

# Episode 14: Empowering responsible AI

## Host and Moderator :

Dr. Andy Packham, Chief Architect and Senior Vice President, Microsoft Ecosystem Unit, HCLTech

## Speakers:

Srini Kompella, Senior Vice President, Data and AI, HCLTech

Robin Lester, Senior Cloud Solutions Architect, Microsoft

0:12

Hello, this is Dr. Andy Packham, Chief Architect, part of the Microsoft ecosystem here at HCLTech

0:19

I'm here today joined by Srini Kompella, who's head of our Data and AI practice, and Robin Lester from Microsoft.

0:28

Why don't you just quickly introduce your roles and then we'll get into some of the questions.

0:32

Yeah, sure.

0:33

So thanks, Andy.

0:34

I am Srini Kompella, I lead data in AI practice globally within HCLTech pleasure to be here.

0:42

Thanks.

0:42

So yes, Robin Lester.

0:43

So I am an architect in the partner side of Microsoft.

0:48

My role at Microsoft is to help partners such as yourselves to create solutions in the AI space.

0:56

Thanks and look, thanks for joining AI.

0:58

Everybody's talking about AI today, but it's not easy.

1:02

And we're finding, you know, in our customer conversations, there are lots of challenges.

1:07

So kind of like ask both of you what are the biggest challenges you're seeing in those customer conversations and how, how are we going about addressing those?

1:15

So if we, if we take it well, there's, there's obviously a multitude of different conversations.

1:21

One of one of the key ones is around misunderstanding about what AI is.

1:28

We sometimes find customers wanting to just take data and AI that data, which obviously is not only not practical, it's also goes against a lot of the our kind of responsible AI principles as well.

1:44

Just throwing AI at something without any really clear goals or guidelines around it.

1:51

So yes, that's sometimes a problem.

1:52

And then on the other side, when we do find people have a a really good use case for AI, they will start down that route only to find that they don't have the base state or the base system set up in place.

2:06

So sometimes, you know, AI is a lead in to kind of develop that back end arch infrastructure

and architecture first to get better quality data to them present the AI because of the AI is will only do what you tell it.

2:25

If if there were biases in the data, if there were problems with the data thing inconsistencies, AI is going to pick up on them and represent them as well as the things you do want it to represent.

2:33

So yeah, a lot, a lot of conversations around data.

2:38

And then we also found a lot of a lot of blockers around organizational buy in as well.

2:45

So it's with any kind of technology that gets that has a lot of hype around it, there'll be people in an organization wanting to do it.

2:55

But ultimately to implement AI, an organization needs to be an organization wide conversation to do it the correct way, to do it responsibly and to have enough buy in for people to be able to be constructing those projects and doing them correctly.

3:15

Actually, it's absolutely, I mean, a lot of, you know, similar themes that we've been hearing from our customers as well.

3:23

And as we are engaging with customers in the AI journey.

3:25

I mean, one of the key themes that we are now beginning to talk about is how do enterprises scale AI at that enterprise level?

3:36

So, you know, I think 2023, there was tremendous amount of, you know, interest, public interest, you know, getting generated with the AI.

3:45

And as Rob said, I think, you know, there's also a, you know, bit of an understanding that needs to be developed further.

3:51

You know, Gen.

3:52

AI is taking over.

3:53

But really, I mean, what is the overall AI, you know, how does it blend together with generative AI to deliver the business value?

4:00

So there's absolutely that understanding.

4:02

But I think 2023, you know, helped people start figuring out how AI can benefit air prices, how AI can benefit employees, how AI can benefit customers, how AI can save costs, etcetera.

4:18

So, but in terms of the true value, it really comes when you put AI into production, right?

4:23

It's really about, you know, deploying it in the systems, using it day in and day out.

4:27

If you're thinking about customer experience, I mean, the customers need to feel it day in and day out.

4:32

It's not just about experiments.

4:33

So one big challenge that we've been seeing is that, you know, a big chunk of experiments don't make it to production.

4:40

It's a very high percentage, right?

