

Episode 23
NoSQL with Rick Houlihan
Released August 10, 2021

DR 00:00:00

I'm Danielle Royston and this is Season 2 of Telco in 20.

Do you guys use Twitter? I think Twitter is the best. If you're not using it, you've gotta try it out. A lot of people think Twitter is like Facebook. That it's about following the people you know, and sharing boring updates about yourself, but actually it's the opposite. It's the best way to find people that you wish you knew, or even better, an awesome way to learn about topics from smart people you've never heard of that are putting out great ideas into the world for free. And if you search right, you can find Twitter GOLD. Something super insightful that makes you go: Hmmmmmm.

Anyway, the other day I was doing just that on Twitter and I found an awesome tweet by a guy named Rick Houlihan at AWS. He describes himself as the inventor of single-table design. In his tweet, he casually mentioned that one of the AWS account managers asked him to give a talk to one of their customers, contrasting noSQL and relational databases.

And Rick shared on Twitter, some immediate feedback he got from his talk. Quote: "Felt like I was back in school, watching a mindblowing lecture that literally shredded everything I thought I knew about NoSQL and DynamoDB." A mind blowing talk about NoSQL and DynamoDB. Sign me up. You know, we decided to use DynamoDB for the Totogi charger. It's one of Totogi's most controversial design decisions because everything in telco is based on everyone's favorite default, the relational database. I'm looking at you, Oracle. Most people think cloud-native equals throwing a legacy app into a Kubernetes container, pairing it with a relational database, and you're done.

When I saw Rick's tweet, I was like, *I need to get that guy on the Telco in 20 podcast*. So today, class is in session, and I'm going to be this student, and we're all gonna learn from Rick why you should be considering NoSQL databases for your hardcore telco applications.

So let's take 20.

Rick Houlihan is senior practice manager at AWS. Hi, Rick. Welcome to Telco in 20.

Rick Houlihan 00:02:18

Hey, how are you doing? Thanks for having me on the show.

DR 00:02:20

I'm psyched that you're here, you know, basically I sort of found you on Twitter and asked if you could be on the podcast and you said yes. And I think I jumped up and down. I was so pumped.

Rick Houlihan 00:02:31

Episode 23
NoSQL with Rick Houlihan
Released August 10, 2021

Thanks again. I really appreciate the reachout. It's a great opportunity to talk to a sector that I think is underrepresented.

DR 00:02:39

Yeah. So I call myself the leading evangelist for the public cloud in telco and sort of one of these few people that's super positive about the public cloud in telco. I think that telco is one of the last industries to move in a serious way to the public cloud. Recently, I launched a new business called Totogi. It's a charging system for telco companies and we built it on DynamoDB. And I think the whole industry thinks we're crazy. It's a real-time charger on the public cloud. I think that's a crazy decision number one. People think that's radical, and then going NoSQL and single-table design as sort of the second call that people think is nuts.

Rick Houlihan 00:03:24

I mean, obviously the single-table design, NoSQL. *Why would I do that? I have relational technologies.* I can only imagine you're kind of doubled down on that when you go into telco, right? Because you're so entrenched in the legacy IT operations. And the relational database is like the foundation of everything. So you're challenging the norms with this type of solution.

DR 00:03:43

Absolutely. And so the inventor of single-table design, let's just talk about NoSQL - and especially NoSQL related to relational databases. Everyone sort of defaults to relational databases as *'of course that's what we're going to pick.'* And I think you have a really interesting viewpoint on that.

Rick Houlihan 00:04:01

It really is about the evolution of technology. Over the years, we've had many technologies for storing structured data, right? The first database we had, we're all born with, right? The one between our ears, you know, when my eyes are opened, it's online, right? It's highly available and your charging system is not going to work if it's my brain power, right? But over the years, we've had many, many innovations such as data pressure, right? It starts to form. And we're seeing a lot of this today. Data pressure is the ability of the system to process the amount of data that we're asking it to process at a reasonable cost or reasonable time; and when one of those dimensions are broken, this is a technology trigger, right? Those are the types of things that drive innovation. We've had so much innovation over the years: paper tape, magnetic tape, distributed block storage, random access file systems.

