic consumption base computing. You want computing to be just along the same lines of where the rest of the utilities is. Like, power and water and stuff like that. You turn on the tap and you use it. Computing should be the same way, which was not the way it was. Even now, it is not quite like it. But it has to get to a point of you just – For plugging on a toaster, you don't go to your utility company and ask them to provide something more. You simply use it, and then the billing comes at the end of the month, and you pay for it. That's the hope.

And in the case of software, there are a lot of other things to it, because it's not just a commodity. There's a lot to be developed on top of it. And so, people want to see certain details being hidden so that they can do their piece. Now, they're left holding the bag of managing a lot of infrastructure, which is not their core competence. And that goes away with cloud so that they can concentrate their efforts on building what will add value to their business. And that's the journey. It's exciting to talk to many organizations, which are also transforming that are newer sort of leaders or the CDOs, and the CTOs, and CIOs are transforming to be people who can digitally transform. And so, these are people who come from a development background, as opposed to people who used to typically be more from sort of IT and procurement sort of background. So those kinds of stuff is changing. And so, it's exciting. And everybody understands that data is that, so to say, has been said a while ago, it's a new oil in one sense, and how to – Everybody understands. And how to sort of take it to the next level of evolution.

[00:06:33] AD: Yeah. Let's talk about that specifically. Like, one of the takeaways I had is lots of respondents were sort of specifically calling out, “Hey, I want my database as something that's in the cloud.” And why is databases, why is data – Like, why is that a focus area for enterprises moving to the cloud?

[00:06:48] RM: Yeah. I think the biggest thing is that databases were sort of – The previous versions, the relationals and the most popular ones were built at the times when network was sort of not even available. They were all built like single node systems and stuff like that. Now, to scale them and run them, many advances have been done, but they all require specialized skills. DBA is a function.

And in one sense, and some of the transformation, that role hasn't gone away, but it has transformed to being like SRE’s, so to say. So now, enterprises do not want to invest more of their time and resources and training people on newer technologies, so to say. I don't want to add more DBAs. I would rather sort of add more on the SRE side of the equation, so to say. Specialization is not required.

So, in the database area, the requirement is to say that, “Well, you take on the database-specific administration skills.” We will have our generic SRE people who can actually do more of the stuff that is relevant to us, not just – It’s almost like using an automobile analogy. We are requiring them to be like both mechanics and drivers.

[00:07:58] AD: Yeah. And one interesting thing I found at the end of the report is like, hey, a lot of people are actually saying, “I want more control over my database and managing my database than my cloud provider gives me,” which is kind of interesting when I think about cloud. And even the points you're saying, like, “Hey, I want to get out of the DBA game and get into more of the driving game.” Like, what do you think's driving that desire for more control?

[00:08:18] RM: That's a good question actually. What has basically happened is that, yes, now I don't have to learn about specific database skills. That's because it's taken care of in the cloud. But then, because of that, now, most of the sort of the movement in the cloud is actually led by developers in one sense. Now, they can go use – But since it’s easily sort of available, you can use whatever now.

All of a sudden, what you'll end up seeing after a while, very quickly, is that now you have a sprawl to manage instead of just that one thing that you were managing earlier, because many of the other options were not available or provisioned. Now, all of a sudden, people are losing 5, 6 different data-oriented infrastructure over there. It could be analytical. It could be a search. It could be a database. But at the end of the day, it's data management as a space, if you will.

Now suddenly, you're left with a situation of the sprawl. And managing the sprawl, there are three main issues that come over there. And so, you're hitting the other side of the equation, “Hey, how do I contain this? How do I manage this?” is the conversation there.

The three issues become – One is just the fact that I have N different sort of database or data applications that are running, which I need to, number one. Understand. Number two is that, as a consequence of that, the threat surface area just became exponential. Now I have to – How is the security of this data? Is the second thing.

Third thing, which doesn't directly come out from this, but people are beginning to understand. If you really think about it, what you're actually doing is that by consuming more of this, you're actually paying more in one sense, because you created the sprawl. Mostly end up duplicating data. And then now you're managing that data. And there is a lot of egress costs that are like you did not think of it. Or maybe it was even okay when you started. But when things mature, it pretty quickly adds up.

And many cloud providers, it’s not by design or anything. It's nothing malicious here. But generally, because you started using, there is no central way of accounting for all of this stuff. That's the second thing that typically happens. How do I even know where are the different places where it is being used? If you go to your CIO or someone to ask the question, they won't be able to give you a concrete answer, because there is no one classic way of rolling all this thing up.

And so, then you suddenly find that you're sort of overspending. And in this report, you saw enterprises saying, on average, like they're spending about $9 million more and stuff like that.

