

# The enterprise is changing. Are we?

Guest: Uli Homann, Corporate Vice President, Microsoft



**ULI HOMANN:** The whole notion of AI and agentic AI is actually not a technology conversation, as you already pointed out, Stu. It's also not even a change conversation. For me, it actually goes far more fundamental, because it actually alters the definition of what an enterprise is, what human work is, how work gets done, and so forth, which means all fundamental questions. And that's where I would start.

**VOICEOVER:** Outwitting complexity starts here. Welcome to The GuideLine.

**STU BROWN:** Uli, welcome to Guideline. This is a podcast I've been doing at Guidehouse.

Actually you're my inaugural one. We're starting. I did one with Ryan Cunningham. We're going to put him next. You get to go first. Want to really talk through the trends in technology. What's happening, what we think executives need to know. And part of what I want to walk through in this with you is, obviously, in our industries, we're very regulated.

And there's a lot of conversation around data and "frontier" and what it means, and all of these things, and some of that can be very confusing to those as it's coming at them from all angles.

So, I want to in our conversation, I want to talk to the executives but also talk to the technologists to understand a little about what's going on. And so, as we talk to these leaders, I think it's really important for them to understand not just the breadth and depth of what's coming, but also the velocity, which I don't think people still completely understand. But before we kick it off, I'd love for you, you know, to give an introduction of yourself and your role at Microsoft.

**ULI HOMANN:** First of all, Stu, great to see you and thank you for the invitation. I think you're super smart to put age before, I don't know, young energy with my friend Ryan. So, I love it. And I'm part of the cloud and AI business unit at Microsoft. I'm part of the engineering leadership team, and I own a function called customer-led innovation.

But that simply means I work with great partners and customers like yourself to think through what we can do for you today, and then work with you to figure out how the portfolio of investments we have made so far perform, and then learn with you how we need to change them, alter them, whatever the mode of modification or strategy change or something like that is required to really make these capabilities work for your scenarios.

Again, you already mentioned regulation or regulated environments as a key concern. That's obviously something that's popping up everywhere and so forth. I've been at Microsoft now, it will be 35 years in about 20 days. So pretty interesting times these days.

**STU BROWN:** Yeah. Pretty amazing. And I think that juxtaposition between the technology and the velocity coming from Microsoft and your partners and what you host on the platform, along with the, you know, the ever-changing environment. And I think, to your 35 years in the velocity, it seems like that the speed is getting logarithmic at this point. Just curious, I mean, how much as you sit in front of clients, and I remember when we were talking about, you know, digital twins and manufacturing and just-in-time, what does just-in-time mean for manufacturing and some of these other companies? Just how far are we away from those conversations now? I mean, the velocity has got to be crazy.

**ULI HOMANN:** Well, so let's first step back and look at why we are here. If you think about it, cloud computing was layered on top of ultimately the investments made by all the cable companies. When the .com craze was going on, there was a lot of cable capacity put into the oceans and land masses of the world. Cloud computing effectively layered on top of that work.

Then you started to see the data work going on. And then because of all these three: network, compute, and data, we got the AI wave. And that is all coming together now in what we have now termed agentic AI. Without any of these foundations, this agentic AI wave would not happen. And the speed or the velocity of innovation we see effectively is based upon this foundation. The speed, I mean, Stu, look, OpenClaw, Claude Cowork, Copilot Cowork, those things were three months ago not even a conversation. Now they are, oh, the future of work is going to look like this because of this. Now, I don't know what's going to happen in the next three months, simply because the innovation scalability is so interesting.

I heard somebody use an interesting word or phrase. I really like that one: To think about the Industrial Revolution, what we were able to do is scale human work through automation and energy and so forth, but we haven't been able to scale human ingenuity. If you now think about AI, we're now scaling human ingenuity, and that is something that

has never happened before. And therefore, I think you're going to see this immense, a wave of new innovation simply because you, Stu, can do things in parallel, where you just supervise and direct that you, in previous times, you had to do yourself.

And that's something that, again, we have never been able to do. Now we can. And you're starting to see results of that work.

