

Albert Chow:

This is IT Visionaries, your number one source for actionable insights and exclusive interviews with CIOs, CTOs, and CSOs, and many more. I'm your host, Albert Chow, a former CIO, former sales VP, and now podcast host.

Dr. Yu Xu:

We are not the first graph database in the market. We are the new generation graph database. So, all of our customers, they tried other things, right? If Oracle worked, if MongoDB worked, if other distributor SQL database worked, if early generation graph database worked, nobody would buy TigerGraph. So, they choose TigerGraph because we created something nobody had before in the world, right? We created the first fully distributed hyper-performance graphic database.

Albert Chow:

So, as companies rely more heavily on analytics to draw insights, how can they be sure that they're even getting the most of other data? How do they know the relationship inside of data? Dr. Yu Xu is the CEO of TigerGraph, a graph analytics software company that has created the industry's first distributed native graph database. Prior to founding the company, Dr. Xu gained valuable engineering experience at major operations like Twitter, Teradata, and IBM. Tune in to hear him explain how the development process behind TigerGraph led to breakthrough technology and why it's reshaping the way businesses approach, analyze, and act upon their data.

Albert Chow:

IT Visionaries is powered by Salesforce Platform and Dreamforce 2022. Catch the news and insights coming out of Dreamforce this year for free on Salesforce+. Just visit [salesforce.com/plus](https://salesforce.com/plus). Content will start rolling on September 20th.

Albert Chow:

Dr. Xu, welcome to the show.

Dr. Yu Xu:

Thank you. Hi, Albert.

Albert Chow:

Listen, we're pumped to have you. You are clearly an expert in something I don't think too many people know about, so I think we got to start there, which is, what in the world is a native parallel graph database?

Dr. Yu Xu:

Yeah, that's a good question. We invented native parallel graph database technology. Basically, before TigerGraph, nobody worked out a distributed graph database. We are the first company that created a fully distributed high-performance graph databases. Native means we build everything from the ground up, from scratch, to optimize everything for graph data storage, graph data compression, memory allocation, CPU competition for graph data, and graph analytics. Parallel is more than distributing. We can use many machines to do bigger data analytics, but also, even on a single machine, we can use many

CPUs, many CPU calls, to do parallel processing, to improve query performance, and do a lot of complex analytics. Otherwise, you would not be able to do it, if not doing parallel processing.

Albert Chow:

Okay

Dr. Yu Xu:

So, we are really excited. We are the first company that made the bigger breakthrough in distributing and the parallel graph processing.

Albert Chow:

So, giving an idea of what it's been able to unlock, because I think this is best told through an example, like, hey, in the past, give me a data set that maybe just couldn't get computed. It would take too many resources, maybe too much infrastructure, to figure out the answer. And then now, if you have a parallel graph database, it can handle that computation. Do you have any examples like that?

Dr. Yu Xu:

Yeah, I can give you two simple example. So, the first one is actually, probably one of the most famous, and it's the most profitable graph algorithms called the PageRank from Google. Before Google, there are many search engines. Google was not the first search engine, but Google used the graph algorithm called the PageRank to rank each website, each content. So, when you search on Google, they give you the best recommendation, but the power really come from the graph analytics. So, they have to look at the whole web pages, links to each other, how important each website is, right?

Albert Chow:

Yeah.

Dr. Yu Xu:

So, they have to do this computation each ration of the situation. Before Google had their own internal graph process system, nobody could do this type of work. So, that's why Google becomes the dominant winner in search engine, because they take advantage of graph engine. It's a building house.

Dr. Yu Xu:

The second example is about the money transfer. Think about just one table, person A transfers money to person B at this location, and transfer the decent amount for money. A simple table everybody can understand. Now, you want to see, does our union buy the guys, professional group doing money laundering, right? So, if we want to use traditional databases, like a relation database, you have to do what is called self-joins. You want to say, if a person A sends the money to B, B sends it to C, C sends it to D, and then maybe many other accounts also do similar transfers through a few steps. Then, a lot of people transfer money, finally, to one account or two accounts. So, that could be a sign of money laundering. So, with the traditional database, you cannot do this. You take weeks to finish your report.

