

# USER FRIENDLY

## Scaling GenAI for the enterprise

**Host:** Hanish Patel, managing director and business and digital transformation leader, Deloitte Consulting LLP

**Guests:** Fletcher Previn, CIO, Cisco  
Irfan Saif, CIO, Deloitte US

**Hanish Patel:**

I'm Hanish Patel. And this is *User Friendly*, the show where we explore emerging trends in tech, media, and telecom and how they impact business operations and the world around you.

As we all know, artificial intelligence, commonly known as AI, continues to reshape the technology landscape at a rapid pace. And technology leaders are tasked with harnessing the power of it at scale for the enterprise. But how can organizations

implement AI and specifically Generative AI responsibly with the proper risk and governance mechanisms in place while creating business value?

Joining me today on *User Friendly* podcast to discuss how organizations can create the adoption of AI is Fletcher Previn, CIO of Cisco, and Irfan Saif, CIO of Deloitte US. Fletcher, Irfan, welcome to the show.

**Fletcher Previn:**

Thank you for having me.

**Irfan Saif:**

Yup. Thanks for having us, Hanish.

**Hanish Patel:**

All right. So, to kick us off, there's certainly unprecedented demand for AI today. And I'd love to start with kind of really understanding from the both of you what are the implications that AI has on the enterprise, and what are you both seeing and experiencing in your role as CIOs? And maybe, Fletcher, I'll begin with you.

**Fletcher Previn:**

Sure. So, within Cisco, we're really seeing three main categories of how we think about AI for ourselves, AI for our customers and partners, and how we embed AI into our products.

So, for ourselves... is sort of the category of things that we're doing to make Cisco employees more productive. How are we bringing AI the right way to help automate, help solve problems, help write software, allow people to spend more of their time doing higher-order work and less time doing the mundane tasks that all of us, uh, sometimes have to tackle. And there's a lot within that that we can get into.

For customers and partners... is how do we leverage AI to make it easier to do business with Cisco? How do we simplify our processes? How do we simplify licensing? How do we simplify the support experience and so on? And then how do we bring AI to bear to make Cisco's own products better? And, you know, if you think about the products that Cisco makes, whether it's collaboration or security or networking, so much of that from an IT perspective and an operations perspective is about finding that needle in the haystack in a timely fashion, finding the anomaly and being able to take action on it before something undesirable happens.

And so, AI is really being embedded across all of Cisco's portfolio whether it's in the cybersecurity products, in the full-stack observability side of the house, network management, or in collaboration products with things like Cisco Webex, where AI is used to do things like summarize meetings and read body language and nonverbal cues or do noise cancellation.

And so, part of all this is the IT department that I'm part of, we're creating AI infrastructure that can be used across the whole company.

It's shared AI infrastructure that allows us to create a blueprint for how other

people can do something similar, but also make that environment available for engineering and product and IT teams across the company. So it's sort of all three of those, you know—for ourselves, for our customers, and for our products.

**Hanish Patel:**

Love it. Irfan, if I could turn to you for that?

**Irfan Saif:**

Yeah. We are also seeing insatiable demand from the marketplace from all of our clients. We're seeing it for our client service professionals as well and how we go about all of the different capabilities within the Deloitte footprint.

So when we think about productivity enhancements, when we think about helping folks optimize on things that really are the higher value, higher-order capabilities that impact outcomes for our clients and for ourselves, those are the things that we're really deeply focused on.

We really started with development of various models, approaches to innovating in areas like IT service management, code development and product engineering for some of our client solutions, thinking about how we can use it for cyber, for network operations, for risk management more broadly.

So there's a whole host of things that we have been focused on. And we sort of look at ourselves, I'm sure like many others, as a reference implementation for our clients. And so, our clients want to know that we have walked in their shoes. They want to understand that we have a good grasp of the problem and understand where we have looked to try and drive the greatest benefits in terms of hypotheses, productivity, gains, outcome enhancements, things like that.

