

00:04

Speaker 1

Hello, everyone, and welcome again to another episode of the Human and Machine podcast. My name is Jaco Markwat. I'm here with my co-host, Lenny Smith. For those of you that have not listened to an episode before or first time joining us, thanks for listening, first of all. And secondly, we are based here in Johannesburg in South Africa. And of course, the Human and Machine podcast is a show where we, first of all, focus on the industrial manufacturing world, including mining. I always leave mining out of that one, specifically with the lens of industrial manufacturing here in South Africa.

00:38

Speaker 1

And what we aim to do is every week we have conversations with role players in the industry, people that we feel are quite influential, that our contributors, that are end users, and really anyone that has a story to tell about this industry that we love, that is industrial manufacturing, take a look at the tech that is available. And we also look at the impact of, for example, that Covid-19 has had on the industry. That's been pretty much the main theme for the last few episodes. I think it has been linear. It's been quite prolific and quite impactful. How are you?

01:17

Speaker 2

Cool, thanks, Jaco. Thanks for the introduction. I'm quite excited. I think this week we've kicked up our podcast with some local flavour. This week we've got our first international guest. So I'm very excited to. That's right.

01:30

Speaker 1

Over Zoom.

01:31

Speaker 2

Over Zoom. I'm very excited to have Jeff Nepo on the call. He's the executive director of business development at Canary. Jeff, how are you?

01:42

Speaker 3

Oh, I'm great, guys. How are you? And thanks for having me.

01:45

Speaker 1

It's an absolute pleasure. We're doing well. And Jeff, you're joining us from Pennsylvania. You based in Pennsylvania, right?

01:51

Speaker 3

That's right, yeah. Here on the east coast of the United States and soon to actually be based out of Texas, where I believe flow software has some offices.

02:01

Speaker 1

Yeah, yeah, that's right. So your move to Texas is imminent.

02:06

Speaker 3

I am 30 days away from being a Texan. I've bought my boots and my cowboy.

02:12

Speaker 1

That's fantastic. Yeah. Jeff, just thank you again for joining us. So, Jeff, you are, of course, with Canary Labs. For the folks that are listening that are not familiar with Canary Labs is. Has been around since 1985, I think, Jeff, it was actually quite a bit longer than I had thought. But for those of our listeners that are not familiar with Canary, it is, first of all, a time series database for industrial automation, but certainly a lot more. It's actually a

plant information management system because that's exactly what it is. It's a system. It's not just a time series database, but, yeah. Canary Labs. Just quickly framing the background of the company. Jeff Canary has been around since 1985.

02:57

Speaker 3

Yeah. What's so funny is our founder was off at university in the seventies and got introduced to a computer as part of a chem lab. And this was back when it was punch card programming, and that's how he fell in love with computing. So before there was a computer science degree, he was at university and just got bitten by the bug. And the next thing you know, he's doing internships with the Department of Defence and developing software that was based around nanosecond computing. And then in the eighties, this need for a database really came about for the first time in the industrial automation space.

03:37

Speaker 1

Yeah. Fantastic. And the heart's progress since then.

03:42

Speaker 3

Apparently. Data is important.

03:46

Speaker 1

That's what we hear every day. That's every day. So, yeah.

03:51

Speaker 2

And, Jeff, a little bit about yourself. How did you fall into this? I know sometimes I relate to you as you are. Time series data. So just how did you get into this industry?

04:02

Speaker 1

That's what you're known as in our office, Jeff. Mister time series data.

04:06

Speaker 3

Mister time series data. That's great. That's great. Okay. I've been known as worse, I suppose. So what's funny is I had absolutely no background in the industrial automation space whatsoever six years ago. My background has been more on the sales and marketing and business management. And I was lucky enough to find Canary, based on local connections, in the fact that they are headquartered in the town I was already living in. And I was actually at a party. I was speaking to some of the members of the Canary team, and they were telling me about what they do, and in particular, they were telling me about the market in which they operated.

