



The Distributed Data Show Podcast Episode #140:

Your Hardest Apache Cassandra™ Questions with Erick Ramirez - TRANSCRIPT

Release Date: March 17, 2020

- Erick Ramirez: Someone asked whether they can use the Java driver and they were trying to do some Kotlin stuff and I was like, "What the hell is Kotlin?" I've never come across that stuff. I was like, "Hey, I haven't got that super power. That's beyond my ability"
- Intro: From DataStax. This is the Distributed Data Show Podcast.
- Jeff Carpenter: On this week's episode of the Distributed Data Show, this is Jeff Carpenter here with Erick Ramirez from down under. How's it going from tomorrow Erick?
- Erick Ramirez: Welcome to the future. I'm doing well. Thanks for having me.
- Jeff Carpenter: Now you've just joined the developer relations team here at DataStax, but there's kind of a long history with you and Cassandra. I know that you're a veteran of the community and some people may know you for a while. Some people may just be getting to know you because of some of the stuff you've been doing recently, but why don't you back us up and tell us how you got involved with Cassandra?
- Erick Ramirez: Well, it's actually a crazy story. About seven years ago, a recruiter sent me a job posting for a role that had Cassandra in the list of things that the company was doing. And it wasn't something that I was familiar with. So, I did a bit of Googling for an hour. At the time I was working as a designer architect for Oracle deployments, Oracle RAC with Veritas clustering and stuff.
- Jeff Carpenter: Oh, got you.
- Erick Ramirez: Yeah. So I come from that background. But after I found out the bit of Cassandra, after about an hour, I was just completely hooked. It was just... I was thinking, wow, this is something that I should have been doing a long time ago. And it's solving all the problems that-
- Jeff Carpenter: Those pesky recruiters, right? I mean, they're always breaking in chaos agents coming at you with new technologies that you should learn.



Erick Ramirez: Yay. So it was solving all the problems that I've been trying to overcome in terms of high availability and resilience, that sort of stuff. And so within that hour I started looking up courses, which back in 2013, DataStax delivered them in a classroom style and it cost them several thousand dollars compared to today where you can just go to academy.datastax.com and just do the courses for free, do it at your own pace and it's still the same instructor led training but you can do it from your own bedroom if you want it to. And so at the time, this was seven years ago. My younger brother Joel encouraged me to reach out to people at DataStax and tell them my story because at the time I wasn't working and he said, it's a fantastic story that you're going across the world to spend money that on a technology that at the time no one really in Australia was using. Hardly anyone has heard of Cassandra and so I did.

Erick Ramirez: So it was a crazy idea, but I did reach out to quite a number of people at DataStax via LinkedIn. And lo and behold, one of those who responded within a few hours, mind you, was co-founder Matt and I was really shocked and surprised that he organized me for to get registered for the course and I registered for free. So just from that time that I found out about Cassandra three days later, would you believe it? I was on a plane to attend the course in New York. So, when I got back from that course, there wasn't any Cassandra work going around in Melbourne at the time. So, just to keep my training fresh, I was helping out on a mailing list and then fast forward a few months after that Aaron Morton happened to be having a conversation with Jeremy Hanna. Jeremy was based in the UK at the time and I was talking to Aaron and saying that DataStax is opening an office in Australia and Aaron from TLP recommended me to DataStax and the rest is history.

Jeff Carpenter: Oh, this is such a great story too of how you just kind of put yourself out there on a whim more or less fell in love with the technology as you started. I mean I think it's a classic case of how you go about doing this, starting out to build your network sooner or later someone puts in a good word for you and the rest is history really. Like you said, one of the things that you worked in our support team here at DataStax for several years, right?

Erick Ramirez: Right.

Jeff Carpenter: And then, so some people probably know you from that if a lot of our listeners might've interacted with you because you were helping to work their support tickets. But then you made the switch recently to developer relations and there's some people that are probably like, hadn't known you before that are like, Who is this Eric Guy? He's all over the Cassandra user



list now and actually some other places as well that... Maybe what we should do is talk through a little bit like where are some of the places that people go to ask the Cassandra questions. And if we have a chance to sort of give some shout outs to some key folks along the way that come to mind, go ahead and feel free to do that, but where do people go to get their answers?