And we've also sort of validated that through some pretty large enterprise studies. So for the last five or so years, we've been running the State of AI in the Enterprise Survey, and we continue to see an uptick in the interest and the engagement around AI. So this past

year, 94% of them said that they believed that AI was critical for success. And it's continued to be on a steady increase.

And they sort of have ambitions around uses, around adoption, around challenges, and what they think the "art of the possible" is. Realizing that when you look at various data points from the outside—the *Wall Street Journal* ran something quite interesting just a few weeks ago, I believe—that the ROI wasn't as clear as we all might like as CIOs and business leaders.

And so, there's quite a bit actually of experimentation. Not everything is going to be successful. A lot of this is not just going to be the technology; it's going to be change management. But I definitely think the interest from professionals, from our clients, from all of our technologists is incredibly high to really push hard to find those use cases that then generate the maximum value for us. So, it's exciting times to be a CIO, for sure.

**Hanish Patel:**

So, I want to stick with something you just said there around that high percentage—I mean, 94%, around seeing the adoption of and the use of. But on the same token the ROI isn't abundantly clear in every use case or adoption, with the fact that there's so many potential use cases for the adoption of AI and, in particular, say, Generative AI as well, how can leaders like yourselves really select the highest priority use cases and, frankly, the ones that are going to give the highest ROI?

**Irfan Saif:**

Yeah. I mean, I think there's a few dimensions to it. And I'm sure, you know, Fletcher will have some similar views. One is around speed to execution. You know, there are areas—take cyber, as an example. Cyber is an area where there's just significant talent shortages in the marketplace globally.

And so, it's an area where the significance, the criticality, of those capabilities, combined with urgency to detect proactively hopefully. But if not, then, in arrears, detect events and remediate them quickly. Those

are all sort of candidates that say, "All right, this is an area we have to invest."

When you've got solutions in the marketplace like the ones that Fletcher talked about earlier, you want to start to explore these and others and say, "Where is the opportunity where even if the ROI isn't completely clear on a spreadsheet, you know from a business standpoint it's an area that requires investment?"

Product engineering is another one. I talked about cyber network operations, similarly. Infrastructure management. There's a number of different areas where, really, we know from how many people it takes the time to deliver an outcome and the efficiency or productivity benefits or uplift to those who are impacted by those capabilities. All of those are reasons to say, "All right, well, how do you focus?"

So, quickly, I'd say speed to execution is one. Certainly cost reduction is important. Reduction of complexity. All those kinds of things are going to be key drivers to say, "All right, across the IT footprint, those are areas that feel right for investment." And as long as we can form some strong hypotheses and have robust use cases and a willing group

of participants to really engage in improving these hypotheses or shaping them, that's what's really important, and that's what we've tried to drive over the last couple of years.

**Fletcher Previn:**

Yeah. And I think just to "plus one" a lot of what Irfan said there, like, in the cyber domain, it reminded me... if you think we have cyber challenges now, just wait. When you can iterate on a threat thousands of times a second using AI, at some point, the only thing that can combat AI is more AI is what the US military calls the third offset.

It has become so inexpensive to be a bad actor that the rate at which threats are being created is sometimes exceeding the rate at which people can patch and remediate those kind of threats. So those are ones where it's just not an option to not participate in how you create safe infrastructure and a safe environment for people and bring AI to that problem.

But I think also, more broadly, historically, the answer to a question or a problem needed to live somewhere in the data that you were looking at. But that's not the case if you can reason it over data. Then you can start to answer questions to which there was

no previously known answer. And so that combined with the ability to generate content, right, Generative AI, that's pretty interesting.

One of the exciting use cases, I think, is in software development. We've always sort of held the belief that there's no compression algorithm for writing software.

Well, now, turns out AI is actually quite good at writing a lot of kinds of software. So, we can shorten the time that it takes to develop software. Use cases like AIOps—reducing outages, collecting system data, running analysis on that, trying to get multiple hours' headlights into a problem or an incident before it happens, and being able to achieve high availability, and over environments, is across network, compute, storage, applications, container, environments. That's a high-value business case to make.

