

AWS re:Invent 2025 - A leader's guide to AI strategy and implementation (SNR305)

62 MIN · YOUTUBE · [HTTPS://WWW.YOUTUBE.COM/WATCH?V=_TV1SP9JJZI](https://www.youtube.com/watch?v=_TV1SP9JJZI)

https://www.youtube.com/watch?v=_tv1sP9jJzI

SUMMARY

Helena Yin Koeppl, an executive in residence at AWS, discusses the transformative potential of agentic AI systems during the 2025 re:Invent conference. She emphasizes the need for organizations to rethink their AI strategies, focusing on identifying opportunities for innovation and productivity gains rather than merely automating existing processes.

- The evolution of AI from generative assistants to agentic systems is expected to significantly enhance productivity across industries.*
- Key productivity bottlenecks that agentic AI can address include unstructured knowledge management, personalized customer experiences, decision latency in complex operations, and faster product cycles.*
- Organizations often fail to realize ROI from AI projects due to outdated strategies and a lack of focus on business outcomes.*
- The RiPPLE framework encourages leaders to redefine their AI strategies by identifying real business problems, aligning incentives, and reimagining processes.*
- The BREAK framework helps organizations challenge existing processes and identify opportunities for AI integration by examining blind spots, constraints, and hidden costs.*
- Successful AI implementation requires a shift in mindset from task definition to objective setting, emphasizing the role of human-AI collaboration.*
- Establishing a robust AI organization involves creating a hub-and-spoke model to facilitate innovation and scaling of AI initiatives.*
- Continuous data improvement and modular architecture are essential for building a sustainable AI foundation that adapts to evolving technology.*

- Good morning. Welcome to your first session of 2025 re:Invent and the first senior leaders track session at re:Invent ever. So it is 9:00 AM, Monday morning in Las Vegas. You must really, really love AI strategy. (group laughing) So let's get started. My name is Helena Yin Koeppl.

I am executive in residence for AWS. So I'm a member of a small group of ex-senior leaders who had run and led transformation in our past life as AWS customers. Now we join AWS to share experiences with people like you. As for myself before joining a... Oh. Before joining now I hear it.

Before joining AWS last year, I led 26 years of data and AI transformation at four Fortune 500 companies in global roles. So today's session, a lot came from my own experiences, but also my past two years of talking to hundred of customers on their AI journey. So how much have happened in the past three years?

We have gone through from generative AI assistance to generative AI agents and now agentic AI systems. And as any technology revolution, we are foreseeing multiplication of productivity gains. So as you can see from steam engine to electricity, to moving to the cloud and the multiplier is getting bigger and bigger and the time it takes to get there is getting shorter and shorter.

What about AI? How much do we see productivity gains should be and how long does it take? And we do hope it can be hundred folds and it can take the shortest of time. Because generative AI and agentic AI should unleash many, many productivity, huge bottlenecks. For example, these are the five productivity bottlenecks the agentic era should finally unleash.

Number one, unstructured knowledge. We have gone through many, many eras of sorting through data, labeling them, a lot of them are structured data. And finally we have unstructured data in PDF files, in emails and critical know-hows in PDFs and people. Answers vary by who you ask and all of this now can have a single sighted front door via agentic AI.

To getting through that institutional knowledge and reduce time to expedite and narrow compliance risk. As an example, one of the AWS customers Moody's, it is a risk rating agency and it takes a lot of time to actually run through the analysis and huge amount of documents to analyze risk rating. What we help was to build a POC with them and reducing basically using multi-agent workflow to reduce that risk assessment analysis from one week to one hour.

Let's calculate the productivity gain here. Number two, we know with agentic AI, one of the key things it can unleash and what we have been talking about for years and years was basically that segment of one customer journey. So we've been talking about personalization, we've been talking about tracking true customer's preferences across platforms and finally with agentic AI's ability of acting and we are able to do so.