Then this is for a simple case of going from on-prem to the cloud. They haven't really fully even leveraged what all they can do.

I think it's in that where you – It's a good compare and contrast that you point. On one hand, you say, “Look, I don't want to manage.” And that's what they're saying. But now you're saying managing in the cloud is becoming a bigger problem. What is it? It's two different problems sort of conflating at a certain level.

[00:11:39] AD: Yep. And I want to push more on that billing point and cost. Now, I hear a lot about cost overruns. It showed up in the report as well. You sort of mentioned how cloud becoming more like usage-based billing. More like a commodity. Like, electricity or utilities. What are you seeing? Do you think people want that sort of pure commodity where it's like electricity, but it's also a little more maybe unpredictable, right? You find out at the end of month you have a big bill? Or do they want some predictability? Maybe you provision an instance and you have that sort of predictability on, “Hey, this is how much that instance is going to cost for the month.” What are you seeing in for people there?

[00:12:10] RM: Yeah. I think this is the other tough part of the equation, which is that the paradigm that has shifted is that, earlier, you will go with, “Let me understand how this application is actually going to do in real life. And let's go provision based on that.” It was a huge discussion, because there's a CapEx element to it and all that stuff. Now, it is the other way around to say that it's okay to say I don't know. I don't know how. Is it going to be a bang? Or is it going to be a bust? I don't know. I'm prepared for both.

That's what the cloud is actually providing. In computer science terms, the pay as you go and all is like at a commercial level. But infrastructure wise, what the core defining characteristic of a cloud is elasticity. It gives you the ability to expand as much as you need and contract when you don't need it kind of stuff, right? That's a very difficult thing to pull off at a certain level.

Having that is not the problem. But like you're saying, how to manage that to have discipline within that? And there are operational issues within that, which is that now that things are more easily available, it's the discipline of what you want to use and not just go use everything that is available. It's like kid in a candy store situation, right? You come back with a bag full of candies,

and some adult has to tell you that can have only so much kind of stuff. So that's the transformation that it is sort of going through. It's a part of evolution. Perhaps the pendulum is a little bit more on the one side of sort of overconsumption. And everybody understands. But it's part of the maturity that people will have to self-govern. And maybe the cloud providers will have to meet them halfway by giving them the tools to govern and within this is how it will be achieved.

[00:14:10] AD: And speaking of maturity, how are you seeing just like the knowledge gap, if that still exists? I mean, is it still pretty tough for these enterprises to find folks with the relevant skills? Is that getting better? What's that look like?

[00:14:19] RM: Yeah, it is getting better, but it is still tough. Because what is basically happening is that the core skills that we had in the enterprise, you cannot cut and paste the enterprise stack to the cloud. That doesn't work. You have an entire – A JSON skill set that you need to build. It's not like somebody – You don't need to go sort of fine – It's not like you went from driving a car to suddenly it's a chopper. So, you don't need to go buy a new set of skills for it. But it's adjacency. So, you have to train some of it. And some of it comes from experience. So, in that lies the resource crunch, so to say.

And so, people have to be deliberate in thinking that, "Look, going to the cloud, what is the purpose of me going to the cloud? What objectives I want to achieve?" And from there, figuring out what should be the technology architecture. So many times, I have seen this mistake, which is sometimes it becomes like a sort of IT project, as opposed to a business initiative.

The one guiding principle people should use it to go to the cloud is what is the business outcome that you want? Because that gives you the discipline. That gives you the constraint that you're actually looking for. Otherwise, it's most like I would rather you do it for the wrong reasons, and then end up sort of in this situation of saying, "I'm paying too much for it."

[00:15:53] AD: One other thing I thought that was interesting from the report, I'm not sure if there's anything here, but just the geographic variation in some of the responses. It wasn't surprising to me US was leading saying 65% of enterprises saying they're making progress in

the cloud. I think given the sort of VC environment and other things here, I think that sort of makes sense.

But the UK, they were the laggards at 46%, which is a fairly big gap when we're talking about those kind of numbers. Any thoughts on I guess why UK, and Spain is another one, are having a little more trouble there?

[00:16:22] RM: I wouldn't put it basically, even if you're sort of many of these companies, which you can see from these, they have a very close tie in with the US side of the thing. And some of the innovation that is happening there is actually perhaps getting to see hitting the ground in the US first and US more, both as a consuming market, as well as – Like you correctly said, the infrastructure being available because of so much funding having gone into it.

Both are more easily available perhaps in the US first. And then other geos. But it's a matter of time, and they will catch up, because it's a lot more democratized now, unlike in the past. So, it'll spread faster. But I would say that no matter where you start, the market that you will tend to serve first is in the cloud. At least it's more US. So perhaps that's where that data is wanting us to do.