And then some of the obvious ones around support. I think that's one that all of us are kind of aggressively getting after: How do you bring AI to customer support? Help desk? Being able to understand natural language and quickly not just get the right answer but sort of complete the process and actually take some action and fix a problem for someone.



And then in the employee domain, one of the things we're working on right now is creating a digital equivalent of a friend who has worked at Cisco for a long time and just knows the answers to anything you need. You know, you take a value of an LLM [large language model] so that you can use natural language and combine it with some technology like, retrieval, augmented generation or reg, so you can feed it proprietary information that is not on the public internet. Data from sales environment, your support environment, whatever the repositories are.

And then you can do things like have a seller, for example, ask it to help you draft an email to the CIO of a large company and maybe explain why a Cisco firewall might be a better fit than a competitor's firewall. Or comb through all the customer support tickets and see which customers might have experienced a given bug, and what was the fix?

So you've really got this gamut of high-value use cases that span employee productivity, seller effectiveness, customer service, security. But I think the key for me is, we want to operate the enterprise as an AI-driven organization, both for cost-optimization reasons but also development of new capabilities.

**Hanish Patel:**

So, based on some of those instances just described there— Irfan, what you also mentioned earlier—given our base of listeners' focus around technology, media, and telecommunications, what are some of the use cases that you are seeing that are creating the most value in that space of technology, media, and telecommunications?

**Irfan Saif:**

Yeah, Hanish. Support is a key one. How quickly—when you think about on the back end, a call center, contact center—how do you empower those there to not just follow scripts but very, very quickly, at their fingertips, be able to pivot to a problem and accelerate time to resolution?

At the same time, the self-service component on the front end, you know, can you eliminate level one, level two support calling altogether by empowering a professional or practitioner to engage with a chatbot or engage with an interface that allows you to do some of the things that Fletcher talked about, right?

Those kinds of things are high on the list right now. We're seeing lots of organizations implement. We are certainly looking at that as well. And again, some of it is down to change. How do you drive that change and adoption?

I certainly think to the other example that Fletcher mentioned around selling, I would even call it, more broadly, around just professional engagement overall. You know, how do you drive summarization? How do you integrate that with things like nudges to drive professionals to accelerate certain behaviors that you know are part of the life cycle of their particular job persona?

And now whether they're in sales, whether they're an administrator, whether they're a marketer, whether they're an engineer—whatever their function might be, there's the integration of things out of the collab environment, integration from the workflow perspective of different engagements, driving communication. Those kinds of things allow you to start optimizing each step of the workflow, which results in pretty substantial aggregate productivity uplift.

And so, those are the kinds of things we're seeing right now. And then, I think there's just an increased expectation across a lot of these job personas that there is going to be some kind of AI companion, assistant, augmentation—you use whatever word you like—but there's an expectation that your SOC [security operations center] analyst has something to accelerate their ability to research an event or an incident. Your engineer has a capability now that's going to help them deliver quality code more quickly.

So there's lots of those kinds of things, I think, Hanish, particularly in the tech space

that we are seeing a lot of. And we're seeing an expectation from technologists in the marketplace, our own technologists, that they are equipped with these things to stay ahead of the pack.

**Hanish Patel:**

Fletcher, anything that you're seeing specifically from your side?

**Fletcher Previn:**

Well, I think in addition to the use cases I mentioned and everything Irfan said, we sort of philosophically approach AI as a force multiplier for human potential. You know, it's an assistant, not a replacement.

The job of securing our enterprises has gotten significantly more complicated, more porous, the attack surface is much larger, everyone's home is essentially a branch office. We need to understand what is happening across networks that we don't manage and what the experience people are having across the public internet and using SaaS [software-as-a-service] tools are.

And so, being able to understand the telemetry that's coming in in a timely fashion and have some help in interpreting all of that, and making sense of it and then taking action quickly, is high value from an infrastructure and AIOps perspective.

And I'll just give maybe one quick example of that. It sounds like a small thing, but you take a lot of small things and it's a big thing. In offices where we have wireless access points, typically around once a year, we would go through and adjust the power that those wireless access points are putting out and attenuate the radios based on what we're seeing from traffic and how it's performing.