04:54

Speaker 3

And what I was hearing from my side was that there was basically three big giant conglomerates that controlled the industrial automation space and that engineers and companies were getting some serious pain points by being forced to work in maybe a polite way, to say, would be an archaic business model. And it felt ripe for disruption. And so my sales brain and my marketing brain started really going full speed. I started a conversation with the founders of Canary, and we found out through that there could be a mutual beneficial relationship. And so I spent about six months just really learning as much as I could about the industrial automation space and then started trying to apply maybe what I had learned to be successful in other business ventures, in other markets into this space. And it really has been an incredible experience for me.

05:59

Speaker 3

Just learn so much. And just to constantly be a student of learning and. Yeah, and so that's my background. But the problem that we saw six years ago is still, I think, as relevant today as it was then.

06:17

Speaker 1

Yeah, yeah, certainly is. And I mean, so that was just under five years that you've been with him. And, I mean, 18,000 plus, I think, installations in more than 50 countries. I mean, it certainly has definitely had a very good, not only adoption sort of curve, but really good growth for you guys over the last few years.

06:39

Speaker 3

I think so. And the problem, essentially, the problem has resonated into every country, every industry, and engineers. They know one thing for sure. They know that having access to good and valid data is crucial for decision making. And ultimately, why make that decision? Why do you need access? Well, so that your decisions can be the best decision to save your company money. Whether that's in time, in efficiency, in avoiding safety issues, it's all about the bottom line. But the issue has been the solutions that were available to you to do this were either outrageously expensive, I mean, millions of dollars outrageously expensive, or the databases were highly unreliable.

07:33

Speaker 3

Or worst of all, it was some type of a homegrown solution that you had to build upon some open source database with potentially hundreds of man hours of working through connectors to make it work inside of your stack. Where's the Roi in all of that? It was a broken model. That's really why Canary has had such a high adoption rate so fast. We had done the work over the last 30 years to provide the reliability in the database, and not just in the database, but the client tools to add the context to the data. All we had to do was change our pricing model and make it easy to adopt the solution into existing stacks. So easy to get data in, easy to get data out fast, Roi generally less than six months. That's the goal when we find a potential partner.

08:33

Speaker 3

And that's why I think. I think that's what the industry has been yelling for. What's your all's opinion? What do you see in the south African market?

08:43

Speaker 2

No, for sure. And I think in the south African market, we also see, we've got a saying in South Africa. So that means, let's see, how much can I do myself, right? And how quickly I can I myself create something out of that?

08:58

Speaker 1

Let's see how much I can fix with duct tape.

09:00

Speaker 2

Duct tape, exactly. And I think there's a lot of, as you said, open source databases on the market. You can go and write, potentially your own connector. You can try and get these things into a database and into, to actually now store the data. But there's critically a few things that Canary does very well when we're just talking about storing data and just the vast amount of data that you guys can actually store in a second, as an example, into your data store.

09:31

Speaker 3

Yeah, it's hard to get your mind around it. We have this problem with big numbers. We can't visualize numbers greater than, generally 100,000, what that looks like, and that's just because our stadiums can be filled with 100,000 people. Right. So some larger systems can have more than 5 million historical tags. 5 million historical tags. And when you think about that loadout, you do not want the choke point of your PIM system. You don't want the bottleneck to be getting data into the archive. And so canary has been engineered for extremely high, inconsistent write speeds. We are able to do over 2 million writes per second in a 24/7 continuous operation. And so that will ensure, on a single server, mind you, so that ensures that the historian just won't be that choke point.

10:35

Speaker 2

And very important to note on that as well, is as we store these amount of data, the migration of the system performance doesn't actually worsen. So you get the exact same functionality and exact same write speeds. Doesn't matter how big your system grows at the end of the, where I think local or custom made solutions really would suffer to try and emulate that.

10:58

Speaker 3

Yeah. Yeah. We have a large team of developers that spend 40 hours a week trying to always make the product better. And just the same reason I don't handle my stock portfolio in my financial on my own is because I don't pretend to think that I'm going to spend more time than portfolio managers understanding the complexities of that system. And that's why we often would recommend, don't try to build this yourself. Just trying to handle reporting around daylight savings time will waste a year of your life.

