

Chris Casano:

Well, welcome back. Good to see you again.

Jack Chi:

Good to see you, Chris.

Chris Casano:

Yeah, how's it going at Iterable?

Jack Chi:

Yeah, it's going great. We have, I think, two plus years under our belt. And like I said, I'm an infra engineer, but my background is really as a backend developer. So, I think it's been trying to make our CRDB deployment just more automated, and takes less time to stand up a new deployment, and just all the intricacies of operating different clusters with limited resources.

So yeah, I think it's been challenging, but, yeah, I'm really excited to be here. I'm already learning a lot, yesterday and today, just code-pairing and all that. Yeah.

Chris Casano:

Yeah. Just yesterday, I know you met some of the other speakers, you were already talking about Kubernetes strategies, so it's great. We're already interacting before even this whole thing started, so it's awesome. All right, so why don't you tell us a little bit about Iterable, for folks that don't know who you are.

Jack Chi:

Yeah. So, when you hear the word Iterable, what do you think? I had to ask this to my founder, because I didn't know what it meant. But Andrew Bonnie, he told me it came from a Scala Iterator concept, you kind of keep iterating on something until you find something that really works. But what Iterable does is they make your favorite brands, and gives them an ability to connect to the right audience.

So, yeah, some of my favorite brands that are on Iterable that I've discovered recently is chess.com. REI is the one I use, and also Imperfect Foods is another one. So yeah, it's a lot of these, it's across-

Chris Casano:

All these, right?

Jack Chi:

... all kinds of verticals. But I really relate to the chess.com part. Just recently became a chess player, so every time I look myself up in the Iterable database, "Oh wow, they know a lot about me now." Probably shouldn't be playing chess during all hands, but yeah.

Chris Casano:

Awesome. And then you guys have grown a lot too since the last time you've been here. Look at this, seeing where you guys are at now is pretty impressive. So, we're going to talk about this too, how things

have scaled and so forth, and how the architecture's helped you along. So, yeah, I don't know. Anything you want to say about the growth, and about how much fun you've been having?

Jack Chi:

I think it's really specifically on the user profile cluster, it's grown quite a lot. The last time I spoke at RoachFest, we were just coming out of being in production. And since then we've migrated some of our bigger customers on board, like Grammarly, NBC Universal. So, yeah, it's really getting people to adopt. But this all happens under the hood. We've made it so that newer projects are already creating user profiles on CRDB, but for the legacy ones, migrating them off without any downtime, it should just work. It's really a collaboration of different teams, really.

Chris Casano:

Okay, awesome. Yeah, maybe let's talk about your application a little bit more. So, it was all around user identity. It's core to your system and what you do, right?

Jack Chi:

Yeah.

Chris Casano:

Give us an idea of the landscape, all the workloads, the behaviors, how users interact with your application.

Jack Chi:

Yeah. So, this is the overall, what I like to call kind of like the life of data, if you will. If think about the top, how data is generated. User can create that data through their interactions with an email, with an in-app, or it can come from third-party data providers.

So our customers use APIs, SDKs, and then they build that capability of sending us the data into the Iterable platform. And once it's stored in there, at a reasonable time, marketers can then really activate a lot of the data and to drive value, and to figure out how to reach their target audience. It could be your simple abandoned cart journey, or it could even be something more real time.

Like we recently launched this feature called Iterable Embedded Messages, and it's kind of like another use case that we're on now. It's a small cluster, but we're just rolling it out slowly. It's to give marketers real time idea of, okay, does this user qualify for our latest promotion? It could be Father's Day coming up this Sunday or Saturday, Sunday, and then it's like, oh-

Chris Casano:

Someone's got to take care of that.

Jack Chi:

So, that's the intelligence layer that we provide... Oh, somebody got on call. And the last red layer is really about the outreach. We started with email, we went into mobile app push, and recently we're really trying to build a global audience using WhatsApp. Because that's a huge channel to be going after. So that's the social audiences part.

[How Iterable Enables Personalized Experiences for Billions of Users](#)

CockroachDB | RoachFest 2024

And, yeah, user identity really sits at the core of everything. All the API calls, all the SDK calls, so it has to just work. And it was challenging when I had my first outage, first P1 incident, yeah, it was nerve-racking, honestly. But yeah, CRDB was there to help support it, it was nice.

Chris Casano:

Cool. Good. Glad we can be there for you. And yeah, just by looking at this with multi-channel, all these incoming data sources, you have users coming that's setting up all these campaigns or events, and so forth. So, it must be hard to think about what the workload is going to be on a day in, day out basis, right? Yeah, there's a lot of inputs, outputs, a lot of activity. I don't know, are things seasonal for you? Are there certain times of day where you have peak traffic and so forth?