It's a manual activity. We typically get to it once a year. And now with AI, that's something that's happening automatically every 30 seconds. And so, you can imagine just sort of multiplying that across your operations. How much time does that return to the teams to focus on other things?



So, the “needle in the haystack” problem, understanding the real-time view of what’s happening, the automating taking action and remediating of things, especially in the security domain, right? Very often, that’s a data overload problem. And how can I take all this telemetry from these different systems that are coming in, rationalize it, and put together a real-time view of how something may be trying to get into my environment and move laterally and what it’s trying to do. And what policy changes or rule changes do I need to take to defend my digital estate? AI can really help with that in a very significant and meaningful way.

**Hanish Patel:**

So as I think about some of those use cases you both mentioned around all the great things that AI can do, and as both of you in your roles as CIOs, looking at “how can I use AI” and in some of those examples as you clearly did there—like to foster innovation, bring more efficiency, but at the same time think about how you’re safeguarding the organization itself through the use of internal and external—could the both of you talk about some other kind of risk and governance structures that need to be in place to really securely implement and make the most of AI?

**Fletcher Previn:**

At Cisco, like everyone listening, I’m sure, we’ve spent a lot of time and energy to make sure that the way that we’re using AI is consistent with our values and compatible with our business model. And so, Cisco has established a set of policies that govern our responsible use of AI, and any new use case for AI that we bring into the company needs to be reviewed and approved.

And the way that we brought in technologies like ChatGPT or Microsoft Copilot, we’re doing that in such a way that none of the prompts or the data that Cisco employees ever input leaves Cisco’s environment, and it’s not used in any of the public training data. You know, that’s a big concern that the questions that people are feeding doesn’t then become sort of public domain. And there are ways that we’ve kind

of brought that in and can audit and put guardrails and filter out any PII that people inadvertently might put in—um, personally identifiable information—mask names, addresses, email addresses, and so on.

And then for any code that is written with the assistance of AI that needs to be reviewed by human beings, that code needs to be documented that it was generated by an AI system. And then we need to be able to have transparency and explainability in any decision-making that is assisted by AI. How did the AI arrive at this conclusion? What were the steps to reverse engineer how it ended up here? And do we agree? So it’s more than a human in the middle. It’s sort of a human as final approver.

**Irfan Saif:**

Yeah. So not surprisingly, very similar to Fletcher, we have developed a framework and been using it for ourselves and with many clients for the last couple of years called the Trustworthy AI™ framework. And it hits on many of the attributes that Fletcher talked about—so security, privacy, transparency and explainability, ethics. And then obviously, the governance around the underlying data itself.

And, I think, that one thing that becomes really important as enterprises start to experiment and then start to move to the right, in terms of maturity and scale and breadth of adoption until it’s really ubiquitous and no one’s thinking about it.

Through that continuum, I think there’s an element of how do you govern this? What’s the executive management responsibility around these parameters? Not just setting the policy but, you know, really looking at how do you verify that these are, in fact, the safe practices that you expect as an enterprise, that your shareholders or stakeholders expect, and then also that your clients expect.

And so, in our case, even going into client boardrooms and making sure that from the board down to the server room, everybody is on the same place. There’s a common



taxonomy, there's a common set of criteria and, you know, some programmatic way to not slow down experimentation. But also make sure that it's being done safely, realizing that we don't have pristine answers to a lot of questions.

The technology is evolving, the regulatory landscape is evolving, depends on industry, it depends on geography. But we've got to operate in this environment. You can't wait for everything to be fully adjudicated. So, how do you safely experiment? And that's the way we've been going about it.

I really do think this notion that we talk about at Deloitte around "humans with machines" is a critical piece of governance as well because I think there are skeptics who do believe that this is not about augmentation. And I firmly believe that there is a great uplift to be had for, you know, many of our professionals through smart use cases.

