Malcolm: So, tell us about Ontotext because I want to hear more about your company and what you do. And then from there, give us the elevator pitch of graph.
Doug: Sure. I'm Doug Kimball, the Chief Marketing Officer at Ontotext. The company's name may be less well known than what we deliver – GraphDB. We've been around for 20-plus years focusing on semantic knowledge graphs.
Graph technologies are a way to store and represent data in a more graphical way. We're so used to relational databases, where we have this nice little ordered table, but what if you have more than one data source? Most likely, anybody we're talking to has multiple data sources.
Graph databases are very good at representing the connections between different data sources and then understanding the relationships between them. So now you bring data together to have information, not just a bunch of different data points. That's one of the key things I look at from graphs.
Malcolm: So, I heard two things there. One is a way of managing, persisting, and storing data as, I assume, this collection of relationships. And the other thing is another way of displaying it or visualizing it, which is a little more node based or hierarchically based. Would you agree with what I just said?
Doug: You've got nodes that describe data and edges that describe the relationships between them. Some graph databases have visualization elements as a part of that and in

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

others, you put visualization tools in, so it depends. There are many different ways a graph database can pull that together. But if we go back to tabular data, you can't represent many-to-many data relationships very effectively. So, when you've got many-to-many, or even many-to-many-to-many relationships, graph databases are the only way to go.

Malcolm: So, when you say relationship, I assume that could be anything. It could be an interpersonal relationship like a familial relationship or anything like how a rivet relates to a pair of jeans or as a car handle relates to a car, right?

Doug: Right. But a rivet could also be part of a building information management system. Then you have to understand where that rivet with its particular measurements fits into that particular part of a building information management system in this particular case. And where else is that rivet used? What other parts go to that rivet? So, now you've got multiple connections.

But, you also have to have the ontologies, the knowledge on top of that, to understand that a rivet used in a building is very different from a rivet used in a pair of jeans.

Malcolm: So, what I've been told over the last few years is that the graph is not necessarily that new. But the use cases that it's very good at have really exploded, particularly around digital transformation. Would you agree?

Doug: Definitely. I think that the easy one is, and I'm sorry to bring it up, ChatGPT. Have you heard of it?

Malcolm: Okay. Good one. Everybody's talking about it. We should talk about it, too. We don't wanna be left behind.

Doug: The reason I bring up ChatGPT and large language models (LLM) is because they are helping graphs to get into mainstream adoption and for people to understand them. When I try to explain what I do I'd say, "Have you used Google, or DBpedia, or Uber, or any of these things that have graph technology embedded?" and people immediately start nodding their heads.

People are beginning to realize what graph technology is. I think things like ChatGPT, generative AI, LLM, and so on raise the understanding that this data stuff apparently is really important. Even for average people, though I hate using that term.

Malcolm: I think you're using that as an example to express the power of relationships and how words or phrases might relate to each other. When it comes to deploying, I assume that this is not an all or none, right? It's not like I need to go flush my entire X million-dollar investment in traditional database management systems. I assume that for a lot of your clients, this is a net new deployment for a net new use case, correct?

Doug: Those databases don't go away. Organizations have invested massive amounts of

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

money in these databases, and that's fine. But can all these databases be connected? Can you get the information out of all these different disparate systems?

Malcolm: Well, that makes sense. Particularly if you were trying to understand underlying causal relationships. This gets back to ChatGPT and LLMs, where you're trying to find patterns in data that wouldn't necessarily be there in a traditional database model. Is that correct to say?

Doug: Yes. Pattern analysis is so important for use cases like target discovery (or drug discovery), for example. If you have treated Doug with a medicine and he reacts in a certain way, what other patterns are going on? Has Doug recently had hay fever? Is he drinking too much coffee? So, now if there's an adverse reaction, maybe you can understand it better. Or if you're bringing a new drug to market or repurposing an existing drug, what are all the patterns of interaction? That kind of pattern analysis you cannot do in a relational database. But if you've got all the information about drugs in several relational databases and you connect them, that's different.

So, one of the things you were asking about also was – is this all or none? What we're seeing a lot of is that companies start off with a project around a knowledge graph. But the best way to do it from an enterprise standpoint is to think that this project will grow into a practice. So, it's not just a one-off, you're not just solving one problem. People realize more and more that if they can connect their knowledge graph to everything else, they can get a lot more power by having all this data brought together, not just as a one-off solution.