Jack Chi:

Yeah, that's a good question. I think about there are patterns to our load. Of course there's seasonal patterns, like Black Friday is a big holiday for us. You can think about that. But other clusters, we have a journeys cluster, and that one tends to be very hourly. So, you'll see jobs that will start and finish on the hour, or start on the hour. So, some tweaks are just simple things like maybe rate limit some of the TTL jobs, or don't do backups on the hour, stagger them, do it at off-peak time.

So yeah, there's definitely some patterns to that as well.

Chris Casano:

Okay. And how do you plan for the scaling of these events? Or is things that you prepare for, or are things automated as far as when you need more capacity?

Jack Chi:

I think we have alerts and monitoring set up.

Chris Casano:

Okay.

Jack Chi:

There are best practices from Cockroach Labs saying, "Okay, you should monitor all the instance CPU memory, connection sizes, and all that." So, we monitor that. It doesn't auto-scale because we run it on EKS, and we actually run... It needs to touch us into the instance sizes, so we don't just auto-scale and spin up, spin down. So yeah, we keep an eye on it, and then I'm the one that's doing the increases. But there are tweaks you can make with memory, with the different EBS, IOPS and throughput. So, yeah, different clusters do behave differently. So, yeah, but you can also like, oh, maybe I over-provision, I can scale down to save cogs and all that.

Chris Casano:

Okay, awesome.

Jack Chi:

And all that... Okay, so yeah.

[How Iterable Enables Personalized Experiences for Billions of Users](#)

CockroachDB | RoachFest 2024

Chris Casano:

I'm sorry, did you want to mention anything else here?

Jack Chi:

Oh yeah, I was going to say on the Prometheus end, so we ingest that data into Datadogs, and it's a rich data set. And then if you need additional data sets to help you monitor and scale, it's pretty easy to add-in the additional metrics as well. Yeah.

Chris Casano:

Nice. Yeah, we're data dog users too, so a lot of our customers are, so cool. Well, that's fascinating stuff. So, we're going to get into a little bit more about maybe a little bit more of the guts of the application, and your architecture and so forth. So, this is essentially your user API, right? And maybe you want to just talk about how this interacts with how you acquire user identities, and then what happens from here?

Jack Chi:

Yeah. So, I think it's worth to note that we started out using Elasticsearch quite a lot. We're still a huge, heavy user of Elasticsearch, and that stores all of the users and the events that they generate. So, it's really like how do we offload some of that traffic away from Elasticsearch? And that's where the CRDB comes in.

So, on the JSON on the left there's two parts. The first one is just a user and an event. Just so you can visualize what kind of data you're dealing with. So yeah, chess.com, here's my user I found, and then here was a recent event, an email saying, "Oh, how can you explore the strengths and weaknesses of your pieces?" So yeah, this is a way to look up the user and join it to find relevant events for the marketers.

Chris Casano:

Cool. And then what have you found on the Cockroach side of things that helped you optimize for access, for lowering latency? Any cool concepts you want to-

Jack Chi:

Yeah, I think one of the coolest thing I discovered going... Because I've always been into PostgreS, MySQL, one of the thing going into CRDB is this idea of hot ranges, and how you shouldn't make sequential keys.

Chris Casano:

Yeah, we appreciate using UUIDs.

Jack Chi:

Yeah, UUIDs. I think there's a blog on your website saying that, if you can, go for the native keys. But then to support that you use something called the hash-sharded index to spread it out. So, in the UUID case, it works. If you're going to pick email, you want to spread it evenly, because I assume it's not evenly distributed.

So yeah, that was useful. That was different for me to learn. Where it didn't work was for our second use case, in the journeys cluster. We started out thinking, "Okay, that's the strategy we should go with, just use the hash-sharded index." But we actually found that led to slower queries, just because we want the data to be co-located. So, using more native keys, your simple hierarchy of things when you look it up. So, that was useful. So it's not... Yeah, it's like I have to think about the business logic, and then tailor the schema to fit. And I think as a DevOps, infra guy, that's a little challenging at times, because I'm not in the code and-

Chris Casano:

As a modeler. Yeah.

Jack Chi:

Yeah, so then you really have to understand the product as a whole. So, it helps to think holistically in that, to really understand how our users are executing.

Chris Casano:

And do you have to educate those application teams that come to create new [inaudible 00:12:29]-

Jack Chi:

I feel like they're educating me more than I'm edu-

Chris Casano:

Oh yeah?

Jack Chi:

Yeah.

Chris Casano:

Yeah. But modeling is a key thing. We talked about this a lot, actually, with even my own team. There's a lot of relational concepts that come over-

Jack Chi:

Exactly.