**Hanish Patel:**

So let's stick with that train, exactly that force multiplier, what you've just mentioned here, Irfan. As organizations continue to adopt AI, and they're just starting on that journey or are somewhat mature on that continuum, what are the implications of the workforce? Because that continues to come up in terms of what does that mean for the future of *any* profession, *any* role? How is that impact on the workforce?

**Irfan Saif:**

I mean, to me, there's a lot of upside. There's some unknowns, but I know I've seen through my career various capabilities, automation, etc. that's come about that's helped make various activities, various tasks, easier.

I'm going to go way back and completely date myself, but when I started out writing code, you were using a basic code editor and everything was on a command line. And you had to automate everything you did through scripts. You know, over time, you saw things like commercial integrated development environments and these things have just evolved over time.

And when I look at engineers writing code today, I'm amazed by the efficiencies, the easy access to code libraries, the ease with which you can, write, test, deploy. It's incredible. And I think that that's only going to evolve over time.

So, to me, it's a very rudimentary example but it's one example of how people's work is going to change, and then combine that with a multigenerational workforce with those coming in now, off campus, about what kinds of augmentation they're going to have across various job expectations in the workforce.

So look at it as a strong impetus for change. We've got an opportunity—as not only for my case, and I'm sure for Fletcher's as well—not only for the internal IT but empowering the platform for what we do for our clients. And that's really compelling because that in and of itself is a different kind of force multiplier.

**Hanish Patel:**

Fletcher, anything from your side in terms of impact that you're seeing or implications to the workforce?

**Fletcher Previn:**

You know, the thing that everybody wants more of and never has enough of is more time. The time to grow your skills, time to learn, time to be curious, time to engage in deep individual work. There's a chronic shortage of time in the workday.

And so, AI, I think, offers a real opportunity to return more time to people. And in our role as CIO and in IT, to a large degree, we're in the business of meeting unmet employee and customer needs. So, to do that job well, as the designers of the future of work, requires deep empathy, which is not something that can be outsourced to AI. But we can use AI to help us unlock everyone's creativity and achieve their ultimate human potential as much as possible. And hopefully do the best work of their lives.

So what is a good user experience in an AI world? It's going to be a different set of skills and a different approach to things.

But it's critical that we get that right as the designers of these systems.

And so, we spend a lot of time, in my role, thinking about IT as a driver of culture change and that your culture is really the only unique thing that you have. And that doesn't change because of AI. It actually becomes even more critical to make sure that we're being thoughtful about how these technologies are leveraged and that the employee experience is really more important than ever.

And that, day to day, so much of work that we do and the decisions that we make can feel tactical.

But taken in aggregate, we are really having as much or greater an impact on the employee experience as HR or any other part of the company. And really playing a very significant role in answering the question, *What is it like to be part of this organization? And are we creating an environment where talented people want to work and enhancing our ability to attract and retain top talent?*

And AI can help us with that and take friction out of the environment and reduce the things that are dissatisfiers and return more time to employees to do the things that they were hired to do and feel fulfilled doing.

**Hanish Patel:**

I'd love if you could share a little bit more around what you're doing at Cisco, what we're doing at Deloitte, and how we're building AI into our offerings that we're taking out to our clients and how that's being permeated through the products or the service offerings we're providing.

**Fletcher Previn:**

Sure. So, from a product perspective, Cisco is really embedding AI into the entire breadth of the portfolio, whether it's security, collaboration, networking, observability. And I touched on it earlier: So much of security is having consistency in a unified policy engine, being able to interpret large amounts of data coming

from different places and make sense of it in real time. And then get help and being able to execute some change quickly in the environment to respond to threats or changing needs or the ability to move traffic around.

And you can imagine if some threats started to pop up that something like an AI assistant being able to say, "I'm detecting this going on, and I can very quickly make a change to the environment to defend against it and go and touch a thousand firewalls in one second. Would you like me to do that?" The value of that is hard to overstate.

In the collaboration portfolio, you have things like I mentioned around noise cancellation or being able to summarize what happened in a meeting that either you got up and stepped away for a while or you weren't able to join and later on come back and just get a quick summary of what happened. But also using AI to build out, you know, we have these large language models, but Cisco is building body language models, right? So much of what happens in a meeting in real life is nonverbal. And having things be able to detect body language and, you know, "So-and-so got up and walked away, and this was happening in the meeting at that time."