As an example, anybody here who is not a Rufus user? Very few, almost none, a few. Rufus is one of the key examples of agentic AI powered and generative AI powered shopping assistant. And because it can aggregate all of those reviews and product information, personal habits and person who bought this also about that and all of this information together, it's really make it much, much quicker actually 4.5 times quicker during prime day for shopping and maintaining low latency.

And we are foreseeing huge amount of value coming out from Rufus creation for amazon.com. And we have seen that not only in shopping in consumer base but also in highly regulated environment like banking. NatWest boosts their click-through rate by 900% and 2 million more higher interest savings account publications by using personalized experiences and recommendations.

So number three, decision latency in complex operations. Huge, huge amount of data, again, unstructured sensor data and all can be aggregated together and leveraging agentic AI to make sense of it. Example, again, coming from Amazon, we have a agentic and generative AI assistant called DeepFleet managing the 1 billion robots we're using in our warehouses and really helping the automation becomes extreme and from handling

troubleshooting to route directing and dynamically deciding what is the best way of passing the right products into packaging and shipping out.

Again, a lot of the agentic AI usage we have seen is in service. In service management in the past and I have experienced that more than 20 years ago. If you want a great customer support, you go follow the sun. You have three teams based in three different continents covering 24 by seven. Now you don't have to do that anymore. DoorDash for the 7 million dashers have that agentic AI system to support all of their inquiries from how much tariff does this product cost today to troubleshoot, again to question and answering.

So that is 3 million savings in just supporting the dashes. Last but not least, faster product cycles. So built-in, learning fast and really sorting through that huge knowledge base. And we see that in healthcare and life science, one of the industries I worked many years in and very, very difficult because normally there's huge amount of unstructured data. One of the examples we have is in Genentech, which is Roche Group company and we help their research organization to automate 43,000 hours of manual work augmenting their researchers' work and that's equal to five years of time saved.

So all of these great examples yet when you're reading some of the headlines, 95% of agentic AI projects seem to fail to deliver measurable ROI. And if you read true and read in the details, what it's really saying is those 5% of the companies who manage to really gain value and actually millions, millions of incremental value in productivities and in business, new business models, what did they do differently is really transforming their strategy and business process.

So what I'm talking about today is stop steering today's AI challenge with yesterday's playbook. There are some of the examples of yesterday's playbook. I'm sorry if the mic is a bit, I mean louder and softer. So yesterday's playbook. As an example, number one, when we are talking about such piece of fascinating technology, very often we ask ourselves what can I do about it?

Instead of starting with the problem to solve and we are actually going into and using a solution, searching for a problem. That's number one. Number two, we don't really rethink about business process. We look at the exact process as before and just find steps of automating and only having incremental gains and sometimes even slowing it down. And talking about the experiences, still we are breaking the experiences down into silos.

Marketing does communication, sales lead generation and none of them actually are working together towards actually that same customer. And there are more in terms of the how and we've been talking about yes AI can unleash productivity gains, but what we do is not to retrain and redesign roles but actually replacing people. This is actually not the right approach and especially not the productivity gain truly the right thing to do.

We have been in terms of organizationally centralized AI community and we will talk specifically about that and of course data which is the new oil but especially to oil that engine of AI and we are still not setting the data up right. So many organizations who come to us who I have experienced before actually have come to this inflection point.

We have been doing a lot of experimentation in the past three years, many, many POCs but we have seen an even ROI and actually scattered of AI spending by function. And really the matrix is how many AI projects am I having rather than the outcomes? And now leaders come to us who are asking the right question which is are we really backing the right top three value pools?

Will this give us the defendable competitive advantage? How do we turn the incremental ones truly into compounding gains? And this is today what I'm going to talk about. How do we do that? How do we actually have an AI strategy and implement to deliver that? So let's think about the true flywheel of AI innovation.

Yes it's great, we have a lot of experimentation and it's great to learn about the new technology. We actually should actually go back to improve one step up to say where am I going to truly invest AI and where is the strategy point and the business strategy that I can support really makes a difference. Start with the opportunity, not the solution. And then of course deliver it but quickly and truly show value and ROI and all the while at the same time, building the right foundation.

