

EPISODE 1485

[INTRODUCTION]

[00:00:00] SPEAKER: Between hyper-financialization and the growth of automation in the world, trade and investments have become an ever more prevalent and important tool for the mass market. But the relative knowledge of professional and hobbyist investors is different, and so is the access to opportunities. Coinrule empowers regular people to compete with professional traders by automating investments without having to learn a single line of code. Zdenek Hofler is the co-founder of Coinrule, and joins the show to discuss what he and his team have built.

[INTERVIEW]

[00:00:35] JM: Zdenek, welcome to the show.

[00:00:36] ZH: Hi, Jeff, it's my pleasure being here,

[00:00:39] JM: You work on Coinrule, which builds automated crypto trading bots, and I've dealt some with trading systems. Usually, the kinds of trading systems that I've dealt with, there's a human that executes complex orders and maybe those orders are combinations of options, and then securities, and you have a position that is a combination of a put and maybe a long position. But it's not entirely automated. You build systems that are entirely automated. Can you describe the prototypical trading bot? What it would do to make money?

[00:01:28] ZH: Sure. So first of all, Coinrule is not typical bot, I would say. We are pretty much a SaaS platform. As a user, what you can do you create a rule, right? It can be a simple rule, like every day buy Bitcoin, whatever. Or it can be even more complex rule, like, if any coin on the market, as price increase by let's say, 5%, and trading volume, let's say increase another 5%, and MA nine, meaning moving average nine, crossing moving average 50, in a timeframe of five minutes, then buy that coin, which is meeting those criteria.

So, this is what you can do as a user. Technically, what we do, we have an engine, and we call it the Arbiter, which is essentially running all these rules in a loop, so all the rules of all the users, they are still

being evaluated based on the complexity of the rule, and of course, based on the market data and market conditions on the market. And then, if the rule or if the condition is triggered, or set of conditions is triggered, then we pretty much place an order automatically, right? So, as a user, you need to connect your crypto exchange, because we don't touch your funds. For example, you need to have account on Binance, let's say, then you need to create API keys on Binance and connect those API keys in our interface to our system. That's pretty much what we do. We are evaluating market conditions, and then we are placing orders on your behalf using those API keys.

[00:03:19] JM: Now, with a combination of rules, I can essentially have a fully automated system that does all my trading for me, and I don't have to do much intervention. I wouldn't have to create new rules over time. Do you have any – well, I know you have some combinations of rules that can lead to successful and profitable trading strategies. Talk to me about some of the combinations of rules that can lead to a successful trading strategy.

[00:03:53] ZH: I think it's important to say that we currently only work with crypto, of course, and only with spot exchanges. So, you can trade spot trading pairs, and also a couple of futures. on Binance, futures, et cetera. That's pretty much the only thing what you can do. But it's not the only thing, right? It comes with a lot of opportunities. Coinrule, was designed as Lego boxes. So, you can build whatever you want to build on top of our engine, on top of our wizard, let's say, and there is really like thousands of combinations of what you can create. And of course, we provide, let's say, a list of strategies, or we call it templates or rule templates. So, you can choose whatever you want.

But of course, you kind of need to know what's going on. I mean, the aim of Coinrule is to become a leader in crypto and not only crypto trading automation because we are going to move to other coins have assets in the near future. But generally, you as a user, you have kind of full control over your rules, it's not a black box. You don't just put money inside when – just hoping that you will earn. You really need to know what you are doing. Of course, we are trying to teach you. We have the learning center. We are still trying to improve that, of course, to provide some video tutorials. We do regular catch ups with our users. We have a telegram community, where actually users can teach each other.

So, we are trying to build the community, and to help users to learn to be better traders. But at the end, what we do, we just provide the tool, and you are pretty much responsible for playing with that. Of course, what we try to achieve, of course, is kind of being a gamified platform. So, you actually learn by

playing, let's say. What does it mean, right? We provide a demo exchange, which is essentially a mirror of Binance exchange, in terms of market data. So, we are watching market data on Binance, and replicating market conditions into our demo exchange, and then you can pretty much learn and build your strategies on top of this demo exchange, and you don't risk anything. But generally, as I said, you need to build your own strategies on your own.