**STU BROWN:** Yeah, I think that's a great point. And I want to get into the agentic frontier conversation and what that model looks like. I was on stage at the Microsoft BizApps and Copilot Summit this week in DC. And the question they said, you know, what's the leading barrier? Is it technology? Is it people or is it change?

And I'm like, yes, but it starts with change, then people, then technology. The technology is not the barrier anymore. And it's, you know, it's the change adoption, and what you're pointing out is a really great, great example of it with the Industrial Revolution. I think the idea of upskilling human intelligent thought and how we iterate and how we design is not innate to change.

I mean, we recognize when something helps us like a Copilot, but to then be able to have digital workers that we then ask to do tasks and OpenClaw is a good example that. I, as you might imagine, play around with it and it's like, you know, they're all out there doing their thing and we're not far from that.

**ULI HOMANN:** Again, I think what's now happening, which is a good thing from my perspective, is we are seeing this arc of innovation, but we're also seeing an arc of responsibility coming in. Again, you were in DC on Monday. Ultimately, we announced Agent 365 GA availability and stuff like that. But I'm now seeing a lot of customers realizing that they need something like Agent 365 to go and govern this, all this, harness, harness that innovation. But do it in an enterprise scale and safe way because we know agents will make mistakes. And now you need to be able to go and register them, have the right identity, observe them, evaluate the work that they're proposing to do before they do it, and obviously record what they did so that you can go and drive forensic value or forensic logic to see, did they really execute what we want?

**STU BROWN:** Yeah, I agree. So when you think about that and you're talking to clients especially adopting the innovation that's coming from Microsoft, in your case, as well as the partners that are building on the platform, and obviously broader than that, the interoperability of what plugs into the Microsoft environment, how do you define agentic frontier, and then how do you look at that operating model? What are you telling them they should be looking at?

**ULI HOMANN:** Well, so for me there are really three components that we have to think through. The first one, for me, is the whole notion of AI and agentic AI is actually not a technology conversation, as you already pointed out, Stu. It's also not even a change conversation. For me, it's actually goes far more fundamental, because it actually alters the

definition of what an enterprise is, what human work is, how work gets done, and so forth, which means all fundamental questions. And that's where I would start. I think everybody needs to step back and say, what do the people that I have in my enterprise, those valuable assets called employees, do in the future? And what does AI do in the future, and how do they coordinate? How do they work together? And do that across all of your business aspects.

Some of it will be higher priority than others, but that's something that is important. Microsoft calls that the "frontier firm," and I recommend at least understanding the framework that we put together, because I think it's a helpful little north star of what an agentic future would ultimately look like. But I'd start with people, processes, and the "what" and "why" of the enterprise, and then the "how."

I think that starts, because it's not about technology. This is about a fundamental change. Then the other piece, again, you hinted at it already. This world is going to be a lot of value-added players working together in an enterprise. There's never going to be an enterprise that just uses Microsoft or Google or SAP or Salesforce or ServiceNow, whatever it is. Everybody will have to figure out how to work together. And then last but not least, you need to start to think about, okay, now, great. I have a set of open standards that are emerging and some of them are established pretty well. I now need to figure out, how do I apply those in my enterprise.

And, Stu, your company is very well positioned to help your customers to say, okay, let's take MCP. So, the model contracts protocol, that's a great capability. But which of your enterprise functions do you want to expose to the agentic layer and what is "read" what's "write," what are change processes, and so forth. Do you need to put extra guardrails in place?

Those are some of the discussions that need to happen. Or agentic identity. What does that mean in an enterprise? Again, Chat is easy because Chat always assumes our identity. Great, easy. But once you have agents, especially autonomous agents working, they need their own identity. What should, what rights should they have and stuff like that. All of these are complicated issues that can be solved but require thought and effective investment of time and IQ to get that right.

**STU BROWN:** So, agreed. And when we think about, and my brain's going to who manages which ones, and there's the compliance management of the agents and the technology and access, and then there's the productivity management of those agents and assets versus with and in concert with the humans and what levels. And I think there's a combined hierarchy there.