Dr. Yu Xu:

Sometimes, you cannot even finish because you don't know the pattern. It's also related to machine learning, AI, because you don't know how people are doing the money transfers, through how many

steps. So, this is where a graph can be really the best tool. It's not just about the performance, it's also about if even you can write a graph for Amazon to solve a problem. With the original base, you cannot even write a SQL query to solve this money laundering problem.

Albert Chow:

Okay. I feel like I understand what you're saying. So, let's double-check my knowledge because I'm going to speak for people who maybe don't know what you're talking about. You can correct me if I'm wrong, all right? So, a traditional database is going to contain, let's say, single records of structured or unstructured data. You can query it. It can find a record, or a group of records, no problem.

Albert Chow:

But where it does not, and where it does break down, is it doesn't understand how the data is in relation to other data inside the same database, or across databases. And so, that's why you used your example of PageRank, which is actually trying to measure, not what a single row or record of a website is, but what all the other rows and websites say about this single record, which is how they created the PageRank. Similarly to your money laundering example, if you were to look at a database of transactions, it's very hard to tell the relationship between transactions, the record on line one and the record on line one billion. But parallel graph database is actually tracking and measuring the relationship between the actual rows of data to tell you these are the things that are interconnected and therefore, possibly, you have money laundering. Did I explain that correctly?

Dr. Yu Xu:

Actually, yes. You expanded it really well, and graph analytics is really about the relationship discovery, pattern match. You are right. Traditional database are really good in storing data and getting them back, right? But if you want to connect the data points, get insights? That's where graph database can do the best job for you.

Albert Chow:

All right, so I need to call my university, and they need to hit me up with a bachelor's. At minimum, bachelor's. I feel I understood a complex subject. When you think about what this does, you gave that prime example of, hey, it unlocks PageRank. Hey, it unlocked possibly, for financial institutions, to identify money laundering activities in a very, very quick way. Give us an idea of how much time this cuts off because I think you mentioned it earlier, but why is the computational power so much faster, and how much faster can it be? Because you made a comment about, if you wanted to query a database to understand the relationship between financial transactions, to see if a money laundering event was occurring, it sounds like a lot of time. And now, you're saying this can happen instantly. Give us an idea of the difference in how much faster this accelerates information processing.

Dr. Yu Xu:

Yeah. For example, one of our bigger customer is Jaguar Land Rover.

Albert Chow:

Okay.

Dr. Yu Xu:

So, in the supply chain business, they're making really good cars, yeah? So, before TigerGraph, it took them more than two weeks to do a key supply chain planning job. With TigerGraph, they cut down that time to about half an hour.

Albert Chow:

Wait, wait. Two weeks to 30 minutes?

Dr. Yu Xu:

Yes, yes.

Albert Chow:

Okay.

Dr. Yu Xu:

Yeah, and also, keep this in mind. This is when your traditional database can do the job, but there are times your traditional database cannot even do the job. For example, we have some of the biggest banks in the US as a bigger customer. Some of those bigger banks use TigerGraph graph analytics to detect fraud to find how many type of groups, customers, they have. This is really, really the two, machine learning and the AI, right? Before TigerGraph, they cannot even know how to do this job. So, it's not just by the performance. Of course, the performance can change a lot of things, but also it's about, TigerGraph can help you to do things previously you cannot even get it done.

Albert Chow:

So, for a company, is this an infrastructure and software problem? So, you mentioned before, hey, you might be on a specific database that doesn't even allow for this type of computation to occur. Therefore, you would need a product like TigerGraph. I'm assuming it's a cloud-hosted software. I have to mirror my database into TigerGraph. How does this work?

Dr. Yu Xu:

Yeah, that's a good question. So, let's know people all understand, one size does not fit it all, right?

Albert Chow:

Yeah.

Dr. Yu Xu:

Customers choose the right database mostly based on workload, based on the applications. So, TigerGraph, we're a full-blown graph database. So, it means if you want to get the performance and the advantages of a graph database, like TigerGraph, we need get your data in, and then we comprise the data. We store relationships, actually, as first class citizen. So, that's why, at the query time, we can be so much faster than traditional databases, and of course our customer can stream the data to TigerGraph. So, our database is always alive. It's always up to date, and we are in the cloud, but we also have a bigger customer prime. I'll give you an example. So, how customers use graph, right? You are right. Initially, customer use graph for some small-use cases, right?