[00:06:43] JM: So, if I pick a strategy, like buying the dip, in a bull market, that alone is not sufficient to create a strategy, that's going to actually earn me profit? Like I need to, like, select multiple strategies? Because you have these strategies that people can select, and just have the bots execute them? How successful are those prebaked strategies?

[00:07:13] ZH: What I would say, like those strategies, they are based on back tested results from trading view. But still, it's based on history. So, there is an assumption that the strategy will work all the time. That's not true. The market is evolving, the market is moving a lot, and it always depends on the entrance. They are typically relatively often fake signals, right? So, it seems that the strategy, triggers the signal, and then, I mean, it actually triggers, but then instead of the price moving up, it's actually moving down for a while, and then maybe the trend is, correct. But depending on the exit side, when, what you said as logic or set of conditions for closing your position, then, of course, depending on that, you can earn or you can even lose, right? Of course.

The basic is you should always think about stop losses. But generally, no strategy is like covering 100% the success. Of course, we are trying to improve those strategies, as you go. We also have – we are running our own fund on Coinrule, using our own strategies, and we share pretty much those strategies as well with our community, and also our really like pro users, users that are with us, let's say for years, they also share their strategies. I mean, they know, it's not like, black and white. It's not like that. Sometimes it's really, really good. Sometimes, especially on the bear market like now, it's a bit worse. They typically do less trades. They are trying to protect the risk. If there is a risk, potential, strange behavior on the market, they just stop using some of their rules, they post them and then they run just some, let's say rebalancing rules for that moment for that while. And then when there is another trend, they can reenable them and that's it's.

As I said, it's not the black box, right? You need to kind of take a look at your rules. You need to follow them, let's say on even like daily basis, like most of our users, they are following the rules every day,

even like more than once a day, and then tweak them, they are playing with them, trying to understand how the system works. That's the funny part.

[00:09:56] JM: So, let's walk through the engineering behind a simple strategy. So, you mentioned I can just give my API key to have trades be executed on my behalf. Can you talk through what the Binance API enables, and just explain what the engineering stack is behind a typical bot?

[00:10:21] ZH: I mean, Binance API is coming with a lot of opportunities, of course. But we try to be, let's say, platform for masses, right? So, we don't over engineer everything, we just do simple things. But powerful things. So, pretty much what we focus on is just simple spot trading, limit orders, market orders, that's it. And then, of course, as I mentioned, as well, like, you can also trade futures. But again, it's very, very limited what you can do, in terms of how we use Binance API. Not only Binance, like all the exchanges we support, it's about like 10 exchanges, currently.

It's also because all those exchanges, they have different things. So, we kind of need to unify them in into one single interface. This is the simple thing. Then, of course, what's happening on our background, that's very interesting and it's very powerful. So pretty much, this arbiter, this engine is checking the market data conditions. We have couple of processes, couple of things, we run on multiple machines, and one of them is the collector. The collector what, it is doing is pretty much collecting data from all the exchanges, in real time, and storing them in cache, right? Currently, we don't actually store historical data. That's actually subject to change. But currently, we only keep the current snapshot of the market data.

In this case, in real time, we are getting updates, and all these arbiters and we are talking about dozens of machines or bots, Kubernetes bots, let's say. Those machines, we run all these rules. We have a queue, actually, multiple queues, and then we processed those rules in those queues in a loop. Let's say we just run the rule, then run it again, then run it again, et cetera, et cetera. Every time we run the rule, we check the market data, right? We check them from the cache. So, in that case, let's say there is a condition, if bitcoin's price increased by 5%. Let's say, in the timeframe of 30 minutes, so we know, okay, we use the 30 minutes candle, in that case. We take a look at the close price of the candle on the previous close candle, which is, let's say, x, and then we look at the current market price, which is changing in real time. And then we calculate the person to track and then okay, the condition is met, the condition is not met, depending on that we continue processing the rule.

Essentially, that's the simple thing what we are doing. The complexity in the system is actually based on the amount of conditions you can create. And also, the combination of conditions, actions and operators because you can combine, let's say, multiple conditions, then you can buy and then anytime, if something else happens, you can sell or you can wait another five minutes. And then, there is a lot of opportunities, a lot of combinations, and pretty much everything is happening on the engine. And then we just sent simple instructions if the condition is placed, met, and then the action should be triggered. We just say, okay, let's send this limit order. You want to buy 100% of that coin, let's say Bitcoin with your USDT wallet. Okay, this is your amount, this is your balance on the wallet. This is what we want to send. So, we send a single simple instruction and that's it.