But when you think about the, so the term “frontier firm” can be maybe not known. I mean, it’s a great term, and people know the term frontier. I mean, so maybe give a quick definition of how you view frontier firm. And then in that context of the, of what you just laid out on maturity curve, where are most companies?

**ULI HOMANN:** So, the way I decide to describe “frontier” is effectively a company that has accepted that AI is not a technology, but a fundamental transformation agent of how the company works. And if you accept that statement or that view, then you can start to say, okay, where do I start in my enterprise? Because you can’t do everything at once. And you start to really put AI into your core business processes, into the way you work, and not just as a technology that smart people like Stu bring into a company, and bolt on to existing capability.

Obviously, we have to deal with brownfield, we have to integrate what’s there. But fundamentally, from a conceptual model, you have to say, look, I’m moving to become an agentic company. I’m going to go and put AI everywhere, make it available to every employee, and expose every system in the right way, governed and controlled to those agents so that they can effectively go and really help my business processes move forward.

And maybe you re-envision your business process over time as well. Where you’re saying, look, we’re doing this because we had Stu and Uli as great people bridging the gap. But Stu and Uli are going to go into retirement at some point in time. So why don’t we rethink how we do these things so that at the end of the day, it becomes agentic.

So that’s my definition. And the second thing is on the maturity side, I see a weird kind of sprint happening. So, on the one side I see companies that are full ahead with various code automation agents. They’re rolling out solutions like Microsoft 365 Copilot to every employee. They understand that the next generation user interface is going to be one of those assistants for applications and stuff like that.

But then I talked to a bunch of companies this week while I’m in South Africa, and they all are super interested. They are very knowledgeable, but they are super slow. For example, capabilities like Copilot is only handed out after a business case per person, and then you, so it’s a very traditional way of thinking about it. It’s not accepting that AI is a fundamental change of what’s going on. They look at it as a technology, as a tool, to improve certain aspects of a job, rather than rethinking of how everything works.

**STU BROWN:** I think it’s a great point, and I love your comment, because I don’t think any organization is immune from this, that they look at Copilot like a procurement conversation. Well, it’s a tool. We’re going to give it to people. And I think because people are used to ChatGPT or, you know, that’s their first exposure. And, of course, they know how to Google things or Bing things. I apologize. They tend to not understand how to leverage AI in the

way you were describing. And in some organizations, there will be resistance. I would expect in some of our government organizations there's going to be resistance to, I don't want to lose my job or I don't do whatever.

And the reality is, we don't have enough talented people to do all the things we have to do in our organizations. So how do we maximize the talent we have and also take away some of the drudgery? One thing you touched on, and I want you to touch on this, because I want to get into the data side, is: I think capturing the knowledge of Uli and Stu, and not Uli and Stu the individuals, but Uli and Stu the best practice, to ensure that those that come after us, and the way I always describe it, and I'd love to get your thoughts on this.

There's the individual AI, the Stu AI, which does what I want to do, how I want to do it, and would answer. And then there's the role AI. And I want to make sure that the role that I perform and how I do it and the decisions I've made, you know, live for the next person that takes the role.

And like, I don't know if I want it living in Anthropic or in OpenAI or in whatever. I want that memory layer to live in my organization and be able to leverage any models I decide to pick in the future. So, I'm curious, as we think about data sovereignty, and I want to get your definition of that, how are you thinking about where the memory layer should live for these AIs, as opposed to in the host models? Is that something you're talking to clients about?

**ULI HOMANN:** So again, first of all, let's answer the first perspective because I think you're up to something. Because at the end of the day, one of the big things is demographic change we have in the West, specifically. And part of this is with prompting, even just prompts, you can capture the style and the questions or the thoughts that experienced employees have. While there's obviously PII data and GDPR concerns and other things, that have to be observed.

**STU BROWN:** HIPAA. All that stuff. Yep.

**ULI HOMANN:** This is a great way of thinking about this. The other way to think about it is when you look at agentic AI, you're moving the human into a supervisor role and the agent or the agents into an operator role. And operators today kind of tweak things in manufacturing environments or in energy systems, whatever, because on their experience and now what we can do is, because AI would propose, and let's say in the first step, here's how I would do the refinery flow or the milling experience for grain or whatever I do.