Chris Casano:

... into Cockroach, but it's a little bit different as far as how you think about keying strategies and so forth.

Jack Chi:

We haven't done a migration where it's an existing schema and we're migrating to CRDB, but if we do do that, I think there's a multi-tool to check the schema compatibility. So, I'm kind of excited to... If there's a need for that as well, to try it out.

Chris Casano:

Yeah, we have an afternoon session on it.

Jack Chi:

Oh, nice.

Chris Casano:

Yeah, make sure you stick around. Cool. Yeah. Anything else that we want to go through here, as far as user APIs, or data modeling, or other things that you learned along the way that you'd want other folks to know about?

Jack Chi:

Things I learned along the way? I think there's just, on the Slack channels and just the public community Slack channels, it's actually a good resource to connect with other people using the Helm Charts or using setting up in EKS. So, that was kind of nice, just to talk to people through Slack.

Chris Casano:

Yeah. We should go into that. Let's go into your architecture. Let's see if I can bring it up here.

Jack Chi:

Yeah.

Chris Casano:

All right.

Jack Chi:

It's a little small to see, but I think the takeaway here is we run just in one region, but it's across, split between three AZs. We run in AWS on Kubernetes, and we're using the operator pattern. And, yeah, I think it's a pretty vanilla setup, but what's nice about it is I only have to install all the supporting cast of characters, like the operator and our Datadog, DaemonSets and audit logging, all that I only have to install it once, and then I can just have add other clusters on top of it. So one operator is now managing four clusters.

Chris Casano:

Multiple clusters.

Jack Chi:

Yeah, four clusters now. And I'm recently making updates so that it can spread more evenly through the AZs, through use of Max SKUs in the pod anti-affinity rule set.

Chris Casano:

Okay.

Jack Chi:

Yeah. And then another upgrade is, it was actually really easy to get it working on the new Graviton instances. The challenge was actually first just getting it to join the EKS cluster. So, we tested out like, "Oh, we can get cost savings by running generic workloads on these Graviton instances." And then we were like, "Oh, there's a binary that works in CockroachDB." So that was nice.

Chris Casano:

So, you saw the cost per performance gain with Graviton then?

Jack Chi:

Yeah.

Chris Casano:

Awesome.

Jack Chi:

Yeah.

Chris Casano:

Cool. Good to know.

Jack Chi:

Yeah.

Chris Casano:

Yeah. So this is, yeah, it sounds like you have a good standard, I don't want to say vanilla, but Kubernetes setup using the operator. So, it looks like you're at the point where this is almost like a platform for you. It sounds like you can stamp out clusters if you needed to, as you [inaudible 00:15:38] use cases and so forth. Is that where you have evolved to?

Jack Chi:

Yeah. The first time when you do a rolling restart, it was kind of nerve wracking. I'm like, "Okay, things, it's under-replicated." But then you're staring at the errors like, "Oh, are the errors climbing? Oh no, it's fine." And I think the key is really working with the Cockroach Labs team to help us make a feature request.

I remember they added in the ability to tune the graceful pod termination period, so that it matches the life of our Hikari pool, which is the client side pooler.

Chris Casano:

[inaudible 00:16:21] pool, yeah.

Jack Chi:

[How Iterable Enables Personalized Experiences for Billions of Users](#)

CockroachDB | RoachFest 2024

And the load balancer just taking the active node out of service, and then it would just let all the connections die and then switch the future connections to the next pod. So, that was really nice. Although there was some misconfiguration on my end with IAM. So yeah, I've made just more... Okay, let's say if other services aren't running, what's the worst? How do I prevent that from failing? So, since I'm running all these on dedicated nodes, I just gave the permissions to the node itself, to whatever AWS permission it needs to update target groups and all that. So, that was nice.

Chris Casano:

Cool. Very cool. So, yeah, if you had to go back and do this all again, what would you do different?

Jack Chi:

If I had to go back? Oh, okay. So, I wish... Okay. The way now we deploy the cluster is using Argo CD. I think this afternoon people talk about that. I wish everything was already on Argo, because I can really version these clusters really easily. The way we were doing it before is I really need to, with every Terraform module, it wasn't easy to version it, but with Argo it's much easier to version a chart.

So yeah, if I were to go back and do it over again, yeah, I wish Argo was there already. Yeah.

Chris Casano:

All right, cool. So, let's talk about what other tooling you're using here. So it sounds like you're using Terraform, right?

Jack Chi:

Yeah.

Chris Casano:

Helm, Argo. You have connectivity to Sumo Logic. Any other tooling that you use here as far as... You mentioned Datadog for monitoring?