[00:17:22] AD: I'd like to shift gears a little bit and just talk about, hey, you're an executive leader at Couchbase. You're providing a cloud-based service in Capella to your customers and get a feel of what that's like. Like, what's hard about running manage databases for customers? What did y'all learn in there?

[00:17:38] RM: I think – So, many of the cloud-oriented stuff, when you go, one of the things that typically happens is that you do a lot of stateless stuff, so to say. Again, sort of horizontally scale, which is very easy if it's a stateless systems. But the database is the thing that is all about state, because you would say, I'll find the state from the database.

Now, to scale this system and offer that – When you say scale, we'll end up always thinking about only one side of it, which is more. But in this case, it's both, that it could be that on Thanksgiving Day and Christmas, it's going to be 10x the volume. The rest of the days, it's like 1/10, right?

This system that we are sort of providing to our customers has to sort of handle it with this gracefulness of wanting to only – Let's say, just to make it physical, you would only one server on 363 days of the year. But for the other three days, you basically have 2.X days, you need 100 times that capacity, or 10 times that capacity. That elasticity to make it happen seamlessly underneath the covers is I think the learning for us. How quickly can you do it? How effortlessly can you do it? And how cost efficiently can you do it? These are the set of problems.

In that at the end of the day, the appearance to the consumer, the developer in our case, is serverless. As in like I didn't have to do anything. I just have the database endpoints for writing data, saving data, doing all my magic with that data. And the rest of the sort of the scaling, the upgrade, the patching, the maintenance, all of that literally is sort of taken care of underneath the covers, which means all the way from operating system upgrades to any software that you're using along with the database, all that stuff, the management of all of that stuff sort of comes into your purview. And you have to give the five nines availability of it and the geographic distribution of it. So that, like you said, the same thing as to work in the US. Maybe more consumption here. But the consumption in EMEA or Asia is not that far behind, but it is.

So, you got to sort of now manage those systems, their capacity differently from the capacity here. Those are all the learnings for us. And, currently, in our journey we are, I always say databases themselves, in our case, they will go in this direction as well, is they came from – Again, use like an automobile example, they came from manual, to automatic, to autonomous. That's the journey of automobiles, so to say. Similarly, where we are in the database journey as an industry is the same thing. Things used to be very manual with databases, with manual specialist skills required, to a place where we have reduced a lot of that by making things more automatic. And where we will go as we proceed in this path of making it more consumption basis, more autonomous, where the system can tune itself, self-heal itself. A lot of that is built in, but we are **[inaudible 00:21:12]** so that it becomes easy for our consumers to see these things happen effortlessly.

[00:21:18] AD: Absolutely. What's your relationship with the major cloud providers? Is that a pretty collaborative relationship? I mean, it's kind of friend and foe in some of these cases, I guess. But what's that like?

[00:21:28] RM: Coopetition is I guess the term that is used in all these places. Because cloud providers definitely want more vendors like us to be available. Only then, they are sort of generally can usable under multiple circumstances, so to say.

In that regard, that's very collaborative sort of relationship with them. But at the same time, as you know, all the cloud vendors also have their own database as a service. At that level, it is the competition part of it. But more choice is always good for the consumer. And we believe we have the strength to offer uniqueness or other competitive offerings. And so, so far, it's been a fantastic journey for us to work with the cloud providers. They have been great partners.

[00:22:15] AD: Yeah. Has it been difficult to build sort of that operational infrastructure on multiple different clouds? Or what's that challenge been like?

[00:22:21] RM: It's definitely a challenge, because multiple clouds have sort of evolved over the last 20 odd years. And that is sort of classic standardization of things over there. In that regard, having – We say we are cloud, or cloud vendor-agnostic. That's one of the advantages you have with a database like ours. So, you develop on this. You don't have to do it three times. We take care of this running on all the clouds. So, you are abstracted from it, so to say. So, in that, we have a classic set of problems to solve, which some are sort of generically solved. But there are a lot of vendor-specific ones that we end up solving because the API's are different, or the security model is different and stuff like that. So that is learning for us there. And, currently, the development, there are a lot more uniqueness for each other cloud that we need to solve. As they mature, I'm sure things will become that much more standardized, which will simplify this all.

[00:23:22] AD: Absolutely. One other question I have, I don't know how well-formed this is, but Couchbase is open source database, always been open source. You can run it yourself if you want to, which is great. Now that you're providing like a managed service, does sort of that open sourceness of it conflict with the managementnes, not like in a revenue way, but just like, “Hey, there are some optimizations,” even like an architectural way. I would say they're like some optimizations that maybe your team could do, but would make it harder to run just from an open

source perspective, right? If I'm running it for a single client, whereas you are running it for hundreds and thousands of clients, is that a problem you face? Or how do you think about that?