4:43

So, and, and that seems to be the limiting factor right now for enterprises to truly realize the power of AI.

4:50

So then if we really double click on, you know, what's really going on there, because there's a lot of effort and dollars that went into the experiments, you know, in the last 1518 months.

4:59

So if we look at it, the number one is, you know, what are the, I mean the business cases having a clarity on what are the business outcomes or the business problems that, you know, enterprises are trying to apply AI for, you know, it's no longer about, you know, proving whether the technology works or not.

5:17

I think the technology is already, you know, quite powerful.

5:20

It'll continue to evolve.

5:21

But what are the business problems then once those are identified?

5:25

I mean, that that's really about translating the use cases to business cases.

5:29

So we are seeing large number of use cases.

5:31

Some enterprises have thousands of use cases, right?

5:35

But how many of them really, you know, get converted into business cases?

5:38

So if when we start applying the filters, that percentage is really low today.

5:44

So then it's a question of, you know, even the detailing of the business cases.

5:48

Rob talked about, you know, responsibility AI.

5:51

I mean, should you really be applying AI here?

5:54

Fundamental question, right?

5:56

So, you know, why are you applying AI?

5:58

That's one of the starting points of being responsible about AI.

6:02

Then the other aspect is also in terms of, you know, figuring out what are the true costs of putting AI into production.

6:08

So a lot of, you know, companies are beginning to realize that they have to think about an AI state, you know, before putting everything into production with the right guard rails, with the right governance, with the right corrective actions, etcetera.

6:21

So, and, and once people start thinking through this, you know, then I think Rob alluded to it as well, data, a lot of the experiments are being run with, you know, sample data or thanks to generative AI, there's synthetic data that is being used, but then the real world data can be quite different.

6:41

So that's the second big blocker we are seeing, you know, in terms of, you know, moving into production.

6:47

Then the next aspect is AI governance.

6:50

So we're beginning to see some companies establish AI councils to drive better governance from business perspective, from an IT perspective, from a compliance perspective.

7:00

And it's really about, you know, bringing in the responsibility AI, you know, as a design element rather than an afterthought.

7:07

So that AI governance is the next big thing.

7:10

The final one is in terms of, you know, having a blueprint for scaling, right?

7:15

So we're beginning to see with all the use cases, you're not getting processed.

7:20

We're beginning to see companies, you know, thinking about establishing AI estates so that they can really scale it in enterprise level.

7:28

So Srimi, I'm not both of you mentioned responsible AI and I think that is that is critical.

7:34

There's massive amount of hype, but it's about driving that conversation about value.

7:41

And for me that comes from becoming trusted.

7:44

So again, how, how do we do that?

7:47

You know, how, how do we kind of establish trust and, and across what dimensions do we need to establish trust?

7:53

OK.

7:54

So obviously as a public cloud provider, we live and die on the level of trust that people give us.

8:00

No one would use this year if we if we weren't trustworthy.

8:03

So we have, you know, we've got Microsoft trust website that kind of goes through all the

different certifications and that kind of stuff that we have to ensure that we do things that way.

8:16

We have, we have a lot of principles as well internally that control why you know, what we're doing what, what data we are aren't allowed to look at.

8:29

But I, I think ultimately trust comes down to having transparency, which one of our fundamental kind of AI principles, just being really open about where the data is, being open about who has access to the data, how it's being used basically is not by us.

8:54

And also having feedback mechanisms.

8:58

So ultimately the, the, the things that we, that we do to, to make sure people trust us at the same kind of things that we put out in our, in our responsible AI principles.

9:12

So, yeah, having transparency, having feedback loops, we've learned a lot over the years about how to do AI responsibly.

9:21

We've got, we have a very long history pedigree in, in that place all the way back from when we released a chat bot that about, oh, it's called Tay about.

9:37

It must have been about 10 years ago, probably a bit longer.

9:39

And it was a really good example of well-done AI that didn't have responsible kind of principles or responsible guard rails in place.

9:52

It was a chat bot that learnt.

9:54

So it was basically we, we released it and within the first day, a lot of people from 4 Chan had gotten and started talking to it taught it actually it was a learning chat bot taught us some really bad things.