Malcolm: Getting back to these complex patterns and correlations, 20 years ago you could figure this stuff out, but you'd be running reasonably expensive regressions in SaaS. And now you can run the graphs and say, "Aha, there may be something here that I need to take a different look at." or "There's a pattern here I didn't know about."

I could see a ton of use cases in the digital world where you are trying to do things like understand buyer behavior or even website optimization – what are people clicking on, or are there patterns that people click on and then abandon? So, the use cases here are endless.

Doug: That's what I see. There's a quote I use from Scott Taylor's book, "Anybody who has data has data problems." And it's true. Anybody who is using more than one set of data sources to do anything to serve their end customer could benefit from using knowledge graphs.

Malcolm: We talked a little bit about how graphs can be used to create, visualize and understand novel relationships and manage and structure data in a different way that's more flexible and more powerful. I'm a CDO and I'm intrigued. I've been given a mandate for digital transformation. And my list of stuff to do is long… So, where should I start?

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

Doug: I would first recommend figuring out what the most prevalent use cases are that get you the most exposure. Because if you've got a long list and somebody's giving you a lot of money to do it, they'll want to see some early success stories. One of the easiest might be looking at something that has an AI or machine learning component to it and then figuring out what the first step to take in that journey is. There are zillions of places you could start. My conversations are usually around what is going to get you the biggest bang for the buck. Then ensure that whatever you're doing is also gonna be part of the same foundation you build for everything else. It's a somewhat generic answer, but, to me, we've got to get them going from "Where do I start?" to "What's a place that makes the most sense to start that gives me the ability to pivot and to change from there?"
Malcolm: Talking about building a foundation is a great dovetail into a recent episode of our podcast where I talk about data fabric. What Gartner calls metadata activation is, not exclusively, but in essence, running graphs against data to understand where these complex relationships exist. And, in the space of two months, I think we've gone from 7 to 10 years from mainstream to maybe 2 because of what we've seen with ChatGPT. This is just my perception. And I think your technology and other technologies like it will be, we could say maybe the beating heart of some of these solutions.
The story that I've been telling is that if OpenAI can use the entire internet as a training set, companies could certainly be using their entire data estate as a training set, right? I would argue that the future will be enabled, not entirely, but partially, by these large-scale graphs. Every company I'm talking to these days is having conversations around AI – what about you?
Doug: Remember, we're a company founded by engineers. These people just love to dive straight into it. Actually, before ChatGPT got the big buzz, we had people behind the scenes working on APIs to say, "Oh, how do we connect that into our systems? How do we do this?" But looking at it from a company standpoint, I think also to your point about data fabrics, if we can stitch them together it would be great. It takes the mindset of SharePoint and puts it on steroids and now you can find anything you need any time just by asking a question.
Malcolm: SharePoint plus Google plus added goodness.
Doug: Yeah. Yeah, I like it.
Malcolm: Right. It's all in the added goodness – a very technical term here on "CDO Matters". I think that's what the original thought behind intranets was and now it actually could happen. I wouldn't have thought a year ago that we'd be having this conversation right now. But I think that's where we are because innovation is happening at blazing speed.
Doug: That's absolutely right.
Malcolm: I've been watching a couple of interviews recently with Elon Musk… You'd think

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

that he'd be the last guy talking about AI regulation but he is openly out there talking about the need for AI regulation. And he said that we're fast approaching the point where we think we're manipulating the machines, but it could actually be otherwise.
Doug: Yeah. That was my point. What happens when AI starts to talk back about us?
Malcolm: This actually happened. I don't know if you recall this. It happened two to three years ago when they started talking back to us, but in a language that they've invented that is more effective than English. I wanna say it happened at Google, where somebody had written an AI protocol to speed up the productivity of something between these two systems. And what they ended up doing was creating an entirely new language that we couldn't even understand because they figured out that there were limitations in communicating in English.
Doug: That's interesting.
Malcolm: Well, we figured it out, but that's the kind of stuff we're talking about here, which has nothing to do with graph, but it's topical. So, we should probably end on that note. Doug Kimball, thank you so much for being here. It's been a pleasure.
Doug: It's been a lot of fun. I really appreciate it. Thank you.

Want to learn more about next-gen graph technology?

Book a Discovery Call

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com