[00:14:29] JM: So, can you help me understand the mapping between user roles and Kubernetes bots? If I create one rule, just talking more detail about how these rules are executed, and at what rate you would scale up? How many new rules would have to be created to scale up your Kubernetes bot count, for example?

[00:14:54] ZH: This is a great question. We actually migrated to Kubernetes quite recently, like three weeks ago. Two weeks ago, actually, before we were using different EC2 instances, but of course, the scalability was a bit more complicated. So, we moved to Kubernetes, and essentially, we are still playing with performance. We want to achieve our SLAs, we defined, which is pretty much the rule is being evaluated up to five seconds on free plan. On paid plans, it's even like three seconds or one second. So, as often as one second, let's say, we want to achieve that your rule will be evaluated, which is normally let's say, on a pro plan, which is our highest plan. Because on that pro plan, we pretty much dedicate a queue and a bot for that specific user. So, all the rules of that user are running on that single pot.

In terms of other users and other plans. It depends on the plan and the speed we want to achieve. But generally, I think, say that currently, we run up to six queues per plan, per region. Because of course, what we need to deal with is like multi regional setup, of course, because there are restrictions, gear location restrictions, for example, on Coinbase. So, if you want to trade USD, you need to actually trade from US servers. So, all these things, we need to take into account. Pretty much five to six queues per each region, per each plan. That's our current setup and we are still playing with that, because of course, we know we are still not reaching the speed we would like to. Pretty much, let's say this, every

100 rules, we need another bot, let's say, which is essentially quite a lot. A lot of bots at the end. But that's something we are happy to do, right? Because we really want to come up with those high-speed rule evaluations,

[00:17:09] JM: If something gets triggered, if the market reaches a certain point where an order would be triggered, there is some latency sensitivity. Because if you hit a certain trigger, and then you want to execute an order, and then by the time the order gets executed, the price has changed again, you can potentially have significant slippage. Are there methods that you employ to avoid slippage due to latency?

[00:17:40] ZH: Generally, I have to say, we are not ready for high performance trading, right? For that propose, there are probably better bots, especially those bots you run on your own machine close to the exchange. What we are doing, and what we need to deal with is of course, the mass amount of users and mass amount of rules. Because typically bots, if you buy a bot and you run it on your machine, it's your bot and you run it on your machine, and you are the only user of the bot, right? That's not our case. Typically, even I mean, it's not even our aim for now, I would say. Because the amount of possibilities and the amount of possible strategies you can create with Coinrule is enormous, especially the any coin option, we are not aware of any competitor that would provide such functionality, because it's really, really difficult to scan all the markets on that exchange. So, that's something that opens a lot of opportunities for us. That's why we do that. So generally, I would say if you need really like high performance trading, we are probably not the best solution for you. But if you really want to set up complex strategies, then I think Coinrule is really perfect for you. Did I answer your question?

[00:19:14] JM: Yeah, I think you basically did. I mean, the kind of trade that I'm describing is something that high frequency trading would be sensitive to.

[00:19:21] ZH: Generally, we don't have and we don't see any problems with latency in terms of – I mean, of course, when you are sending a market order. So, you have two options, as I said, market orders and limit orders. That's it. Market orders. Okay, at the moment, we evaluate the rule, the price is, I don't know, x, and then we place the order, and course, we are using market orders. So, on the exchange, the price is slightly different. But that's just fine. Our user base, they don't care much about this problem, because they can still earn and benefit it from the rule itself, from the strategy they can create. And of course, when you are using limit orders, the problem is still there, it's a little bit better, of

course, because you control the limit price. Of course, there is another risk, because we would limit orders, we use fill or kill option, if what I mean. Generally, meaning if there is no one at the other side, the order is canceled immediately.

So, there must be an ask or bid for your bid or ask, otherwise the order is canceled. But generally, users, they don't have problem with the that even with the lack or so. There are, of course, different problems that we need to work on. But I said nothing.

[00:20:47] JM: Did you mention you have your own fund that you're trading to test out the platform?

[00:20:54] ZH: Yeah.

[00:20:54] JM: Can you talk about what are the rules that you execute in your own fund?