11:34

Speaker 2

Jeff, I think when we talk about PMs and we go up to the second part of it's good, and it's all good and well, getting the data in and actually being able to connect to all of these different points on a manufacturing floor. But I think that something that steps apart between just being a normal time series historian is the ability to add context to your data. I think that's very crucial for any PIMs kind of implementation, is to have the ability to assign these different contexts to the real time data that you're actually storing into your historic yeah, Lenny.

12:07

Speaker 3

That's a great point. So part of my job is I get to travel a lot. I find myself, particularly inside of Fortune 500 companies, working with process engineers, and it's usually on more of a enterprise corporate data centre level. They are collecting data from all of their sites around the enterprise, bringing it back to a central location and from a very high level, analysing that data. The one thing that has struck me as odd, and it could just be, again, my outside view coming into this industry, is that here we have what I like to call automation superstars, men and women who have cut their teeth on creating automated workflows and processes, but when I'm watching them funnel through the data, it's still very manual. It's very manual.

13:02

Speaker 3

They are trying to find that needle in a haystack to find what could be going wrong. And one of the things that's become really important at Canary is to try to give automation superstars automated workflows. So let's just. I don't know, I always use boilers because it seems like it's a bit of a universal when talking to engineers. Right. If you've got 100 boilers within your enterprise, you don't want to check 100 boilers every three days to check their temperature settings. Instead, you want to be able to say, show me the boilers that are operating outside of parameter.

13:40

Speaker 3

And so being able to add that type of context to a large data archive where you're, instead of asking for everything, you're asking for the exceptions to the rule, really gives the end user the power to be able to take raw data and turn it into information that's actionable and that's what's important.

14:01

Speaker 2

No, definitely. And I think it's not just giving them access and being able to kind of add these parameters on top of it, but also you might need to be able to create calculations on top of this data to actually steer you into the right context as well.

14:17

Speaker 3

And you're exactly right, the calculations are key. And so I think it's a great segue to talking about, actually, Lenny, I think I'm going to steal this one from you. I believe, if I'm not mistaken, we co presented maybe almost two years ago now, and I watched you talk about the dashboard on your car, right?

14:39

Speaker 2
Yes, yes.

14:41

Speaker 3
Yeah.

14:42

Speaker 2
Still one of the best visualization techniques that's out there, and it's something everybody look at every day.

14:48

Speaker 1
And I guarantee you Lenny still has the same car as well.

14:54

Speaker 3
So, Lenny, as you, I think, very succinctly demonstrated, our dashboard gives us raw data it gives us the data that we're used to consuming on an operational level in the field. What's our current speed? How many ticks on the odometer? What's the level of fuel in the tank? But then, Lenny, am I right in saying the second piece you talk about is the condition based?

15:20

Speaker 2
Correct.

15:21

Speaker 3
What happens if a check engine light comes on? Right? So, right. We're defining rules around the raw data or around multiple pieces of raw data and saying, here's a calculation that's more situationally aware. And then what was your third?

15:36

Speaker 2
It's not like I can monitor each and every temperature sensor that the car is able to generate for me in each and every oil pressure for me to try and understand. Listen, is this actually something going wrong with my car? There's much better ways to visualize it. There's eventing systems that does it, like the event engine light on the car. And that's something that you can build into the PM solution to actually do and do that. The other point that I had, Jeff, is you only know where you're going if you're looking in the past. So, yes, you know what your current speed is, you know what your current odometer is. But how much have you actually travelled on this trip, and how long can you actually go before you have to fill up?

16:19

Speaker 2
So that aggregates a type of data on, based on your trip and how fast you're driving to actually determine the gallons that's. That's left. It's. It's crucial for you to actually make an informed decision on when you need to stop for gas. And it's exactly the same in our plant information solutions.

16:36

Speaker 3
Yeah, you're exactly right. And so I think typically. I think typically, if we would look at the stack, right. An engineer would have to go back to the PLC and program those calculations in. And not every organization touches their PLC's. That might be calling in your integrator, and that's got an expense to it. Canary thought. Let's give you an engine within the canary system where you can build out these calculations so that if you have a temperature that's giving you a real time update, but you also want to know what the running average is for the last 15 minutes. Well, you don't need touch the PLC to do that. Let's just build it off the tag that's already in Canary.