Albert Chow:

Yeah.

Dr. Yu Xu:

For department use cases, it's very similar to the history of relational bases. When Oracle started initially, it was used by HR department for payroll system, for inventory control system, right? But over time, people realized the power of a database, and then relational database become a corporate-level infrastructure, right? So, we are seeing exactly the same things happening to graph database. When we come out of stealth mode four years ago, when our customer bought TigerGraph, they only bought TigerGraph because nobody else can do the job, right?

Albert Chow:

Yeah.

Dr. Yu Xu:

Not relational database, not a documented database, not a distributed SQL database. So, we started with one-use cases, two-use cases, like money laundering, like fraud detection, like a real-time recommendation engine for e-commerce, right? But over the last couple years, customers get a huge ROI, since they add more use cases, they add more data, then graph is becoming the central IT structure for a lot of companies. I'll give you an example. UnitedHealth Group is one of the largest healthcare companies in the US. They're actually a Fortune 10 company, right? So, they already built the largest healthcare graph in the world. They put more than 50 million patient data, a lot of doctor and hospital information, insurance name information. They build a single, unified patient graph. Then, on top of that, they're building many, many applications. So, really, using TigerGraph is the next generation database housing to power new applications, especially machine and AI type of applications.

Albert Chow:

So, what I've done is, as you're talking, I'm pulling up these stories and examples, and I'm looking at your work history. I'm hanging on one thing that you just said. You mentioned you just came out of stealth mode just four years ago, and already all these enterprises, they were already encountering massive problems, of which no database infrastructure system, calculation system, AI system, whatever. They couldn't solve their problem.

Dr. Yu Xu:

Exactly.

Albert Chow:

So, they gave you a shot just because they've tried everything else. Nothing else works. So, they said, "Hey, Doctor, give me your system. I want to try this out." That's how you got your first wave of customers just four years ago? Just like, hey, they were already stuck. They couldn't figure out the solution. You came forward and said this is mine.

Dr. Yu Xu:

Exactly. You can imagine, when we come out of stealth mode four years ago, really, nobody wanted to use TigerGraph. Nobody wanted to bet their career on a new startup, a new product. And we are not the first graph database in the market. We are the new generation graph database. So, all of our

customers, they tried other things. If Oracle worked, if MongoDB worked, if other distributor SQL database worked, if early generation graph database work, nobody would buy TigerGraph, right?

Albert Chow:

Right.

Dr. Yu Xu:

So, they choose TigerGraph because, like I said, we're so proud. We created something nobody had before in the world. We created the first fully distributed high-performance graph database. So, a lot of our customer, before TigerGraph come out of stealth mode, they already love the graph concept, right? It's visionary, so you understand, the graph really matches up. The business model can communicate with the business advantages, right?

Dr. Yu Xu:

But the problem before TigerGraph is that there was no enterprise-level, high-performance distributed native graph database. So, when we come out, our customer give us a shot. They give us a chance to prove the exact thing that you mentioned. Of course, it's not an easy process because bigger companies have a lot for process. You have to do rigorous proof of concept. For example, if you listen to the talk from UnitedHealth Group, they look at 30+ company-slash-product before they choose TigerGraph. They clearly rank 10 company product, TigerGraph was ranked the number one. It's not just number one. Number two cannot even do the job. If number two can handle the job, they're more famous, bigger than TigerGraph, [inaudible 00:15:34]. They would not just choose TigerGraph. So, that's how we started four years ago. Of course, right now, we focus not just on performance scalability because it already is the absolute leader. Now, we really focus more on easy to use for developers, for data scientists, and also business guys who don't want to be a graph algorithm expert.

Albert Chow:

So, we got to dive into that in just a moment, but I want to walk people through this timeline. So, if we go on LinkedIn, we can see that you started the company in, it looks like, 2012, 2011 timeframe.

Dr. Yu Xu:

10 years ago.

Albert Chow:

10 years ago. So then, 2018 is when you go live. So, for six years, you built this in stealth. I'm curious, you obviously have a lot of patents, and you have a lot of knowledge in this space. I'm curious, why do you think the VCs or whatever funded you? Why did they fund this project? Because they gave you a long runway, so they clearly knew that this is an opportunity, but you also mentioned many had failed before. So, what is it that you showed, maybe, way back in 2012 where they said, "Dr. Xu, I believe in you, man. I want to give you some money."