[00:21:00] ZH: It's not a question for me, I would say, unfortunately, because I don't have the view. My co-founders would probably know a bit more and I'm not even sure how confidential it is. But what I think I can say, we use definitely rebalance strategies a lot, and they are very performant actually, even on the bear market, rebalancing strategies, they are really good. And then, of course we do some random trades, typically. Alec, It's actually managing that, my co-founder, Alec. So, I don't have that view. But I think that he's running some other rules. When there is an opportunity to earn, right? On some new coins, let's say, or maybe some favorite coins of him. But generally, those rebalancing strategies that are really, really good and performing well.

[00:21:57] JM: Can you describe what a rebalancing strategy is?

[00:22:02] ZH: So, generally, I mean, we have some templates about rebalancing strategies. But generally, in the first part of the rule, you take a look at the market, and you want to actually sell the best performance coin that you own right. And then, at the same time, pretty much you buy another coin, which is currently a bit down, but has an opportunity to be up. You do that on a regular basis. So, let's say once a day or two times a day, depending on your taste. But yeah, that's generally it. You sell something with profits, and then you buy something which can be profitable, potentially.

[00:22:50] JM: Got it. So essentially, you have a portfolio of different assets. And at certain moments, you're going to want to sell maybe, Solana increases a ton, and then you sell some of your Solana profits, in order to keep your portfolio at a certain percentage of Solana.

[00:23:10] ZH: Yes.

[00:23:13] JM: I mean, you're trading with real money, and I imagine the danger of a software bug can be quite severe. You don't want to end up like one of these horror stories of trading platforms that went haywire and lost traders a bunch of money. How do you test your trading systems?

[00:23:31] ZH: That's a really, really good question. I mean, we do a lot of smoke tests. That's the first thing. We run our own rules. Even on the fund, and we have also test fund, so we also run quite a lot of test rules, even random and not randomly. Regularly and every time there is a new functionality, we run our own rules on that functionality as well, to be sure that everything is correct. That's a small test. But then, of course, I mean we have unit tests and integration tests. It's not perfect, we know that, because we started as really small startup a while ago, and we were trying to move forward. And we did it. We've grown a lot. And now we, let's say, I think, we have a team of 12 people, 12 developers, and we are moving forward. We move to Kubernetes. We are trying to refactor a lot of things, actually. Because not everything is perfect, but it really works. And then, of course, with every new code, we do unit tests and integration tests. We really tried to be in that place, but it's correct and it's true that we didn't have those tests from the very beginning, right? Because we were just moving forward, and I think we were doing a proper startup journey, in terms of, we tested the actual product, manually, like a few years back, and then we were just, "Okay, this is making us profit, let's just put it out, right? So, it's incredible journey I can say. What achieved for couple of years, I think it's really great.

[00:25:22] JM: Do you have any horror stories of misconfigured platform you have ended up executing some damaging trades?

[00:25:31] ZH: I mean, generally, we don't typically have problems with damaging trades. I mean, the only thing what can happen, pretty much is that something is not triggered. I mean, I want to be honest, so it happened. It happened few times and we always analyze what's happening. We know pretty much, very well, because we have a lot of locks, like, really like terabytes of locks, from all the rules so we can analyze what's happening. Yeah, we used to have problems with the collector, let's say, the part

of the system collecting market data. And really, like months back, we really didn't have much alerts on that. Now we do. And with every incident, we try to be better, of course, and try to come up with some measures, preventing that particular issue to happen again.

But yeah, recently, I need to say that we had some problems with the migration to Kubernetes, especially because it's really, really difficult to test it on staging environment. But generally, what we do we communicate with users, we are very open. And if there is a problem, typically, our users get to know very fast. Because when we introduce something, within few minutes, they know. So, when we introduce something, and it's not working, as it should they know that and they say that on Telegram channel very quickly. So, we are able to respond very quickly, as well. If there is something we miss, sometimes it's happening, because of course, even with tests, you don't always cover every single scenario. That's the beauty and the problem of our rule wizard and the rule engine, that you have really thousands of possibilities. And we are typically testing rules, say like 5, 10, 15 rules, different scenarios. But we can't easily test every single scenario. Even like, let's say every week, from users, we get to know a scenario or strategy that we have never thought about. It's really amazing how different people think about our system, how they try to build strategies that we were not aware of. So, it's not definitely possible to cover everything. But yeah, we listened community. And we really tried to react promptly if there is anything.