10:07

It turned into a racist, over sexualized 14 year old girl very quickly.

10:13

We took it down.

10:14

We learned the lesson.

10:17

And then and, and from that, you know, we, we developed all the principles of, of what you need to have in place to make sure that you're not doing harm to yourself, you're not doing harm to your customers, you're not doing harm to any third parties and you're not taking away anyone's rights.

10:38

Yeah.

10:39

And, and so we, we kind of, we kind of boiled that down for into Microsoft Research releasing a paper that is called the, the responsible AI maturity model, which basically goes to all the levels of, of responsibility in, in different areas from, you know, organizational wide to how you're developing your applications.

11:03

And obviously we have lots of lots of tooling's and stuff in this place to help partners and customers make sure they're doing, I think, creating things responsibly.

11:13

Yep.

11:13

So I mean, yeah, I mean, thank you.

11:15

We all remember the story.

11:17

And I think that that, you know, what always impresses me is the transparency there.

11:20

Yeah, this is a story we learn.

11:22

We all learn from those kind of things.

11:24

So I mean, thanks for that, Srin.

11:26

I mean, that trust story again.

11:28

Oh, absolutely.

11:29

I think, you know, the key words are trust and transparency, you know, and you build trust with transparency.

11:35

You know, you, you, you know, they kind of, you know, feed off each other, right?

11:39

It's, it's essentially how we trust individuals, right?

11:41

I mean, it's, it's constant interactions.

11:44

You know, there's credibility that goes from transaction to transaction.

11:47

So it builds trust.

11:49

So I mean, you know, so that that's a simple, probably the simplistic way to look at it, like trust and transparency.

11:54

Trust would verify, but then there's a tremendous amount of hard work that goes behind to make this happen.

12:01

Just doesn't come easy.

12:03

So, you know, there are regulations in Europe, you know, that create a, that provide a really good framework to start with, right?

12:11

But you know, so it's comprehensive.

12:12

There are best practices that are, you know, published by economic institutions.

12:17

You know, Rob talked about Microsoft.

12:19

I mean, you know, we, we, we've got some checklists and templates, so to speak, right, all to aid enterprises to, you know, establish, you know, the trust and transparency.

12:29

But if you really double click down, right?

12:31

So one is, you know, how do you take this and create an enterprise level, you know, policy framework, right?

12:37

So it's got to be there saying that, you know, I mean, we're going to use AI, but we're going to use AI responsibly, right?

12:44

This is how we're going to do it.

12:45

This is how we're going to, you know, create the transparency for both the creators and the consumers of AI, right?

12:51

The creators, you know, if you look at it, I mean, they've got to like, I mean, do it by design, not as an afterpath, right?

12:58

So that then gets cascaded down everywhere.

13:01

I mean, previously we were talking about it, right?

13:03

Even on a business case.

13:04

Why AI, right?

13:07

And as part of that, I mean, some of the things that Rob talked about, you know, need to come in.

13:11

Is it harmful?

13:12

You know, is it prejudicial, right?

13:14

Is it, is it creating a negative impact?

13:16

So then I mean, it's got to be a question mark saying that, OK, we've got to flesh this out more from business case standpoint before you move forward right now.

13:25

There is some other element as well.

13:26

When we were talking about AI mean responsibility AI, right?

13:30

I mean, the models, the people who are creating the models, you know, yes, there could be bias and you can control it, you know, and, and you can actually, you know, control it a lot better.

13:40

The bigger problem is essentially, you know, around the data that you're training the model on, right?

13:46

A lot of the bias because at the end of the day, I mean, the AI model, you know, it is garbage in, garbage out.

13:52

So if you're feeding, you know, a wrong set of data, you get the wrong set of outcomes.

13:57

So when we are talking about responsibility AI, we got to consider the data sets.

14:02

So if you're training the AI model on prejudicial datasets, bias datasets, the bias will automatically creep through AI and in, you know, you see the output by AI, but the real issues are in data.

14:14

So that's where when we are talking about responsibility AI, we got to look at it holistically across both data and AI, right?