Dr. Yu Xu:

Yeah, I think, first of all, I'm a database guy. I have been working on database for more than 20 years. That's my passion. That's also probably the only thing I really know about, right? So, like you mentioned, I did a PhD in database. I've got many patents. I work at Teradata for many years in San Diego. Teradata,

of course, is needing enterprise reliable data housing company before Snowflake becomes a dominant play in the cloud. Also, I work at Twitter. Of course, I joined Twitter pre-IPO, and I left, quit Twitter before my first year anniversary. So, people invest a lot of confidence. I'm really committed to the success of TigerGraph. So, back to your question, yeah, it was a long journey, but building a database is super hard. Building a distributor graph database is even harder because nobody made this happen before.

Dr. Yu Xu:

So, I have the writer background. I got the writer team. So, that's why, for the first six years, we have to make the product work. We had some pitfalls. We had some issues. It's never easy, but we worked with early customers. We tested a product. It's better tested. So, take a lot of time to build a really good product, and also, like I mentioned, we didn't want to take any shortcut. We don't want to build a graph database on top, like a distributor storage engine, like a [inaudible 00:18:14] database like [inaudible 00:18:15] or MongoDB or other databases. Because we tried. We found, if you do the data nodding, if do the large-scale graph analytics, performance, it cannot even meet our own internal expectation, right?

Albert Chow:

Wow, okay.

Dr. Yu Xu:

So, we tried. So, now, just still startups. They're under time pressure now. They want to build a graph solution quickly on top of other distributor storage engine, but we tried many years ago. So, we finally decide we want to do the right approach, even if it takes time. So, we choose C++. It's not easy to hire C++ experts, but it's the best system program engine, program language. If you look at operating systems, major databases, they're all built using C++. So, we did this hard work. We make sure we polish our product. So, that's why I take a lot of time to build a good product.

Albert Chow:

So, the people that were involved in TigerGraph from the very beginning, the people that backed it, they had to know the opportunity in front of them. Because it sounds like Google clearly figured it out, but they also weren't providing that service to anybody. So, we knew that people wanted that service. Like you said, it was your building. So, they definitely had some type of conference and said, "Hey, Dr. Xu and his team, they know how to get this done."

Albert Chow:

You mentioned in the later answers that you are now really focused on making it easier, more accessible, and give us an idea of what that's been like. Because one of the big challenges I've always found with really extremely intelligent people such as yourself is, you have a curse of knowledge. It's very hard for you to understand what I don't know, and so you're trying to make this an easy system. People are probably like, "I don't even know how to use a parallel graph database, so how do I use it more easily?" Give us an idea what that's been like. What kind of team members are you needing to hire? Who needs to be on your team to make sure those elements, ease-of-use elements, come into play in the product?

Dr. Yu Xu:

Initially, even four years ago, everybody at TigerGraph is essentially an engineer. So, we have no sales market team. We don't even have a product manager because we're building a new product. So, we think from developers' point of view, because we're developers.

Albert Chow:

Yeah.

Dr. Yu Xu:

So, it started to work initially. So, we created our own query language, we created our own visualization tool kit. But at one point, we decided, we need to get a separate product management team. We need to get our customer advisors. So, I think initially, we only focused on performance scalability, because it's just the only differentiation we have. Again, we're not the first graph database company. So, we focus initially on performance scalability. Then, our customer also tell us, "Yeah, you're doing great in performance scalability." I'm asking more team members, more department, to build more application on top of TigerGraph.

Dr. Yu Xu:

Now, we have a problem. We need three more people using our language. We need to get more people to understand the graph. Can you lower the bar? Can you make sure more people in my company can build applications or get advantages of using your graph product? How it can do this more quickly, right? You can still do innovation in terms of performance scalability, but that's not a bottleneck anymore. You prove yourself already, right?

Albert Chow:

Yeah.

Dr. Yu Xu:

So, that's why, about earnings this year, we started a bigger change internally. We still do innovation in performance scalability. We're still doing a lot of patterns, but now the whole team can understand, it's really about making graph available for everybody. Not just for the best engineer teams in the largest companies. It's really about making graph available for more developers, more data scientists, and more business users.