17:23

Speaker 3

If you needed to build condition monitoring in, let's just use the event service and tell you exactly when this has happened, where it happened, and during the duration of the event, what some other key parameters were around your different tag values. And so I think you're right, Lenny. That context build out has really been one of the secrets to our success, because you're not just purchasing a database, you're not just purchasing some. Some data collectors to log data to an archive. You're purchasing the tools you need to actually grab more information than just raw data out of the archive.

18:01

Speaker 2

And just on that raw data, it's also lossless raw data.

18:06

Speaker 3

Correct.

18:07

Speaker 1

We should explain what that means.

18:10

Speaker 3

Yeah, so that baffled me. Again. Outsiders perspective. I come to find out that most of the solutions don't actually keep the data values in the original format, that they would run some type of algorithm that would end up being a compression algorithm using, like, a swinging door, or worse. Worse would come back after three months and take all of the raw data and just roll it into a five minute time average, and some type of interpolation was getting applied. And so when I learned that was prevalent in the industry, we really started speaking about Canary's lossless data compression. So our algorithm never changes the raw data values. I don't care if you store data for ten years or ten minutes, it's the exact same data value as what it was the day you wrote it to the archive.

19:04

Speaker 3

We still pick up a three x compression ratio. So it's one of the smallest footprints in the industry, and it is by far the most advanced algorithm in the industry for storing raw data.

19:16

Speaker 2

That's insane. I mean, we're talking about, as you mentioned, data that change every second, and it's millions of tax potential, millions of values that we store, and you guys don't lose a heartbeat over it. So each and every one of those raw data points gets stored, and it's really something that's. That's truly remarkable about the technology.

19:36

Speaker 1

Yeah. Jeff, quickly, you mentioned something earlier, and maybe if, for the listeners that have been listening to the last few podcasts, there's been a few themes that have been sort of quite prevalent throughout the podcast we've done so far. There's definitely a theme around pricing. You spoke about sort of almost locked in sort of pricing that we've seen in the industry, in the market from some of these sort of large companies that have been around for seemingly forever. The second one is definitely the ROI, and probably more than just the ROI, but probably, how does that benefit the bottom line, which is ROI, obviously, but what is the business driver?

20:20

Speaker 1

And I think some of the discussions we've had over the last few podcasts is the importance of understanding how a piece of technology improves a process or enable people to eventually get me a better bottom line. That's definitely been a theme that's been coming through the conversation, and the one that I just want to spend a little bit of time on is theme of, or the notion of data, or information as a currency and the value of data. So what we've been observing in South Africa over the last little while is probably with the introduction of cheaper devices, cheaper networks, the ability to very quickly and easily now add multiple disparate kind of data sets

into your system has become definitely less of a barrier than what it has been over the last, I would say probably compared to five years ago.

21:11

Speaker 1

So when we're talking about these massive amounts of data, we sort of classify it as almost industrial IoT, I suppose. What are some of the technologies? I know there's a lot of talk about MQTT, sort of in layman's terms. For those listeners that are not familiar with that space of mass data collection, what are some of the observations and technologies available in that space?

21:38

Speaker 3

Yeah, that's a great question. I want to address the problem that makes this question so important, the problem you touched on, it's so easy now and affordable to go buy a new piece of hardware and plug it in somewhere and connect points to it. The issue comes in managing all of those pieces of hardware and in moving that data through your existing stack. So you've got this model where your instrumentation is one level, then your skate is on a level, and then the historian typically goes on the next level, and then you have MES or an ERP system, and then cloud and advanced analytics, machine learning, AI, and it just keeps stacking. And organizations have been having to manage all of these connections between all of these different components of the stack.

22:31

Speaker 3

And every time you change something somewhere, you have to go through and reconnect more systems. It becomes extremely expensive, not in hardware and licensing, but in manpower, and it really slows project development and launch down. And so that has been the problem. Now what we are seeing for the solution is a moving away from that traditional stack and moving more to an actual broker centric architecture. So whether it be an OPC UA server or a MQTT spark plug server, like what Cirrus link offers, the idea is that we get away from the poll response and we move to a sub pub model where we have publishing clients, we have subscribing clients, and everyone is interacting with the same, almost a hub spoke architecture where the broker is central. And that is really the push.