Erick Ramirez: Yeah, so the main portal that people should go is community.datastax.com. You can ask questions and Cassandra experts from the community and from DataStax respond to your questions. Most of the time you'll get me responding to your posts and if there's anything that I don't know, I'll reach out to other experts, particularly inside DataStax to get the answers that you're looking for. But there are other avenues there as well. There's the ASF Slack, the Apache software foundation, Slack channel for Cassandra. So that's the official Slack channel for open source Cassandra.

Jeff Carpenter: Did that kind of replace the IRC chat that was the kind of the home for a long time?

Erick Ramirez: Right. I wasn't particularly active on the IRC channels at the time, but since the move to Slack, that's when I started getting into it. And then there's also the Cassandra user mailing list. So, that's fairly active as well. There's quite a lot of traffic that go through there. So I think more people know about the user mailing list just because it's been around for quite some time and I think it's a natural progression to join the mailing list just because when you do a Google search, the mailing list archives come up. So a lot of people know of its existence compared to the IRC and the Slack channel. Now obviously IRC has been retired now, so it's really just the Slack. And then there are other sites. They're like Stack Overflow. I've dabbled a bit into Stack Overflow but at the moment, in terms of my workflow, I haven't been able to dedicate time to Stack Overflow mostly because we prioritize the traffic on community.datastax.com and of course I'm very active on the ASF Slack and the mailing list as well, so...

Jeff Carpenter: Right. And you are just one man.

Erick Ramirez: Right.

Jeff Carpenter: At least as far as I know. Yeah. And I think, Stack Overflow in particular, there's been a good community of folks on there. I know a couple of years ago Aaron plots passed Jonathan Ellis as the top question answerer and that's a title and a credential that he'll never fail to mention. I'm just kidding. Love you. Great job Aaron. So, I mean there's plenty of places, I mean, I guess you go to Quora. I haven't heard as much about like the DBA site of traffic there, but I think there's a few questions that pop up there, but yeah.



You're definitely concentrating on a few of the main sites there. And I guess, so if you want just to drill in on something real quick there. I know that you prioritize and recommend the community site at DataStax. Is there any particular reason that you would prioritize that technically?

Erick Ramirez: It's mostly because I want to give the best user experience. So, it doesn't make sense to be doing everything and doing it badly. So whereas if people go to community.datastax.com, as long as I'm awake, they're pretty much guaranteed that they'll get a response within a few minutes, like if you have to wait an hour or two. But there are other engineers as well that are available during my nighttime, but I'm pretty much on there 16 hours a day.

Jeff Carpenter: Gotcha. And I also think that the community site, because it has that article format like you see on some of the other places like a Stack Overflow or a Quora, it tends to be a little bit more searchable than something that's just kind of like a threaded, like an email list or like a Slack. So there's a trade off there of like, I need an answer now versus I'd like to search through and do some research. I think there's probably a good balance there. Use the right tool maybe for the right circumstance.

Erick Ramirez: Yeah, and even if, for example, I'm in transit, I would tend to respond to a question and say, "I'm not in a position to answer right now, but I've seen your question. I'm going to respond within the hour," or something to that effect anyway, just so users know that there's someone there.

Jeff Carpenter: Yeah. Okay. Good. Okay. Yes. I hadn't thought about it that way. There is, I guess folks, there is an SLA for the community.datastax.com.

Erick Ramirez: It's not a strict SLA. It's more the fact that I understand we don't like users and developers getting stuck on something. We want their experience with Cassandra to be as frictionless as possible.

Jeff Carpenter: Cool. Okay. So, as a database company and hopefully a data driven company, you might think that we run word analysis, word counts on our forums. But we probably should do that. But what I'm going to do instead is just ask Eric. So are you ready? I'm going to fire off some questions about questions at you. What are the common questions that you see popping up again and again on the various sites that you support?

Erick Ramirez: Oh gosh. It seems like it's, I think the biggest friction for new users of Cassandra is data modeling. So the most common thing that comes up is from an operations perspective, how do I get rid of tombstones? And in a lot of cases DBA's are looking for an operational or churning fix to what is a design problem.

Jeff Carpenter: Right.