As we invent new technologies, we tend to push the boundaries of those technologies with new applications, right? And telco, interestingly enough, is one of those sectors that drives this data pressure. And over the last several decades, you know, some of the biggest data processing engines have been built to support telco workloads, and specifically charging. And if you look at the pressure on the systems, it's tremendous. We've worked with a lot of different telco service providers at AWS, obviously. And we talked to a lot of these guys about cloud migration. The reality is that when it gets to looking at cloud elasticity, the ability to kind of extend into an area

Episode 23
NoSQL with Rick Houlihan
Released August 10, 2021

where I can turn things on and off when I need them and when I don't, right? This is where you can get the real value from the cloud. So even though we're kind of breaking the boundaries of the relational database with the workloads that we're driving today, and this is why we're starting to look at NoSQL databases, it's not just a technology choice that you have to make to satisfy these workloads of tomorrow. It's also a platform choice. It's.. *'Am I going to move to the cloud where I don't have to worry about owning that infrastructure?'*

DR 00:05:45

Yeah. When I talked to telco execs, I mean, they're obviously very nervous about it not being under their roof. They're nervous, especially in charging about latency, right? If it has to leave the network and go off to AWS, this far away place and come back, will it do it fast enough? And then being - *'Is it compliant?'* is something that's thrown in my face often. I think picking the right tool is really important. Being purposeful in your technology decisions, I think that's what we try to do at Totogi. Yes. And so I guess, talk to me about your thoughts on when is the right time to use NoSQL.

Rick Houlihan 00:06:24

Absolutely. It's only the right time to use new technologies when you know how to use them, right? And not before. And this is the biggest problem I see with NoSQL today in the market. And people run to NoSQL because they got a big data problem, right? Or they've heard it's the next new thing. And they try to use it at the same way they use the old technologies. And you look at the relational database. You have problems with those relational databases, and people do when they scale and the charging and telco, it's tremendous data load. Now we have data, voice, text, all the multimedia messaging, different charging rates, the choices need to be made across large, huge data sets. And that's where SQL servers break. Right? You start to shard that data across multiple instances, the instant you start doing that, then you should be talking to a new SQL database because that's what they do at their core.

Rick Houlihan 00:07:06

They're built to do this. And when it comes down to is *what type of app is it?* Is it an application that always asks the same questions? If it always asks the same questions, then that's kind of the definition of an OLTP application, right? Online Transaction Processing. What is charging? Charging is an OLTP process, right? The decision tree is applied to the information that comes in, your buy rate plans and codes, whatnot, how much you decide to bill for, here you go, here's the answer, right? But it's always the same question. And if that's the case, then we can always structure the data in a way that's going to be optimal for that particular pattern. That's where NoSQL databases differ from the relational database. The relational database has a normalized data model that is agnostic to every access pattern. And as a result, it's optimized for none of them.

Rick Houlihan 00:07:50

So you're automatically paying more overhead when you're using a relational database than you are when you use a NoSQL database. And the funny thing, is that all data is relational, right? So

Episode 23
NoSQL with Rick Houlihan
Released August 10, 2021

when we say relational database, NoSQL database, it's funny because the data itself doesn't change just because I'm using a different database. It's always going to be some set of relational data set that I'm going to be interested in. The collection of items that match certain entities are defined in my relational model. And we need to be able to express these relationships in a NoSQL database, the exact same way that developers do in a traditional relational database, because grouping the objects based on the access patterns is how we're going to efficiently work with our data. And when we get down into the modeling of the data, this is where the magic happens, so to speak. Everything becomes an index scan, and that's a much more efficient operation than searching across tables, joining data, and creating the ad hoc views. Again, a certain set of applications need that, but not OLTP.

DR 00:08:46

Yeah. I think analytics requerying your data set and asking different, interesting questions and trying to get insights might be better for a SQL relational.

Rick Houlihan 00:08:58

Yeah, absolutely. Depending again, on the nature of those analytics. If they're that class that I call, which is the most common class of Operational Analytics, right? This is something we actually found at Amazon that we were doing with our relational databases instead of running those summary queries against the operational data sets, which was way too expensive. We were running them on a regular interval and updating summary tables that would then be selected off of to get the metrics for the current period of time. So running them once, accessing them many times started leading us down the road of, *'Hey, you know, why don't we go down the NoSQL path'* because we're kind of creating a NoSQL database out of a relational database and that's not a good path.

DR 00:09:36

Right? And so I guess Amazon pretty famously made a decision in terms of, *"Hey, wait a second. Everything's on a relational database and maybe it shouldn't be,"* and you started to really move things to DynamoDB.