[00:28:08] JM: The back end, what is actually deployed to Kubernetes, tell me about the programming languages and other systems that you're using, any web services or abstractions, you can describe.

[00:28:22] ZH: Sure. Let's say from the very beginning, we said, back in 2016. I said, actually, let's use NodeJS, right? So pretty much everything is written in Node. Why is that? I mean, generally, it's actually cheaper, right? We could use Python or Go Lang or so, but it's much cheaper from the company perspective and easier to find NodeJS developers. Pretty much everything is written in Node. And relative to the other stock we use, so we use MongoDB. That's our primary database. For a couple of services, we also use Postgres on RDS. MongoDB, currently, we have our own cluster. We don't have any third-party service, but actually we are thinking to use managed service for that. So, it's subject to change. We also use Redis, for caching, for streams, for market data updates, using the Pub Sub channels, and all the stuff and we use Redis quite a lot. It's really big part of the system. At the very beginning, we actually started all the system. We started with lambdas on AWS. So, even the engine,

API's, everything was on lambda. Then we, of course, very soon realize that it's slow for us and it has some limitations, a lot of them. So, we moved to easy to sync instances running on PM2.

And then the API was still running on lambdas, and cron jobs, and on all these SQS consumers, for example, because we still use SQS, as well, for a couple of things like historical things. But generally, we are slightly moving to Kubernetes fully. So, pretty much the staffers, Redis, MongoDB, and everything running on Kubernetes currently, even those jobs and couple of lambdas, historical lambdas. What else? What else? I think that's pretty much it. It's relatively simple. It was very complex a few months back. But with new power, which came a few months ago, we just decided to refactor these things, and simplify the architecture. Because before we were using three different regions, US, UK and Europe, and then managing that those clusters and replicas, and everything in three regions, it was very difficult, very complex.

We used to have problems with that, because for example, we were monitoring European replicas, with alerts, but not that those US, UK replicas, and there was an issue and incident, two months back, that actually US replica was disconnected from master and we were not getting data. And yeah, users started to complain and we were like, okay, but everything seems to work. And then, we realized it's not working, actually. So, there were a lot of issues because of the complicated in architecture, infrastructure. The aim was to simplify everything. So currently, we run from a single region, even if we need actually multiple regions. So, we use cross regional proxies to actually send those API requests to exchanges, from different IPs from different regions. Yeah, I think that's pretty much it.

[00:32:02] JM: You mentioned something interesting there about the geolocation of trade execution being relevant. When I think about platforms like Binance, or Coinbase, I imagine them as having API gateways that are all across the world. So, what exactly is the importance of geo locating your execution? Or can you talk through that more detail, the proxy that you're discussing?

[00:32:30] ZH: Yeah. So, this problem is probably related only to two exchanges we support currently. One is Coinbase. Typically, if you live in UK, and you want to trade pounds, British pounds, then you need to trade from UK servers. If you travel, you need to use VPN, because otherwise, you are able to connect to your account, you can trade, but you can trade pounds. You can trade, I don't know, using USDT, but you can't trade pounds. The same with Euro and US dollars.

So, there was a need, and it started with Coinbase back to I don't know, 2019 or so, that we needed a solution to be cross regional. Right. So, we were using database replicas, and cash replicas, and rule engine machines, in those three regions to be able to actually trade those fiat currencies. Then Binance US came, and you can trade on Binance US outside of the US, right? So, there was a second need, and we are from Europe, like all of us. So, generally, it was another platform or another exchange that we needed to solve. I mean, we already had a solution for Coinbase. So, it was pretty much simple. But maintaining the infrastructure, in those three regions, it was very difficult.

The latency, of course, was very problematic there. Because every time, for example, like maintaining a replica of Redis in the US, even using the VPC peering and these things, when that replica restarted, it required to load data, synchronized with master. The master was in Europe, and then it took even like, oh, I don't know why 30 minutes, 40 minutes to actually load those data. We are talking about a lot of gigabytes of data and then it was difficult. So, when that happened, that US engine was using data, which were, let's say, 30 minutes old, when that incident happened.

So, the need of using a proxy was really important. Actually, what I need to say, that need came a little bit later better, because there was another need before. Exchanges, they rate limit your requests, right? And what is very difficult in our work, or in our service is to maintain pool of IPs. You can't send everything from a single IP. They would burn us or they would slow us down. We need to back off, if they say, "Hey, you should back off, we need to back off." In that case, we can miss some trades.