23:33

Speaker 3

It's the push from the vendors, because the vendors, we're tired of having to manage custom connectors to other pieces of software. That's one of the reasons we love flow and we love inductive automation so much, their ignition product, because we don't have to manage those connectors. We can just use MQTT.

23:54

Speaker 2

Yeah, in South Africa we always had a term about the spaghetti integration. That's us.

23:58

Speaker 1

Spaghetti integration.

23:59

Speaker 2

Spaghetti integration. So yeah, the hub and spoke model really does cater for that. And the nice thing about the hub and spoke model is you're 100% right. I don't have to go direct through all of these layers anymore. If an ERP needs a signal, or an MES solution needs a signal directly from either the historian level or the PLC level, doesn't matter. It doesn't have to go through all of these hoops. It can just go to the broker. And everybody that needs to consume that data point can consume it. Definitely. It's been an eye opener for me, and the amount of data and devices that can be handled by that is really impressive. So it's really been an eye opener for me.

24:36

Speaker 2

Coming from the traditional kind of automation pyramid, like we always used to know it, that hub and spoke model is really something exciting in the industry. And Jeff, also, could you guys from Canary are part of the MQTT symposium or drive, maybe you can explain a little bit more on that.

24:52

Speaker 3

Yeah, absolutely. So MQTT was co invented in the mid nineties by Arlan Nipper, and I believe Andy Stanley Clark. I might be off one of on his name, but either way, they invented it to solve a problem. It got deployed with Philips 66 on their pipelines and it really started to take off. And so now we find ourselves in a space where IIoT is driving the universal adoption of this standard. So Arlin and team did something very wise. They wrote a specification around the MQTT transport and said this specification will allow anyone in the IIoT space to understand how to read and how to package payload using the MQTT transfer protocol, so that we can immediately all start speaking the same dialect of MQTT. It's enabled, essentially plug and play IIoT projects just like now.

25:59

Speaker 3

You can connect a driver to your home, I'm sorry, connect a printer to your home network without having to go download the driver and install everything. The printer just kind of shows up. That's what Spark plug has allowed to happen. And then Arlen and team very wisely gave it away. They didn't want to own it. They did not want to force people to only play in this space, only play with their company. So you think about the difficulty with the big three on a vendor side is you have to start purchasing your hardware and your level two and level three software solutions from the same vendor because they hold your feet to the fire on their protocols.

26:45

Speaker 3

Well, by giving Sparkplug to the Eclipse Foundation Tahu project, it has essentially made it open source and has put the priority of maintaining the solution or the specification to the community that's going to use it. And so Canary immediately got on board, became a founding member of the Eclipse Tahoe Spark plug working group, along with inductive automation, along with Cirrus Link, along with Chevron, who's really pushing, from an end user's perspective, demanding their vendors support this protocol, as well as a few other vendors and end users as well. So we're very excited about this project, as you can tell from my ramblings.

27:35

Speaker 2

Yeah, that really is impressive, Geoff.

27:38

Speaker 1

Absolutely. Probably a massive requirement just for the future of everything, of the industry as a whole. It's crucial for this technology to be open, to be multi vendor, to be accessible to most and all. It's what we need right now to overcome this gap and this requirement from industry. So that's great. That's exciting. It's very exciting.

28:01

Speaker 3

Well, and to the credit of Emerson's and Alan Bradley and Wonderware that they are recognizing this and they are making the right moves to shift. They're making the right moves to shift. I equate it to. Canary's always been very lucky in that we have been the small boat on the lake with a two foot rudder, and we can pivot very quickly to address needs in the market. These large companies are huge cruise ships, right, a big cargo ships, and their rudders are very long, large, and their turn is much slower because of the size of the company. But they are hearing, they are responding. And I really think you're going to see these types of standards and protocols, these universal standards and protocols do what Modbus did for the industry so many years ago for the empowering of Iiota.

28:55

Speaker 1

And we probably see that in pretty much every kind of technology industry is the ability for things, systems and people to connect and speak to each other. Interoperability seems to be one of the magic source words when you talk about technology, regardless of the tech type, interoperability is quite an important one.