And then in the networking space, using AI to be able to help people use natural language and express their intent of the outcome they want even if they don't necessarily know the commands to do something—that the ease of use that that can bring to the products. And then on the customer service and support side of the house, just making it easier to understand what you're consuming from Cisco simplifying the licensing experience, simplifying the support experience. All of that is work that's underway to just make it easier to use our products, to buy our products, and to get support from our products. And AI is really front and center across all of those efforts.

**Irfan Saif:**

Yeah, Hanish, for us, obviously, the solution set is a little bit different. And so, I would

put it in a couple of buckets. One is there's platforms and capabilities that we use to deliver to clients.

And so if you think about things like in the tax business or in the audit business, how does a tax professional go about the end-to-end approach to doing a corporate tax return? How does someone go about the process of delivering an audit? Those platforms are starting to increasingly have various AI capabilities built into them to augment the tooling that a professional uses in their responsibilities.

Then you think about capabilities that we have where clients ask us to help them solve a problem. It can be anything from help us drive an AI program and do all the kinds of things that Fletcher and I talked about today. How do you set your data up? How do you pick the right spots to experiment? How do you build business cases? How do you actually go about driving organizational change management adoption? We help in all of those regards as well. And that's a little bit more bespoke.

So there's a lot of methodology, there's a lot of tools and techniques and capabilities that we can bring to bear and use in helping an organization stand up a program like that and run it.

And then in the middle of this, how do we go about the business of client service—if we're implementing the solution for a client needing to write code, if we are doing a managed solution for our clients around cyber to do threat detection and response. Things of that nature, we've got our own models, we've got an integration of commercial tooling. There's an AI overlay on that to help us drive client outcomes faster. So there's multiple dimensions outside of our internal productivity uplift as IT professionals to serve our clients in different ways.

**Hanish Patel:**

I love that. Again (laughs), my mind was racing as both of you were talking about

the way that you're embedding AI across various offerings in that. And probably the best way that I can sum it up after kind of hearing what both of you have talked about as your roles as respective CIOs is, I mean, it's certainly clear that organizations need to center their AI strategies and programs on creating value.

And both of you talked about some of those use cases as examples of where that value is being created around efficiency, returning time to people, improving the experience. It's really exciting to kind of hear about those applications of AI in the TMT industry, and actually beyond, and how it will evolve the workforce, as we mentioned, and really kind of reimagine how work is performed by us all.

But at the same time, it certainly is not lost on me that it's critical to develop risk, the governance structures, that you both talked to, to ensure frankly that AI has been implemented in a trustworthy manner. So with that, I do want to thank you, Fletcher and Irfan, for joining me today and really having an inspiring discussion around AI.

And I wish could have even more time because I'm sure there's so much more that we could dig into in terms of just opportunities that lay ahead, but also some of the guardrails that need to be in place to really kind of implement AI in a way that we all want it to be implemented. But I can't wait to see what the future holds just thinking about what both of you described. And to all our listeners, until next time, happy listening.

Thanks for listening to *User Friendly*. To subscribe or listen to more episodes, search for Deloitte User Friendly in your favorite podcatcher or find us online at [userfriendly.deloitte.com](https://userfriendly.deloitte.com).

If you like what you hear, please rate us on Apple podcast and leave a review.

This podcast is produced by Deloitte. The views and opinions expressed by podcast speakers and guests are solely their own and do not reflect the opinions of Deloitte. This podcast provides general and educational information only and is not intended to constitute advice or services of any kind. For additional information about Deloitte, go to [Deloitte.com/us/about](https://deloitte.com/us/about).

Explore more episodes of  
*User Friendly* at:

[userfriendly.deloitte.com](https://userfriendly.deloitte.com)

**About Deloitte**

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States, and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see [www.deloitte.com/about](https://www.deloitte.com/about) to learn more about our global network of member firms.