14:21

So I think some of, you know, people are calling it as data poisoning because there can be willful, you know, potential bad actors who can, you know, push in, you know, poison data, which will impact AI, you know, negatively.

14:34

So how do you filter for that?

14:35

And how do you create the transparency around what, you know, we call as observability of data, right?

14:42

So I think it's holistic, you know, policies, regulations, frameworks, keywords being trusted in transparency.

14:48

And the big aspect is to look at data real closely, you know, not just AI.

14:55

So final question, Yeah, none of us can do this alone, right?

15:00

It's incredibly complex.

15:03

And you know, Microsoft will always been stringing the partner ecosystem.

15:07

Let's just talk a little bit about why partners are important, if you would.

15:10

Oh, yeah, absolutely.

15:11

So I mean, we're, we're obviously quite a big organization, but there's no way we could scale to the size that we need to support the customers if we didn't have partners there.

15:23

So absolutely partners are really just an extension of Microsoft.

15:30

And so with our partner ecosystem, obviously we have the division that I work in, in GPS global partner solutions to that essentially to support those partners to help them understand kind of the principles that we have to make, make help them with the technologies.

15:51

So ultimately we create a lot of tooling.

15:54

We don't really create a lot of, we just create some solutions, but we don't create a lot of solutions.

15:58

So we create a tooling, we give those to our partners.

16:01

Our partners create amazing solutions with that tooling and then they feedback to us what they need and so that it turns into a kind of circle of help and creating like amazing news things.

16:15

Yeah, and I think so, and I think that's really important that the ecosystem.

16:19

So you know, the final question, how do how are we evolving that ecosystem?

16:23

Oh, absolutely.

16:24

I think he said it right.

16:25

I mean, you know, that's exactly how we are working with Microsoft.

16:28

So they're built fantastic technology.

16:31

They continue to evolve it.

16:33

Our job as system integrator is to, you know, integrate those technologies and build solutions and deliver those, you know, help deliver those outcomes for our clients.

16:43

So for example, I mean, you know, we've created this entire, you know, initiative around scaling AI at enterprise level.

16:49

And to enable that, we're calling this as the Enterprise AI Foundry, right?

16:54

So that's really bringing in a lot of the elements that we were talking about.

16:57

You know, how do you bring in the data?

16:58

How do you bring in the AI together?

16:59

How do you bring in the infrastructure in the right way?

17:03

So you know, how and in effect, and you know, all these coming together help clients, you know, scale AI and deliver better outcomes for their customers, for their enterprises.

17:12

So as part of this, you know, we, we are stitching together the technologies from Microsoft.

17:17

We're also working with several other partners and you know, we are stitching together their technologies as well, because, you know, our customers are going to need a lot of these technologies to make things happen.

17:27

And they have the choice, right?

17:29

And, and we enable the choice right.

17:31

And of course, we work closely with Microsoft in Co-developing, you know, some of the assets, Co-developing some of the propositions.

17:38

So that at the end of the day, I mean, look, we are in the age of AI, right?

17:42

Whether you know, we want, I mean, we, we really are.

17:45

It's a question of, you know, how fast can we evolve in a good way, in a responsible way.

17:51

So applying the technology.

17:52

So that's our job.

17:53

So with the Foundry construct as an example, we bring in the Microsoft technologies, other partners technologies, we've invested into building some assets, all the sole purpose being how do we help our end of companies customers scale the eye and at a price level.

18:08

So the entire partner ecosystem is vital for us because we are not building all technologies.

18:13



We are building some level of accelerators to, you know, on the journey, but we are relying heavily on technologies, cutting edge technologies coming in from partners such as Microsoft.

18:24

So Robin, thank you very much.

18:27

I think you mentioned all the key things.

18:28

You know, it's about partnership.

18:30

Yeah.

18:30

It's about establishing, it's about thinking about the data and it's about doing that in a way that's kind of responsible, holds ourselves accountable for delivering this in a way that is beneficial.

18:40

So thank you very much.

18:41

Thank you.

18:42

Thank you.