29:18

Speaker 3

Yeah, absolutely. And if you look at the success of some of the other platforms that are out there, that's why they've been successful. Why has ignition, by inductive automation, really started to catch the world on fire? Why? Why, in less than eight years, have they gone from new to becoming a standard? It's because they don't hold you down to any one thing. They've embraced MQTT for interoperability. Look at flow software. Between flow canary and ignition, we can share data in with less than 30 minutes of configuration because of MQTT.

29:56

Speaker 2

Yeah, that's true. Maybe steering the point a little bit back to Canary before we talk about the great ways that we can actually analyse and interpret the data. One thing for me from the Canary side, Jeff, is the way that we can make all of this data available to anybody in the enterprise or in your enterprise. I mean, we're talking potentially about your, I don't know, 5 million points. Not everybody is interested potentially in all of those points. Certain people only interested in certain assets. So there's a great way that we can actually add another piece of context to our data, which is the concept of an asset, and make that available to any user pretty much, that needs to see that data.

30:46

Speaker 3

Yeah. So let's talk about that. It's a great point, Lenny, to your 100,000 subscribers and listeners to this podcast right now. Raise your hand if working from home on two days notice. Right. Four months ago, I guess three months ago, if it created connectivity issues with getting access to your data, or at least pain points. And everywhere there's hands going in the air.

31:12

Speaker 1

Probably what's kept it teams exceptionally busy over the last few weeks and months.

31:17

Speaker 3

Yeah, no one in an IT department anywhere got laid off. You saw operation guys requesting to become IT guys. Yeah. So there's two parts to Canary that really have made those that had already adopted it king of the hill, if you will, during COVID and for future listeners, we're speaking right now, still in the height of it. The first part is what, Lenny, you just alluded to, and that's our asset modelling feature. So the archive, the historian that we write data to, it's always the same. The points come in. They're named just like they are generally inside of the SCADA system or at the PLC or the OPC server, MQTT broker, etcetera. And we, and we want you to write them that way. But your clients, your end users, your engineers, they probably don't need to consume the data that way.

32:13

Speaker 3

And certainly the folks on the business side of your organization don't want to see engineering level tag naming, and we want to be able to offer tag aliasing and then tag grouping or restructuring into the things that those tags describe, like assets. And so we have a gatekeeper service within our solution called views that allows you to create as many virtual views of that historical archive as you would like. So simply put, I can take a group of 100 tags out of my 5 million or my 5000, whatever, and I can organize those 100 tags into a special grouping, change the naming of those tags, alias them, and provide them to a certain group of clients without ever impacting the historical record.

33:01

Speaker 3

So now when my clients come in and want to have access to production data, maybe they're looking for bottle efficient bottling efficiency. So I want to know the good count versus the production. So I know what total count. So I know what my production was for today. And I want that for all of my lines across the organization. They can just come in and look at good production tag for every line without having to wade through pages and pages of tag names. And that has been so powerful, it's made it so much easier for people to get access to the data they need and avoid the data they just frankly don't care about.

33:44

Speaker 2

One smart guy, I wonder who that is. Once told me that it's like looking at the world through different lenses.

33:53

Speaker 1

Yes. Very quickly, before we lose track of this topic, I very quickly want to speak about the sort of licensing or pricing that you mentioned earlier. So when you're talking about this virtual view, my understanding of how those sort of models and associated pricing works, isn't that an expensive way to get your asset views if you pull that specific tag into every one of those views? Or is canary licensed differently?

34:24

Speaker 3

Yeah, so it is absolutely licensed differently. In fact, the only thing we license is the number of tags you want to store in the actual historical archive. You can place one of those tags into as many virtualized views as you would like and there is no impact on the business model. In fact, we stop licensing tags at 40,000 tags, the entire system goes unlimited at 40,000 tags. And so that gives you ultimate scalability to start with a ten or 15,000 tag historian and just go crazy with it because everyone is getting more IO. No one is getting less IO at this point.

35:05

Speaker 1

Yeah, yeah, 100%. You good?

35:12

Speaker 2

Yeah.