Erick Ramirez: And that's a tough one. So they talk about should I do a forced compaction. This stuff that might alleviate the symptoms but it's not a cure. So there are techniques that we might do, some of the stuff that a lot of people... It's not really documented, but if you do say a scrub of the SSTables or upgrade of the SSTables, it has the side effects that it will evict the tombstones for you. That's not what they were designed to do but it just has that side effect.

Jeff Carpenter: Right. Good. So, I mean I see people have kind of dug themselves a hole maybe with a bad data modeling decision. They're starting to realize there is a problem when they see a tombstone warning like you cited and you're kind of throwing them the rope, like okay. Here's a... While you should consider going back and maybe looking at your data model, that's more of a long term thing in the short term. Here's the rope to kind of pull yourself out of the hole with go run a scrub and get on the right track.

Erick Ramirez: Right. Or when they end up because they've previously been doing forced compactions then they end up with large a size tiered compaction strategy SSTables. So, you need similar sized SSTables for those files to get compacted together but if you've got like one giant SS table, then it takes a really long time for that to get compacted again. So, at least in the modern Cassandra age, like from Cassandra 3.10 you now have tools that allow you to split particularly the new feature in nodetool compact. When I say new, it's been around for two years now where, when you do a compact, if you specify the split flag, it won't generate just one giant SS table. It will generate smaller SS like in terms of instead of just one file. So at least it avoids that problem.

Jeff Carpenter: Okay. Any other common questions that pop up regularly?

Erick Ramirez: This one keeps coming up and I noticed that it came up during your previous webinar as well, is how do I do joins? I don't know. It's like-

Jeff Carpenter: All the time.

Erick Ramirez: Yeah. Well you don't because as you know, each query has to have its own table. So there really is no reason to join. Joins are really expensive, particularly when you're doing things at scale. So, in a relational world, you don't think about this stuff because you typically don't have a scale problem. You don't need your queries to perform in single digit millisecond response types typically.



Jeff Carpenter: Yeah. So, where do you send people when they come at you with that join question? We kind of want to route them in the front door, I guess to the learning about things in the right way.

Erick Ramirez: Yeah. So unfortunately there's no quick solution for those because if it means that they didn't get the data model right first stop. So there are some ways around it. So, depending on why you need the join, you can probably use Apache Solr. So in DAC you've got DAC search that allows you to do joins between course, which Solr course are backed by Cassandra tables or otherwise you'll have to use... If you're doing like an analytics workload, then you'll have to do it in Apache Spark. But otherwise there isn't like an out of the box solution of apart from getting your data model right, meaning that you need to design your table correctly so that you can satisfy your query without doing a join.

Jeff Carpenter: All right. Well, let's switch over. I went to the most common questions. What about some of the hardest questions that you've been asked? What are the real stompers that you've encountered?

Erick Ramirez: Maybe you had to go learn to get the answer out. Look, I'll confess. So, we've got it on record now. Someone asked about whether they can use the Java driver and they were trying to do some Kotlin stuff and in particular they said that the suspend functions weren't working and I was like, "What the hell is Kotlin? I've never come across that stuff." I was like, "Hey, I haven't got that super power that's beyond my abilities." So I had to reach out internally in DataStax and had to figure it out.

Jeff Carpenter: Right. Hit up the drivers team and they sorted you out.

Erick Ramirez: Yeah. But in terms of hard question to respond to, not because it's technically difficult, but it's more the fact that it's hard to diagnose is when someone has an SSL encryption problem. So say by design, SSL stuff is difficult if it's not working because you know like it might be because they've created their certificates wrong or they've sourced their certificates from somewhere and it's not compliant or they haven't built their certificates correctly and in a question answer forum format, there's a lot of back and forth that's required to figure it out and possibly the question answer forum is not the best place for that. So, that in that respect I struggle with those kinds of questions and I'll reach out to the users and say, "We really need you to log a support ticket so that a support engineer can work through it with you."

Jeff Carpenter: Sometimes you really do need like a little bit more of that hands on. And I would say that that SSL certificate handling all that kind of area is not

uniquely a problem to Cassandra deployments. Like you said, it's just kind of hard period. Are there some things that make it harder when you're working with a distributed database?