Jack Chi:

Yeah.

Chris Casano:

Any other key components you thought that were useful for setting up your environment?

Jack Chi:

I think the key thing to call out was in the enterprise cluster, we had this requirement to get data user ID at near real time to Snowflake. So, we found this feature called Changefeed, and it was relatively painless to set up CDC into essentially S3, and then there's a Snowpipe staging ground where they would receive an SQS notification, and then they would essentially do a read from the bucket, to get it into the right table in Snowflake.

So, as a DevOps engineer, that was pretty simple to set up. So, I think that feature is kind of cool. It's not on this diagram.



Chris Casano:

Yeah. It's one of our fun features. You don't have to set up anything ancillary for CDC, it just works out of the box, and point to where it needs to go and off you go. So that's great, I'm glad you're making use of changefeeds.

What else? While we're here, we're talking about your setup, we're talking about the operator, what would you like us to improve upon? What could Cockroach do better as far as making your setup, administration, resiliency, anything else that you'd want us to do?

Jack Chi:

Yeah, I think with the operator, some of the pain is, for me, it's a lot of manual, DB ops kind of work, once the cluster gets going. Things like configuring these cluster settings, setting up users, permissions and tables.

Chris Casano:

So, like bootstrapping.

Jack Chi:

Bootstrapping, yeah, bootstrapping. I recently have seen how Elasticsearch ECK gets deployed. For example, if I want to have human users create it, essentially I mount a secret, like a Kubernetes secret, with the password that's hashed. I wish there was a way for Cockroach to do that, is I would mount an external secret, I could sort it in Vault, or some other secret providers, and that gets audited and that gets created automatically. It's more declarative than... But we've made this other service to work around it, but I wish that was more out of the box, and that can play well with the other cloud native way of doing things.

Chris Casano:

Interesting. Okay.

Jack Chi:

That's one. The second one I think you guys closed the loop on is row-based SQL audits. It's like if you only want audit logs coming from certain users, I think that wasn't done before, but now that's available.

Chris Casano:

Oh, yeah. I know one customer really wanted that, so I'm glad it benefited you too.

Jack Chi:

Yeah, because some users can be quite chatty, right? If it's a service user, they want to... Yeah.

Chris Casano:

That was exactly their use case. They don't want all these service account things logged in the audit log.

Jack Chi:

Yeah. The way we got around it is we just wrote some custom Logic that we filter out certain users. But that's... Yeah.

Chris Casano:

Awesome.

Jack Chi:

Other things I wish we had..

Chris Casano:

It's a pretty good list.

Jack Chi:

Yeah.

Chris Casano:

Cool. Yeah. So, I guess what's next? What's next for Iterable?

Jack Chi:

Yeah, I think it's more about reducing the time to production with a lot of these clusters. Let's say that if another team at Iterable want to use CRDB, can I make the infrastructure more self-serviceable? That's one.

And personally for me, it's understanding some of the stuff from the COGS perspective, between AZ transfers. I don't quite know exactly how to go into the cost explorer, and really find that out, and plan for the cost. Is it in the schema that maybe that they're not optimizing for, that's leading to a lot of cross AZ chats? Is that why, or... Yeah, so that's really for me to figure out.

Yeah, but I have to say, I think we're doing a good job keeping up with the newer versions, just because doing a rolling restart to upgrade is kind of like, you can do it at business hours, really, rather than waiting for that off-peak. So, yeah.

Chris Casano:

That's great. So, upgrades are fairly painless then?

Jack Chi:

I would say so, upgrades are like, yeah.

Chris Casano:

Yeah, database upgrades are always fun in the past, so it's nice that they could just roll with them.

Jack Chi:

Yeah.

Chris Casano:

Awesome. Yeah. So anything else? Any other words of wisdom that you want to give some of the folks out there that might be even getting started with Cockroach, or even with user identity management? Any other key tidbits?

Jack Chi:

Well, I was really surprised to find I can just spin up that server list and connect to it. So, if you're in the application developer, spinning that up, just testing it out, and putting some load on it to see how it would perform is good. Yeah, I would say that's it.

Chris Casano:

Awesome. Well, I know we're between lunch, which is the best session of the day, so I think we're a little behind schedule too, so maybe we'll just [inaudible 00:24:10] a few minutes early. But Jack, really appreciate you coming to the stage again, talking about all the great things Iterable's doing.

Jack Chi:

Sure. Happy to do it.

Chris Casano:

So, thank you, really appreciate the time.

Jack Chi:

Thank you so much, Chris.

Chris Casano:

Thank you.

Jack Chi:

Thanks for having me.

Chris Casano:

Yeah.