Albert Chow:

So, how has that process gone? Because I'm sure you had to recruit new talent. Because you mentioned before, just four years ago, TigerGraph's all engineers, everyone understands how to code the language. Everyone's more technical, so did you just reassign people and say, "Hey, you got to figure this out," or did you go out and get some outside expertise, like, "Hey, these are some people that have a lot of skill and expertise in simplifying complex subjects."

Dr. Yu Xu:

It's both, but mostly, we had more external expertise. So, we built a new dedicated team called the developer relationship team. So, their goal is to sync the perspective of developers. So, tell the engineers, PM, we have a gap in ease of use. We have a gap in this feature. So, we also built a dedicated team called the customer success management. So, not technical people by training, but they



understand the use cases. They talk to customer all the time. They understand the customer requirement. Then, they translate the requirement to our engineering and the PM. So, we build different teams, so we can be more focused on customer success, be more focused on ease of use from a customer perspective.

Albert Chow:

And from your perspective, how do you know? This is one of the things I always find fascinating by extremely technical founders, and I love understanding more about how you're thinking. It's how do you know what success looks like in domains you're not familiar with? How do you go about measuring this? Because I've learned that it's very difficult to understand what you don't know. It just is. It's just very hard to know what is a good job when you don't know that domain. So, here you are, you're clearly a brilliant data scientist. You understand databases, 26 patents. We already hit all that. You got the company, it's working. You mentioned a POC of 30 companies. You were the only one that could even do the job, right? 29 companies couldn't even do the job. You, TigerGraph, could do the job. How do you measure the people and the teams that you don't know about? How do you know what success looks like?

Dr. Yu Xu:

You're right. That's one of my biggest challenges. So, I think it's really about learning every day. Look at the other success for founders, and also listen to customer feedback, and also get advisers, some advisers, who I really trust, who have done this in the past. So, it is always learning. Not easy. Yeah, I agree with you. It definitely is really one of the biggest challenges to a technical founder.

Albert Chow:

Now, one of the things that also is true about a company like yours is that it's very difficult to predict what people are going to do with the new capabilities that you've unlocked. What are some of the use cases that people have done where you're just even surprised by it. Well, I didn't even think that we could do that, but here they are solving problems that you could have never even predicted existed.

Dr. Yu Xu:

I have so many data points to share. Initially, when I started, my background is really in database. I have never worked in any other industry. Initially, I saw major customer use cases would be around social media, but it is all social media companies, like Twitter, Facebook, they all build their in-house solutions. They don't want to buy, and the other smaller social app companies are not doing really well. So, we build the best graph database, but our problem is to say where we can make the breakthrough in customer use cases, right?

Albert Chow:

Mm-hmm.

Dr. Yu Xu:

So, actually one of the surprising things to me is healthcare as an industry. Healthcare is really, all the industry, really conservative. They figure out a lot of problems already, but I mentioned the UnitedHealth Group use cases. So, one of the application in production is actually an application called the call center. You talk about this call center, it's not that exciting.

Dr. Yu Xu:

You would think people already figured out many years ago, but it turns out, they have so many applications where a patient or nurse make a phone call to the call center. They need to understand everything about this customer. You don't miss any appointment. What are your upcoming appointment? What treatment have you been through? Did you get a good quality in comparison to other people, patients, who are similar to you? So, before TigerGraph, the customer service representative had to open many applications, connect to different databases, get information. It takes them a lot of time. It's not a really good customer experience, and also, the customer lost money. Every minute they spend costs money, and also, every actual minute they spend with a patient, it's not a good experience because you have to ask stupid question. You don't sound like you're doing a wonderful job for your customer.

Albert Chow:

Yeah.

Dr. Yu Xu:

So, with the TigerGraph, to build a single unified healthcare graph, when anybody calls in to know everything about you from last year, this week, what medicine you took, family, doctor, history, anything. So, they give it the best customer experience. They can also recommend to you immediately, are you actually getting really good quality treatment, by looking at the patients who are similar to you? So, that's machine learning. That's AI. That's really, what is it called? The real-time care, plus recommendation. So, that use case was really amazing to me because it's a older industry, it's an older use case, but it has so much social impact because we all can relate to healthcare problems.