Rick Houlihan 00:09:49

That's right. And that was my team. Actually, I led that team. Made a lot of sense, convinced the business, and drove an initiative internally we called "Rolling Stone," which was migration of about 11,000 critical business services, including hundreds of what we call tier one services. These are the revenue makers, right? The ones that when these guys are down, we're losing money. One of the biggest reasons we chose to go to DynamoDB of all of the NoSQL platforms we had available was because of the mission critical nature of the service, right? It's what we call two year zero service. You don't stand up an AWS region without DynamoDB. If DynamoDB goes down, the region goes down. We've never lost a region. So.. DynamoDB does not go down. I wouldn't say that there hasn't been availability issues across individual tables. Of course there always are these issues and cases, but DynamoDB gives you a four nines availability

Episode 23
NoSQL with Rick Houlihan
Released August 10, 2021

guarantee out of the box. If you just turn a table on; if you add global tables, it's five nines. A lot of people talk to five nines, but do they deliver it? No, they don't deliver it.

DR 00:10:46

I can't tell you how many RFPs are on telco, like "You must deliver five nines," and everyone responds on those RFPs: "Yep, I deliver that."

Rick Houlihan 00:10:53

"Of course we do. Of course we do." Right? And how do you measure it? Right? DynamoDB measures it. And not only that, it's an SLA guarantee. I mean, it's there and on every single one of our DynamoDB tables, you can look at your availability metrics and it's all monitored for you. These are the types of things that I think when you get into mission critical services at tier zero operators like telcos, right? These are backbone service providers. They can't go down. Right. When they go down, real bad things happen. Okay. So I want to deploy that type of service on a framework that is built to be elastic and be able to respond to any demand. Right? I mean, I have examples of tables and DynamoDB that have spiked from millions of requests per second to 13-20 million requests per second, and done it instantaneously. What are you going to do with your relational database when I need twice the throughput? Right? You're not, you're going to go offline.

DR 00:11:45

This is what's really interesting, is that so many telcos are pre-purchasing expected, compute and storage needs, and database needs. I was CEO of a publicly traded company where we had a dimensioning team where we were guessing. And so it would bump it up by like 10 points or 20 points; we were in our case using Oracle, which is expensive, of course, being unused for most of the time.

Rick Houlihan 00:12:08

That's exactly right. The average enterprise application services running it, you know, 10 to 15% utilization. Albeit, the telcos might have a little bit more of a window of time that they're active. So we'll say that there were 30%, twice the industry average, what does that mean? That means 70% of your investment is sitting there on the floor of your data center doing nothing. And that doesn't happen with DynamoDB. With DynamoDB, it's just-in-time capacity provisioning every day, all day. So you turn it on when you need it, you turn it off when you don't. Heck, you don't even pay for storage in DynamoDB, until you write the item, I wouldn't have to pre-provision disk space. This is unheard of in the database world.

DR 00:12:47

I mean, this is the beauty of it, right? Where it has all of your scaling, it's elastic as heck up and down. You don't really have to manage it. And what's so beautiful, and not just for DynamoDB but all the services that AWS - you're paying by the API call or you're paying by the right. Paying for the usage.

Episode 23
NoSQL with Rick Houlihan
Released August 10, 2021

Rick Houlihan 00:13:02

All the serverless stack. Yeah. And that's it. So when you get people who do like TCO analysis and they say, *well, yes, if I have this workload that runs it, you know, X capacity and it does this 7/24, 365 then, oh man, DynamoDB is going to be really expensive.* I said, okay, why don't you show me that workload? Because I haven't seen it. And I've seen, I don't even know. I would say 10,000 or more workloads on dynamo at this point. And I work with all of our largest customers and nobody has that workload. In fact, the guys that come closer to it would probably be like Samsung galaxy sync service, right? That is probably the closest to high utilization, 7/24, 365 of anybody I've ever seen. And the reality is they try to own it themselves too. But their scale was so huge that the investment of personnel that they had to make to run that cluster, we're talking 560 Cassandra nodes, right? I mean, that was 35 people, 7/24, 365, just to run the cluster, just to keep it online. Right. They said, this is way too expensive. They turned it over back to us and they were happy to pay the overhead because the computer TCO and their actual TCO were way out of whack. Right. So they actually tried to take the workload back from us and they, that didn't last long, three months later, they pushed it back.