35:12

Speaker 1

Sorry, it looked like you're going to ask a question there. It feels like we're to get in a question. Jeff, I wanted to quickly, you mentioned earlier about. So obviously the objective is to get the right information into the hands of the correct decision makers and the more context that we have there, the better the decision making capability. What is, what does that look like today in 2020, that delivery of that information? I know, I know that, you know, it's potentially, first of all, it's got to be accessible. Those dashboards of that information in that context of contextualized information has to be accessible from anywhere, any device. I know that preferences to get that kind of view in Excel. A lot of people sort of still digest and work with that draw, not raw data, but that contextualized raw data within Excel.

36:05

Speaker 1

What does that look like today? What is your gut feel? Or what's the feedback? Or what have you been seeing in terms of the delivery of that information?

36:14

Speaker 3

That's a great question. So first, let's divide our information recipients into two groups. We have our live ongoing control process group. They need to consume real time data with very short historical viewpoints. So this is my control operator who is watching his pressures in real time with a 30 minutes historical trend. Yeah. Now that side alone, typically we're taking our data and writing it right back into a SCADA system, visualizing data onto HMI screens, and then sometimes offering maybe daily or shiftly feedback reports to operation supervisors. That is much what you would expect it to be, I think. Note that all of our tools tend to be based in HTML, or you can reach from excel and pull processed or raw data or event based data right out of the historian locally on the control network. So there's nothing that's probably too special there.

37:24

Speaker 3

If it is special, then you're in a really tough spot, right? Yeah, but the second part, and really the highlight around the COVID side of things is this idea of remote access. I'm not on the control network, I'm not as concerned with real time, last 30 minutes type of data. I want broad picture data. So how do you get the data and make it easily accessible to those types of individuals? Well, we have found that decoupling the reporting module, which we call Axiom and our view service be able to get Axiom away from the historian, not require that it sit on the control network, but instead let it sit on a public facing web service, gives you the ability to give

third party access to the data that's important to them. Again, virtual views, asset modelling, only see what you need.

38:17

Speaker 3

But even more important, build out these HTML dashboards and automate the reporting of those dashboards into individuals inboxes and never require them to have access to the system at all. So it's a PNG file, an image file that shows up and it's a condition based report that shows them, you know, a seven day, a 30 day, they're able to get all the information they need and with zero effort. That has been really successful, especially right now.

38:50

Speaker 2

And Jeff, I also presume that obviously, if I want to make changes to these dashboards, I mean, I don't really need specialized skills to understand how to pull data out of these stories and write queries and all of these things. It is a very simple drag and drop kind of environment to actually build potentially your own dashboard. If you want to now look at a specific asset or monitor something specifically through the action tool, you're right.

39:15

Speaker 3

The design tool is built right into the browser, and it's drag and drop. So I'm a young dad. A young, heck, I guess not anymore. But my kids are ten and 13, and they are pros at axiom. So they've taken the online canary academy, which is the equivalent of, I don't know, maybe 2 hours worth of videos, and they can build their own little dashboards. And maybe that's, you know, maybe that's not actually recognizing a low learning curve because our kids are all smarter than us at this point, but no, it is.

39:53

Speaker 2

And Jeff, obviously one way to consume the data is for us as humans to make these decisions. But potentially, there's other systems and other business solutions as we spoke about, that actually need the data as well. How easy is it for Canary or other systems to make this data available? Can anybody just get the data out of a canary system? Is it a black box? Is it open? What do you guys have that for there to actually take this mass amount of data and pass it on to potentially other expert systems or ERP solutions?

40:29

Speaker 3

Great question. And we certainly don't want to be casual around access to the data because security is so important. But canary, everything about canary is built on top of .net dot. And so we are being installed on Windows systems, Windows platforms. We incorporate active directory to be able to qualify clients and access immediately. But then it's really important that we make sure, as you're hinting at Lenny, that it's not just Canary tools that have access to the data, but that our proprietary NoSQL database actually can have SQL queries written against it via ODBC connector. But the most common two formats now are really using our publishing service to get real time data on a regular interval published out to third party systems via JSON, as well as our web APIs and our MQTT publishing service.