Dr. Yu Xu:

So, that's the first time I really think growth can really help our customer, to improve the world. So, that's just one use case. It's that many other use cases, like cancer patient care, senior member care. So, because healthcare system is huge, and also, the healthcare graph try to have the doctor to do more accurate diagnostics, more affordable treatment for the patient. So, this is just one use case. The graph can so fundamentally change the everyday life of people. So, it's happening right now. We also have use cases in power grade, in extra grade. They use TigerGraph to optimize the power supply, so that also has a huge social impact.

Albert Chow:

The audience members who listen to IT Visionaries a lot know that I care very much about healthcare. It was one of the key things I wanted to study and learn about when I went to grad school at Emory, and that was really disappointing to hear, when I got to school, just how siloed data really is. Every time you go to a doctor's visit, that doctor won't talk to another doctor. There's HIPAA compliance, all kinds of things that prevent a business, or a doctor, or whoever is on the other side trying to help you, from understanding what you've even gone through. Every doctor, every hospital, every office has its own system.

Albert Chow:

So, if you're able to bring this together, a great portion of improper care is actually just misdiagnosis. Not knowing that you already got a test. Not knowing that you already got something that might

disprove an outcome. So, I can see that right out the gate of how powerful a system could be. If you're telling me it could be fully distributed, it can say this is the relationship and build back a graph of what I've gone through in regards to healthcare. And you said, instantly, it can fundamentally transform how I receive care overall.

Dr. Yu Xu:

Yeah, I can imagine you really studied hard in your college because what you said is exactly what the expert from UT told us.

Albert Chow:

I would love to say that I studied hard, but yeah, I did not do that great in school, but here we are talking to people that clearly did. When you think about yourself, one of the things that's unique about TigerGraph specifically, and I'm really fascinated by it. This company's four years old, but you're one of the few people that we've had on the show that is in a CEO role, that it doesn't appear that you've... How much time have you spent leading people? Sometimes, people have to go through a CEO, president type role a couple times before they figure it out. It doesn't look like you had that. You were mostly on the engineering side. How's that learning curve been for you, leading people? Would you say it's a strength? Are you getting better? How would you rate yourself?

Dr. Yu Xu:

Yeah, but I have one advantages because TigerGraph is really a technology product-driven company, right?

Albert Chow:

Gotcha.

Dr. Yu Xu:

We don't need to invent some new business model. We didn't need to invent some new application because the database business is really mature business, right?

Albert Chow:

Yep.

Dr. Yu Xu:

If we talk about sales marketing, talk about a pricing model, we don't need a to do much innovation because we can learn from the best database company in the world. Even that, we are different type of databases, but the models are the same. So, I have advantages. If we really do a good job building the best database product, then everything else can follow. So, back to the question, I think being CEO is really tough, especially you come from engineer background. So, you know engineers. Typical engineers are shy. They're really nice. They don't want to have a conflict. It's harder to make a tough decision. So, it was not easy, but I learned over the time and also learned from other people. Yeah.

Albert Chow:

There you go. Well, it's pretty fascinating. This company, when we first got approached, this idea of having TigerGraph on the show, we started reading about you. This is a super fascinating story. The

quick ascension, of course. Any quick ascension story is always fundamentally very interesting, and the fact that it's engineer-driven makes it even more so. It just goes to show, I guess it's like what Jeff Bezos talks about, which is, focus on things that aren't going to change. So, people wanting faster, better performing databases is never going to change. No one's going to be happy. No one in 10 years is going to be happy with their performance unless it's faster, cheaper, something, right? So, you're entering a market. I like that concept of what you're doing of saying, "Hey, I'm going to bring my expertise to a market that I understand, and then let my engineering and product shine."

Albert Chow:

Dr. Xu, it's been really fun having you on the show, but before you go, it is time for the lightning round. The Lightning Round is brought to us by Salesforce Platform, the number one cloud platform for digital transformation of every experience. Dr. Xu, this is where we ask you questions outside of the realm of work, so our audience can get to know you a little bit better. Are you ready?

Dr. Yu Xu:

Yes.

Albert Chow:

Okay. When you were studying at UC San Diego, how much time did you spend at the beach?