[00:23:57] RM: No, not at all. Because in one sense, if you go in that direction, what you're implicitly saying is that, “Look, then I want to run my own database instance, which means you need to have those skills and you have to – You're taking on the responsibility of the uptime, the management, the upgrades, all that other stuff that comes with it. So, you need specialist skills. And you're vested in sort of developing that, so to say. There is a class of sort of people.

But it's not an open source discussion. That is open source is merely to ensure that developers can pretty quickly put things together and they have a way of sort of trying before buying kind of situation, right? That's where that helps work.

Now, the delta is like this, which is that we give you the choice to run our database platform. In your data center. You have the option of able to run it in your data center, or in the cloud, or we will manage it for you. That's the management spectrum of it. How much hands-on or hands-off you want it to be? Because there are certain use cases where you want to be pretty hands on, because that data is so dear and I need to be controlling that. Many enterprises are still there. In fact, there are enterprises who would actually – Who have been asked to show an off-ramp from the cloud before going to the cloud. How would you achieve that? It's a very valid question, because there are government regulations, or there could be security issues based on which you may decide at some point that this data should not be out in the open. It needs to be totally in our control. Like, there are many industries of that nature, finance, healthcare and stuff like that, where that's a huge.

Why I'm saying that is, then, the really mature companies will have all these three models, because they are big, they’re multinational. And because it's got a geo component too. Which is certain geos, I don't want to run it in certain clouds, because for geopolitical reasons, I would rather learn it. It's going to be mixed.

But there are a lot of places where small to medium and other companies who they don't want this. For them, a managed service has a lot of not just – When I say small to medium, it's not about a company thing. When you want to try something, when you're building something, you

don't want only when it is mature and it's called a lifecycle. And only when you go through that, you realize that I need to have specialized skills. Until then – Or your dev test prod kind of stuff. The way we build it, you can offload. You can get to that point of investing in it only when it is really need. As opposed to right now, you're to start with that investment, which is where many people – To your original point, that many people are saying, “I don't want to deal with that. You deal with this.” That's where it begins. And there are some cases where they want to deal with it. That's also available to you. That's how we do this.

[00:26:54] AD: Yeah, absolutely. This is an exciting time, as you're mentioning, for databases and exciting time for Couchbase with having this managed service. I guess, like, what are you excited about over the next, I don’t know, couple of months or a couple of years? Whether it's features, or technologies, or anything like that? What sort of gets you excited?

[00:27:10] RM: Yeah. I think there's a big – And we have a long way to go. And there are a few trends, which are very interesting from this perspective of being in the data space is very exciting for us for what I see. First is this decentralization of the cloud, or more – Things became more edgier, if you will. And the definition of edge itself is sort of evolving, which is a network edge versus the real industrial edge, or domain-specific edge is also sort of evolving. And so, in both these places, it's about distributing the data. So, people have come to the conclusion, “I've understood that.”

If you have a centralized data, then the lack of that data being available to you at the moment of engagement, that's a little too much these days. Because there's a lot of data and there's a lot of compute to be done. Yet, you're carrying in your back pockets phones, which have amazing processing and storage capabilities. So, why not sort of use that? That's one side of the edge. That is, basically, even the cloud, it cannot be sitting in these – Wherever these data centers are. They need to be close to where the network's sort of edges.

And so, in that, there is a lot of innovation. And we have a distinct edge over there in terms of distributing the data. We have ability to do sort of with a click of a button in microseconds move that data to the new frontiers of edge of the cloud, as well as we have the mobile database and the sync capabilities to move into that. So, a lot of amazing IoT use cases that we see. And the excitement of what can happen there. All the way from processing the data edge at the edge,

finding insights and intelligence at the edge, and taking action at the edge without you having to come to the sort of do a centralized way of doing it.

Amazing amount of work with a distributed way of doing stuff. And this is one element, which the next generation of applications are also going to just call the WASM model and all that stuff. So, there's amazing stuff that's happening where we can participate. That's one side of the equation. Other is this AI bit.

What has really happened is that – Just use that same automobile example, is that when you go from manual to autonomous, it's actually the car that's driving itself. You're there just to supervise, right? That's the flip there. So, what that really means is that the data that is mainly generated on which you're doing earlier, it was to assist humans. Now, machines are producing the data. Not the humans. And so, that volume, and variety, and the speed at which this data is getting generated is something which some of our old databases were not built to handle. They were built to handle the human scale, not the machine scale.

[00:30:02] AD: Awesome. I love it. Yeah. Ravi Mayuram, CTO of Couchbase, thank you for joining us today.

[END]