And then the supervisor can say, eh, based upon what I'm seeing here, let's tweak this, tweak that, and this, which has two interesting side effects. The first thing is, we now captured Stu's experience as an operator because he now is able to tweak what the agent or the AI layer proposes. And then we can run simulation that says AI proposed this.

Stu changed it. Which one was actually better? So, you can not only capture what Stu's experience is all about, but you can also say, oh, is AI just didn't have the right data. So, Stu is right? Or was Stu actually all wrong all along, which is a bit naughty.

**STU BROWN:** It's possible. It's not. It's not improbable.

**ULI HOMANN:** No. It's again, it's a bit naughty, but it's a good conversation to have, especially when you think about manufacturing, energy management and those kind of things, where this efficiency and so forth is really important. So, let's talk about memory because again, Microsoft agrees with your perspective that a per agent or per agent subsystem memory model is just wrong.

A memory should, is, in the business context, an enterprise function, and all the memories that you have in your business work are actually property of your company. So, you should be careful what you put into your business tools, because the company should have the right and does have the right to retain all of that content and use it.

And having that managed in a central place with the right governance, with the GDPR perspective and so forth, instead of every agent having to do that or not doing that, even worse, then you get into trouble. And my view is also, as I said, this is enterprise data and enterprise asset. Therefore, it needs to follow the same rules as every other data asset that is hosted, no matter where you host it.

It's under your control. You effectively put it where you want it, meaning location-wise, you secure it, only you, meaning the company, has access to it. If the operator needs to help you in certain things, there has to be established processes, protocols like lockbox and others to open up those capabilities so that at the end of the day, you are sovereign in your management of that critical set of data.

**STU BROWN:** Yeah. So, that gets to Satya's comment around firm sovereignty, right? I mean, it gets to this idea of you establish that governance framework. I guess one of the questions I would have, and it's going to be a challenge for us in organizations that are regulated where, you know, retention of chats and other things are much shorter than they need to be.

This is an area where, and I think if I understood what you're saying, we need to take this, we need to retain it, we need to put it in that lockbox and we need to, if we have strong rules around who and how they can access that data, and that's the governance structure that you put on top of it with the security structure and Entra and other stuff that you have, that it allows us to say, we can retain that information because it's not going to get out to the wrong people, but we must retain it as an organization because that is our IP, that is our secret sauce.

And so the question becomes, how do we make sure we take that data in those proprietary systems along with our unique data, because things like procure-to-pay or supply chain or whatever, those are, that process is a process that everybody has, everybody's got procure-to-pay. Now the difference is how we do procure-to-pay, which is, I think, where you're going on the manufacturing side. I want to make sure that we apply how we do it to ensure that we're always applying, whether you're the least skilled or the most skilled employee or customer or whatever, you're getting the best answer we would give to anybody at that level all the time. And I think that's the promise of what that data, the data governance structure, can bring.

**ULI HOMANN:** The way I would rephrase it because I completely agree with you. If you think about the future of ERP and CRM, they are ultimately ontologies. If you really think about it, what a data layer is, it's an ontology. And now what we need the ERP and CRM system to do is to expose that ontology with the underlying data, processes, and stuff like that, into the agentic layer so the agents can reason over these data sets, as well as the connections between the data sets, to really go and do the interesting logic, because again, I see a lot of people building agents still as directed flows.

And while that is okay for certain scenarios and maybe the right thing to do for certain scenarios, it is not the default because agents need to, should be able to reason over things, and they come up with a plan to execute, whereas directed flows, if you really need to direct, fair enough, but be thoughtful because agents have reasoning capabilities and those are getting better and better. So, if they have access to the right ontology, to the right data that grounds them, you can, you will see remarkable results of how business processes are changing.

**STU BROWN:** Yeah, I agree. Look, I mean, I think advice to people out there: if you're using agents in a direct flow way, you're basically doing glorified RPA and that's okay. That may be all right. But the reality is agents are intelligent. The whole idea is they understand the context of what they're dealing with. They learn from the history and how you train them.