DR 00:14:10

This is something I've been talking about and probably not something that you talk about too often, but I'm going to sort of pivot us a little bit into like an HR topic. And that is, you think about that use case with Samsung and all the people and money that they have in their organization, focused on keeping the cluster up, which is super important - gotta to do it for your customers. But is it really like their core value proposition to the market? Like their ability to manage a cluster? Why don't you give that to AWS who knows how to do it at scale, automated, best in class - and then that expense, and the time and energy of those people can now be focused on coming up with the next best idea for Samsung,

Rick Houlihan 00:14:50

For Samsung, that was Bixby, literally almost half. Those guys got repurposed and put onto the Bixby, which is their verbal assistant; and talking about core of their business. Yeah. That's exactly what you're talking about. You don't want to be in the business of running infrastructure. Even a telco provider, they don't want to be in the business of running infrastructure. That's not a core value for them. Okay. Core value is developing the IT infrastructure that powers their application services, all of the code that they write that makes that stuff work. That's your core value. You take all those infrastructure skills and you apply them to application service and API development. You want to talk about how you're going to leapfrog your competitors. That's how you do it. You repurpose that dead HR infrastructure that is really just focused on the undifferentiated, heavy lifting. That's all it is, right.

Rick Houlihan 00:15:34

Keep the servers up, patch that server, make those security patches on those routers, threat assessment. This is the other thing I'm finding now is I'm working with more and more CIOs and CTOs who are realizing that the risk that they haven't been acknowledging is the actual human infrastructure that runs their data centers. Right? Look what we just went through. Okay. How

Episode 23
NoSQL with Rick Houlihan
Released August 10, 2021

many data center operations were impacted by COVID? Is that a risk that you really want to maintain as a business? No. You don't want to maintain that. You don't want to have to manage that risk. And a lot of these CSOs and CTOs and CIOs, you know, it's all about risk mitigation, right? What are the risks of my business? How do I mitigate those risks? Let's execute those plans. As soon as that risk pings high up on the radar, guess what they mitigate. So I'm seeing a cloud migration acceleration in the Fortune 500. I think they're moving faster than we ever thought they would. And you're seeing it.

DR 00:16:23

I think COVID absolutely accelerated that in telco. I mean, it may have moved the ball forward 10 years, a decade. Right. I mean, I don't think they ever thought they couldn't visit their own data centers. They couldn't visit their stores. Right, right, right. They thought that they were digital and they weren't. And they were just like, holy cow, we're not even close. And so I'm having much more interesting conversations about the public cloud. It's been great. But yeah, you know, sort of this DynamoDB concept of people are like, *'Oh, it's a cloud database. What if it goes down?'* And people don't really talk about DynamoDB is a tier zero service. Charging is monetization and the heart of a telco, right. This thing -- if it's not up, they're not making money.

Rick Houlihan 00:17:06

Yeah. The blood's not running. Right. Yeah, exactly. The patient's dead.

DR 00:17:11

And they're using human infrastructure to make sure that this thing is up and I'm over here saying let's, let's use a tier zero service that promises five nines. Yeah.

Rick Houlihan 00:17:20

And that's exactly right. I mean, not just 'promises' but delivers. Right? Yeah. We didn't make that service level guarantee lightly. Dynamo went through its growing pains. I mean four or five years ago, I think it was like right after I joined the specialist team, we had a major service event in IAD. it brought down like half the internet. I remember I was on the plane, listening to somebody behind me talking about how they couldn't make their Amazon order and something was wrong. And I was like, *oh, that must be some internet problem or something.* And I land. And then I go to the customer site, talking about DynamoDB and they tell me, *well, didn't you hear?* I'm like, *'What?'* You know, I mean, not the way you want to find out, but honestly it was a come to Jesus event for all of the service team.

Rick Houlihan 00:17:59

Right. We cannot have these types of events that took down a major chunk of internet infrastructure. And it caused something that I've never seen in 30 years of engineering experience. I've never seen a team, or a leader of a team, stop the release train in its tracks and reinvent; basically write off all the technical debt in that stack. And that was Jim Sharp, the best engineering leader I've ever worked for in my entire life. A real powerful peer at AWS, and now in charge of the IAM infrastructure. But he basically rebuilt DynamoDB into the service it is

Episode 23
NoSQL with Rick Houlihan
Released August 10, 2021

today. And when he was done, we delivered features that were years in the waiting: encryption at rest, global tables, backup, restore, point in time restore - it just keeps on coming, right? Because he'd built a nice modular framework that's easy to maintain, releases in chunks. And this thing is just - today, it's a beast. We have hundreds of thousands of customers. We might be well over a million customers at this point. I have no idea.