Dr. Yu Xu:

I wish I spent more time. No, I was at the beach. Yeah, yeah. Especially now. I mean, it's San Francisco. Yeah.

Albert Chow:

Yeah. Cold water, man.

Dr. Yu Xu:

Yeah.

Albert Chow:

When you were studying, did you ever think that you were going to run your own company, or were you pretty happy, thinking, "Hey, I'm going to be a database engineer?"

Dr. Yu Xu:

Actually, my first ever job was in San Diego with a startup founded by one of my PhD advisors. That was the first time I'm working full-time. I worked with people from India, UK, Greece, Ukraine, Russia, and other countries. Flat structure, we worked really closely. We have a big vision to change the world. So, actually, from this time on, I decided and know I want to create on my own company someday.

Albert Chow:

That's awesome. So, from very beginning, you thought this was going to happen one day, just not clear on what the product was until parallel graphs became the thing.

Dr. Yu Xu:

Exactly.

Albert Chow:

What do you do for fun when you're not working?

Dr. Yu Xu:

I like to go hiking.

Albert Chow:

Yeah, where do you like to hike?

Dr. Yu Xu:

[inaudible 00:34:27] Mountains in the San Francisco Bay Area.

Albert Chow:

There you go. When you are hiking, what do you think? Do you ever think about work, or do you try to zone out and think about something else?

Dr. Yu Xu:

One beautiful hiking, actually help you to [inaudible 00:34:42]. So, yeah, when I'm hiking, I still think about a lot of things related to work and businesses.

Albert Chow:

What do you see for yourself in the next five years? What do you think is going to happen with TigerGraph or yourself outside of TigerGraph?

Dr. Yu Xu:

We want to make TigerGraph a bigger company, a big success for everybody. And I told the team, graph is the next bigger category. Actually, we are going to make a new international query language, called GQL. So, that is the only second database language to be standardized by ISO after SQL. So, we are on a bigger thing. So, we want to make sure graph database is going to become the first choice for applications, for developers. So, we have a big innovation. We have to work harder to deliver that.

Albert Chow:

So, I want to let you know that that is a big, hairy, audacious goal. I love it. So, for those listening, Dr. Xu is saying he wants to create his own language, GQL, graph query language. I want you to know that we've only had a handful of people on our show that have actually created a commercially viable language. The last person that was on our show was David Heinemeier Hansson. He created Ruby on Rails. So, will we, in the future, meet other people that have built their businesses on GQL?

Dr. Yu Xu:

Absolutely. Actually, I want to say, GQL is going to be out actually in about one and a half year. So, we are on the steering community, working with Oracle Core and other companies. So, it's going to be

This transcript was exported on Sep 21, 2022 - view latest version [here](#).

international quality standard. In some way, we are making the standard for the world. We're making a huge impact for the world. So, it's really exciting.

Albert Chow:

Well, I'm going to tell my son that this is what he needs to focus on. He needs to learn GQL. Forget all the other stuff. Just learn that.

Dr. Yu Xu:

That's really excellent early advice because there's a lot of opportunities, and actually GQL is really going to be very productive for a lot of data scientists, developers. Yeah.

Albert Chow:

Yeah, I can believe that. Congratulations on success. Anytime I hear about a product-driven company ascend to what you've done in just four years, it's always impressive. For anyone that's not sure about TigerGraph and wants to find out, you can check the show links below. It's also available at, if you go visit their website, TigerGraph. It's spelled exactly how it should be, tigergraph.com. TigerGraph. Dr. Xu, thanks for joining us today on IT Visionaries.

Dr. Yu Xu:

Thank you so much for having me. Thank you.

Albert Chow:

Thanks for listening to IT Visionaries. We want to hear from you. Let us know how you like the podcast by leaving us a rating and review on Spotify, Apple Podcasts, or wherever you listen. IT Visionaries is created by the team at mission.org and brought to you by the Salesforce Platform, the world's leading low code platform. If you love the thought leadership on this podcast, Salesforce has even more IT thoughts to chew on. Take your company to the next level with in-depth research and trends right in your inbox. Subscribe to a newsletter tailored to your role at [salesforce.com/newsletter](https://salesforce.com/newsletter).