Erick Ramirez: Yeah, so the thing that I struggle to answer the most, I guess is when I get asked, We're currently using Cassandra 2.1 or 2.2 and we're doing some testing on Cassandra 3.x whether it's 3.0.11 or whether it's 3.11 whatever. And they're comparing the performance footprint. And for me that's a really difficult question to answer. I might put it to you because you're not comparing apples for apples. So, for me it's more like, okay, so you deploy on Cassandra 3, you do all your testing, does it meet your requirements? And then if it doesn't, is it because you haven't used the right configuration in some parts over they override.

Jeff Carpenter: You copied over the same YAML that you had or something like that without kind of looking at new options.

Erick Ramirez: But comparing 2.1 or 2.2 isn't a fair comparison because you've refactored the storage engine, the performance of the JVM is going to be different. And I don't know that that's a fair comparison to do a side by side in terms of the results from your testing.

Jeff Carpenter: Right. And this makes me want to do a whole other episode on the right way to approach an upgrade and updating your configuration and retuning and retesting at scale and all those things. But maybe we will just save that for a future time. All right, well we're starting to run short on time. One of the things that I wanted to ask you kind of in conclusion, having had the journey that you had from exploring a technology, falling in love with it, taking training and then, joining a support organization to answering all the questions that you are now, because you've kind of seen the whole spectrum of everything. So what advice would you give to a new Cassandra user about how to get up to speed?

Erick Ramirez: The very first thing that I tell anyone and everyone is that ask questions. There's no point getting stuck on something. There's different places where you can ask questions. Oh, I'll say again about community.datastax.com because there's no point wasting your time trying to figure something out. It really doesn't make sense just go and ask. It doesn't cost you anything. The next thing is that I'd say that maybe spend a bit of time on looking at docks.datastax.com and understand some of the concepts. But if you're a type of person that you prefer, like a video type content, then we always send everyone to academy.datastax.com. Learn the basics. And it doesn't mean that you spend hours and hours, maybe even just the primary. So some people learn by trial and error. But if you just learn at least some of the basics, look at some of the basic concepts



and for those who are a little bit advanced, feel free to browse through the code because you'll learn a lot looking at the code on git hub too. So those are the things that I would say.

Jeff Carpenter: Oh, interesting.

Erick Ramirez: Yeah.

Jeff Carpenter: Okay. I love that last one. Browsing through code, the actual Cassandra source code.

Erick Ramirez: Right. Because you get to have a feel for... Sometimes there's internal comments in there and then you discover that, Hey, someone's made a note about a particular JIRA ticket, and then you look at that JIRA ticket and then you get to see the conversations and the motivations, why things were done a certain way, that sort of stuff.

Jeff Carpenter: Right. And I don't think that we're suggesting that that's like your point of entry where you start getting your questions answered there. I've seen this debate come up in different places too because I got some feedback on my Cassandra book where I have references in certain locations to the actual Java class that implements some functionality and suggest that you go and look at that. And some people have come at me and said like, "I don't care. I don't want to know what the implementation is. Why are you including that?" And then I've had other people come up to me and said, "That is so cool. Thank you for doing that because now I went in and I looked at the code and now you know, I'm getting myself smart on distributed systems by going through it and looking at some of this stuff." So I think according to your personality, your results may vary there.

Erick Ramirez: Right. Everyone's different, everyone learns a different way. Someone wants to do like a tutorial that might be like a three day tutorial or something like that. Someone else wants to do it. Just jump in, make a few mistakes and then look up how to get around those mistakes. Some people like reading pages and pages of documents or books. We're all different. We all learn differently.

Jeff Carpenter: Awesome. Okay. Well I think that's pretty good run through for today. I want to thank you, Eric for joining us on the show and I want to say thank you to you and all the other people who we didn't have time to name today that are just kind of grinding it out day to day answering question after question, whether it be on Slack, on email list, on our forum's site, thanks to you all and you make the community a better place and help grow it day by day.



Erick Ramirez: Thank you.

Jeff Carpenter: And we'll see you next week on the distributed data show.

Outro: Thanks for listening to the Distributed Data Show. Please subscribe with your favorite podcast app, give us a rating and don't forget to share with your colleagues.