DR 00:18:51

Well, the opportunity to stand on top of those shoulders, right? That kind of engineering and that kind of battle testing. Right? Battle testing of - you sort of eat your own dog food. Everyone loves that descriptor. Certainly Amazon's eating their own dog food with DynamoDB, but all the millions of customers and all the workloads, the benefit that the applications that are built on top of that get from that investment that I don't have to solve as a vendor...

Rick Houlihan 00:19:16

That's right. Nobody has operational tooling like DynamoDB. Nobody has anything like this. And even the other cloud providers, I talked to those guys, and they all know DynamoDB. It is the shining holy grail of cloud. (DR: It's the mic drop DB.) It is the mic drop of cloud, exactly! It really is. I mean, other cloud providers can compete in certain areas, but they cannot compete in the distributed backplane database service area.

DR 00:19:42

At Totogi, we're super psyched about DynamoDB. We think it's our secret sauce. (Rick: I agree.) We try to architect our applications using AWS principles. We want to take advantage of the advancements you guys do in your custom chips, or you get 40% price performance overnight. The services, the pricing. I mean, our ability to lower TCO of a charging system by 80-90%?

Rick Houlihan 00:20:10

Right. That's your core value. That is what your business is. And one of the things about NoSQL that people kind of, they perceive that there's some sort of vendor lock-in to NoSQL, but the reality of it is that with a distributed service like DynamoDB, it's fully elastic. What locks you in is your data. Nothing else. How much data do you have in the system? How much data do I need to get out of the system? And how fast can I get it out? That's what locks you in. Because the design patterns that I use, the data modeling that we talk to, I've actually met with your team. We've talked about your tables. You know, we went through the modeling and they'll know this as well. They can take that exact same data model and stick it in MongoDB, Cassandra, CosmosDB, Cloud Firestore..

Rick Houlihan 00:20:48

It doesn't matter because it's an indexed object collection, and they're all indexed object collections. They all fundamentally do the same thing. So how fast can I pull it out? And how fast can I write it in? Now, DynamoDB, I can always turn up the RCU, the read capacity to, *'What do you need? Do you want to get that terabyte of data off my system in 10 minutes? No*

Episode 23
NoSQL with Rick Houlihan
Released August 10, 2021

problem. I can do that for you, Mr. Customer. You're not locked in at all. The problem is going to be, where are you going to put it? How fast can you get it there?

DR 00:21:11

Yeah, no, that's epic. Well, Rick, it's been a pleasure having you on the podcast. I appreciate your time talking to us about this awesome database that we're basing everything on. And so thank you. Thank you for giving it to the world.

Rick Houlihan 00:21:26

I think you made a great choice, and absolutely, let me know if I can help you guys out. I think you're groundbreaking in your market area. And I think this is an awesome service that you're offering.

DR 00:21:37

Well I might take you with me to a couple of tier zero tier one telcos. Please do.

Rick Houlihan 00:21:44

Please do. I'll be there.

DR 00:21:46

Stick around because we're ending each podcast with a Telco in 20 takeaway. I have 20 seconds to tell you something you need to know:

A lot of folks in telco talk about five nines and claim they're meeting that standard. The thing is, most vendors are not REALLY hitting five nines. If their solutions are not fully distributed, redundant up and down the stack from human infrastructure to the actual infrastructure, they aren't.

And as you just heard Rick say, Amazon's DynamoDB does. And that's because it's a tier zero database. Everything at Amazon is built on it. And lucky for telco, Totogi's charging engine is built on it too. For us, it was a no-brainer decision. You see, charging systems have known access patterns and that's because they are OLTP applications.

When you have known access patterns, DynamoDB is database gold. It has five nines of availability, is lightning fast, redundant, and secure. It has built in scalability and elasticity. All managed by one of the best technology companies in the world: Amazon. And even better, you're charged only for your usage. Pay for what you need when you need it. And nothing more. That's the way the public cloud works.

Can your on-premise #fakecloud relational database system do all that? No, it cannot.

While you're processing that gem, I bet you're like, how do I jump on that Totogi train? We're ready to have you come onboard.

Episode 23
NoSQL with Rick Houlihan
Released August 10, 2021

Call or text me at +1 925 TelcoDR. Don't forget to hit that subscribe button, share our podcast with your colleagues, and if you liked what you heard, leave us a review. If you're a telco leader, making your move to the public cloud, I'd love to have you on the podcast. Let's connect on LinkedIn and on Twitter at @TelcoDR. And sign up for our awesome email newsletter on TelcoDR.com. Later, nerds!