41:24

Speaker 3

If your application can benefit from a SQL query, from a web API call, or from an MQTT subscription, it will be no problem to get data out of canary. And if it doesn't, we still publish via OPC HDA as well.

41:42

Speaker 2

Awesome.

41:44

Speaker 1

Perfect. I can't believe we've been chatting for so long. It doesn't, it feels like it's been ten minutes. Jeff. Yeah. Thank you. Thank you for your insights. Fascinating conversation. It feels like we can chat about the topic for at least another hour. There's so much to talk about and certainly quite topical at the moment. Your prediction for

the future of data and information, are we close to predictive? Are we close to prescriptive? What is your prediction for the future of data collection, analytics, and decision making as a whole? That's probably not a fair question. What are your predictions for the future?

42:29

Speaker 3

Okay, so that's a great question, and I will throw answer out there without any confidence, but I think one, we are still a bit away because I know firsthand there's the common struggle for organizations is still to collect data at a high granularity. So what has been missing has been the bandwidth to pull things at extremely fast pull rates. We can read data down to ten milliseconds continuously. There's very few organizations that are taking advantage of that because they don't have the throughput. So that's the first hurdle. Maybe 5G helps solve that. We'll see. But let's assume that we've got the granularity and everyone's collecting data, and they've been collecting data long enough to run some of these algorithms.

43:26

Speaker 3

On the machine learning side, I think what you're going to see is you're going to see a lot of third party companies, as they already have started emerging to try to do things with this data. And it's going to be asset specific. If you're in the wind turbine business, you're going to be looking at contractors who have some special machine learning algorithms around wind turbines. Right. But all of that over time, I believe, is going to move. I think you're going to see, and this is where I come back and say the big three. Emerson, for instance, is so important. They are going to acquire this technology.

44:08

Speaker 3

They are going to then move this technology to the edge where the asset sits, and it's going to be hardwired and built in and ingrained in the assets you're purchasing and putting onto your lines, into your process. That's my prediction. I don't think we're going to have to do it post fact I think it's going to happen in real time at the edge. I think we'll see that loop, I hope.

44:31

Speaker 2

Great. Well, one thing that I can take away from this session is one. Okay, sorry, it's two words, but that's definitely the words purpose built. And I think the Canary Labs solution is really purpose built to really give you, at the end of the day, a complete PM solution to help you to improve your analytics and definitely to help you make better decisions on your assets for your day to day manufacturing environment.

44:58

Speaker 1

Yeah, for sure.

45:00

Speaker 3

That's a great summary and I really appreciate the opportunity, guys, to talk about it today and to your listeners out there. Share this, not because of canary, share this because of what the guys at eliminate are doing and help them grow this. I think this is going to be a great podcast, not just for South Africa, but for the global industry, and I appreciate the effort you guys are putting into it.

45:24

Speaker 1

Fantastic. Jeff, thank you for your time. I know you were due to be in South Africa sort of over this period. Obviously, a lot of folks plans changed quite significantly the past few months, so we hope to see you in South Africa fairly soon.

45:37

Speaker 3

Oh, I hope so. Guys, my last trip there has been my favourite travel experience around the globe. There's just. There's something about South Africa that's just captivating and I think, and I'm pretty sure it's the people. So I appreciate and look forward to seeing you guys again soon.

45:54

Speaker 1

No? Fantastic. We like to believe that we're some of the friendliest people around here in South Africa. We certainly enjoy hosting folks from all over, but, yeah, we definitely look forward to seeing you in South Africa soon when the dust has settled, so to speak, around this horrible pandemic that's affecting so many lives and economies all over the world. Thank you for your time, Geoff, and good luck for your move to Texas. Exciting.

46:18

Speaker 3

Thank you. And I agree, I think most hospitable people, but also the best rugby, right? Is that.

46:23

Speaker 1

Absolutely. Well, officially, according to the World cup standings and last win, definitely the best rugby playing country in the world.

46:30

Speaker 3

There we go, guys. There you go.

46:32

Speaker 1

Thanks for remembering that.

46:35

Speaker 3

All right, guys, well, thank you for having me. And cheers. Cool.

46:38

Speaker 2

Awesome.

46:38

Speaker 1

Thanks, Jeff. So next week, we're going to have a look at the.