So, there was a real need we saw in our logs, that it really happens sometimes that some of our engines on our servers, they really had problems with that. And we asked Binance, we are in kind of direct touch with Biden support and they told us the only possible way, because we can't increase those limits for you. The only possible way is to use an IP pool, right? So, you realized, okay, we need to use that IP pool, we need to use some kind of proxy. And then the other need came again, the other needs about the infrastructure. So, we said, okay, so we should actually run everything from Europe, as Europe is our primary data center in Germany. We need to run everything from here, and we need to use a proxy, we need to use some kind of proxy system to be able to first trade from different IPs from different regions, even if everything, the full engine is running in Europe, and the second need was, of course, we want to scale the number of IPs up to the maximum.

What was quite funny, is the Binance, there is a limit of 70 IPs. It's a hard limit, 70 API keys. Because as a Binance user, you need to whitelist those IPs. It's quite a new thing that came, I don't know, like six months ago. You will need to whitelist IPs, otherwise, your API keys will be working on it for 90 days, I think, and then you need to re enable them somehow. That's not user friendly, right? So, we needed to whitelist those IPs and there is a limit of 70 IP. But what was actually a bit funny is that there is another limit, they didn't tell us about that. And limit this actually 250, I think, characters on that field.

So, as a user, if you put all the IPs separated by space, it's a one field, which has 250 characters. It's not 70 IPs. It's actually like something like 36 IPs or so, depending on the length of those IPs, right? So, we were like, okay, so the limit is actually not 70. It's like 35 for us. Currently it's okay, so we are three regions, 35 IPs in each. It's relatively good. At some point, once we will have to scale because of the amount of users whatever, then we will probably have to come with different regions, or, of course, different set of IPs for some users, because it's like on a user or API key level. That's actually a funny story.

[00:38:28] JM: We're nearing the end of the time. But I just want to close off by zooming out and get a sense for where do you see the platform going. I think people could obviously build sophisticated, automated trading on top of this. I could also see people building things like Robo Advisors on top of your platform. But that's probably not what you're focused on. But maybe you could tell me how you see the platform evolving over time. And if there are any analogues to the past, I don't know other popular trading platforms you could compare it to.

[00:39:08] ZH: Yeah, so there are two things to mention. Technically, we really want to move forward quickly with features, et cetera. So, we do a lot of refactoring now and nowadays, which unfortunately, leads to a fact that we don't deliver new features as quickly as our users would wish. It's something that we kind of have these days, these weeks, let's say. But we are moving there, and the idea I mean, division of Coinrule is of course, when I start with simple things like introducing more and more indicators, like a lot of users are asking for, let's say, an increase. That's not what we currently have. But it's relatively simple to do. It's just that we, as part of that work, we want to refactor the bits behind, so it takes a bit more time.

These are simple things. The long vision is, of course, to move to other more traditional assets, because crypto is fine. But to be able to run our strategies or rules that are similar to crypto rules on

actually stocks, commodities, that would be massive. There is no such tool, or if so, it's very expensive. On Coinrule, you just pay simple monthly fee and that's it. It's not like very expensive. It's not something that is targeting to really professional traders or –

[00:40:49] JM: There's also the element here that I think Robin Hood has really opened up how many traders there are in the world. And so, there's lots of what you might call semi pro traders, and they want, they might want a tool that has some of the functionality of the really big expensive old platforms that enable automated trading, but they want a better user experience, more modern user experience. It seems like you're targeting that direction.

[00:41:17] ZH: Yeah, yeah, I would say so. But generally, I mean, Robin Hood is doing their job really well. But they don't offer what we can offer, right? We want to be more – even more user friendly, and to bring something that is not here. Really, like using any coin feature, or all the other opportunities and the variation of what you can create, that's really powerful and that's the way we want to go. Of course, we want to open it for even like decentralized exchanges. We already play with some prototypes, and then of course traditional assets, as I said. That's something that we really want to go there and I believe, we are on the right track.

[00:42:08] JM: Cool. Well, Zdenek, thank you for coming on the show. It's been a pleasure.

[00:42:11] ZH: Yeah, it was my pleasure. Thank you for inviting me, Jeff.

[END]