And then they start to react and they start to know where to go, just like a human would say, well, I need to go get this information to answer this question. And that's where it gets really interesting to me as I start playing around with some of the open source stuff, is being able to, you know, watch them interact and see how they govern each other, which is sort of fun, right?

And it's going to be unique to see how we do that inside of the Microsoft environment, because you could see that we're going to have all these different roles that, you know, you're going to, when you take over an organization, back to my role based AI, you're going to have an org chart, that org chart, I assume when I pull up an org chart in Outlook and I see who works for Uli, I'm also going to see eventually the agents that work for Uli. I mean, that's going to be built into the org structure. And I can ping it, if it's in there, and it says,

this is Uli's assistant for this, I could ping it and say, I have a question I need to ask Uli and this is what I'm thinking. It's going to answer me. So, I don't have to ping Uli. Because it knows how you would answer. You get the information, and that's probably this far away.

**ULI HOMANN:** I want to go back to something you said, which is, I think, something really important where we also need to think through what's going on. So, if you think about the future of work, one of the big questions that's in the market is, lots of junior work might not be required anymore. Again, we will see how that pans out.

But one of the changes that are happening is that, even an entry level person coming into an organization will actually have to start to say, hey, I'm managing resources, which is not something that they are used to. They're not trained for, because of the agent work that you said. Again, I have, let's say, ten agents working for me, that also means that I'm responsible for effective usage of the resource.

I'm yield management. All of these things that you and I have dealt with over the years because we manage organizations and other things. But if you're an entry level person, you don't know anything about this stuff. So, I think we need to start working on onboarding, training, but also with the universities and so forth to help people understand that organizational management of resources and so forth is something that every person will have to now learn.

**STU BROWN:** It's actually really a brilliant statement. I think, you know, as I think through, you know, how we learned coming in as junior engineers, junior analysts, whatever. You know, we went through an osmosis process of being in the room and learning. And, you know, we've gone away from all the in-office, even though that's come back a little bit. But there was so much we learned through observation and osmosis of being in the room and things went down, which we don't do anymore.

Well, look, it's always great to talk to you, my friend, and I look forward to getting together soon when you get back in-country. Or maybe I'll come visit you. I want to wrap every one of these with sort of, you know, what is your guideline? In other words, like, if you had to offer one guiding principle to our listeners based on what we talked about, who are trying to move their organizations, who are trying to understand, and what they should adopt.

What would be your guideline to them to say, this is what I would recommend?

**ULI HOMANN:** I would recommend two things, and I have to say two. The first one is, understand the frontier firm framework. You don't have to believe it. You don't have to use it as-is. But think about it, apply it to yourself, and then use it to really change the organization. And then the second thing is, really have that governance model with respect to identity, observability, evaluations, registries, logging, all of these things, in place before you scale agentic AI.

**STU BROWN:** Yeah, I think that's a great point. And I think it's, you know, for those organizations out there, I've been a big advocate of the frontier construct. Why? Because everybody wants to be smart. And they are, and they know their organizations. But take the framework that's been given to you and build around it. It may not be perfect, like you said, but it gives you the constructs and the building blocks to say, how does this apply to our organization?

And then I view it as, as bookends, and then start filling in your chapters between the bookends. Uli, my friend, thank you very much. I appreciate you taking time while you're in South Africa. I appreciate you coming on my inaugural podcast here at Guidehouse called GuideLine. So, thank you, my friend.

**ULI HOMANN:** Thank you, Stu. Fantastic to be with you and good luck with GuideLine.

---

The GuideLine is a podcast brought to you by Guidehouse. [Subscribe](#) wherever you listen to podcasts so you never miss an episode. To see more from Guidehouse, head to our website, [guidehouse.com](https://www.guidehouse.com).

## About Guidehouse

Guidehouse is a global AI-led professional services firm delivering advisory, technology, and managed services to the commercial and government sectors. With an integrated business technology approach, Guidehouse drives efficiency and resilience in the healthcare, financial services, energy, infrastructure, and national security markets.

© 2026 Guidehouse Inc. All rights reserved. This content is for general information purposes only and should not be used as a substitute for consultation with professional advisors.