So that is the new AI strategy playbook. There are six chapters I'm going to talk about, how do you actually define a AI north star? Starting from the AI strategy that mirroring your business strategy. How do you reimagine business process and personalized experience? Then how do you actually build that motorway not just about technology or platforms but truly about people, both in terms of everyone who are leveraging AI in your company but also your AI organization.

So number one, let's start with AI North Star. The successful AI transformation, the biggest mistake organization very often start is let's have AI strategy and let's look into how do I leverage AI. The right question we should be asking is what is your business strategy and how do I, and where do I translate the business strategy into AI strategy?

I've talked to organizations who wants to translate business strategy into AI strategy. What they do is listing all of their seven year business strategy and looking into and wondering how do I do this? How am I going to translate that very high level vision into a very detailed AI roadmap? The number one question is you want to ask yourself where do you have opportunity to automate?

Where do you have opportunity to innovate, creating new product new services and where do you actually have that disruptive advantage? So keep that three questions in mind and still think about how am I going to translate? And the answer is we should ask different questions. Not the same question, not the question that we've been asking for the past 50 years, but I'm going to introduce a new framework which is called RiPPLE.

These are a new mental model which help you to ask new questions, to translate, to define the must-haves on where your AI can play a role to create that competitive advantage as comparing to nice to haves. Okay, so. So number one.

R, R means rational pause. Let's look into where is truly making a difference in your business, where your and actually the uncomfortable truth, where are you bleeding customers and where you are losing revenue to your competitors. So the existing market defended and protect your competitive advantage. And how do we ask different questions?

Here are the different questions that you can ask. So think about yourself comparing to competitors. Where are the nimbler competitors serving customer faster than us? What are the high volume decisions that costing us a lot of revenue through delays or errors? And that's your opportunity. What routine work prevents our best people from winning new businesses and where do we lack real time visibility into customers or competitor moves?

So that's number one. That helps you to define where is this problem space that AI could play a major role. Number two, incentive mapping. Incentive mapping is all about realizing within these organizations, we often talk about organizational silos and how our organizational silo created. It's not just because we are in different departments or functions or even in SAP teams, it's because our incentive systems are not aligned.

So we should actually ask different questions to discover that. So you have conflicting incentives across different departments. Where are departments measured on matrix that work against each other? Where do department silos prevent us from serving customer better? And what handle bottlenecks and data process are preventable? And if you're discovering those opportunity space and that's where we go into the second sector which is re-imagining business process.

We have another framework to introduce. Number three, perspective divergency. This is quite interesting. We do know that we do have problems that we try to avoid because it's too complex, because it's always, always been there and we have tried to solve it for years and we couldn't solve it.

But right now ask yourself differently, what if you have a AI native competitor trying to get into space? A place I worked before and I led AI product innovation before and it's about it's legal space. And legal space has many, many years of labeled data and we truly think that we have the competitive advantage.

However, when we are, and there has been quite a few successful new entries with AI first approach. At first they were small but because they have the valuable insights very quickly generated with AI and generative ai, at one point they're truly becoming threatening competitors. So where does it make them nimble?

Where does it make them competitive and how can we define and how can we actually leverage our own, for example, proprietary data to actually become more successful than them by leveraging AI? So here are some of the questions you can ask yourself to identify that area you want to (indistinct) AI and you want to try with experimentation and moving into production.

Okay, so you've discovered some opportunities and a next natural actually important question you want to ask

yourself is not should we actually use AI here? Not just could we? And that's very, very important. That is the contrarian truth, right? So you can ask your question as the process come into question, what if it is already the best fit for the context?

If we introduce AI into this process, could it introduce new risks? What are the mitigation? And especially when we have perceived bottlenecks, is it real or is it actually misunderstood? So all of this, think about your business, think about in your area that not could we but should we as well. The opposite.

And very important question you want to ask yourself is also, if we build AI, if we truly invest into resources, if we truly invest into this problem area with actually transforming business process, can we build sustainable competitive modes through AI? Maybe it is, maybe we can truly create that data that are unique that competitors cannot and will never have access to and we can embed AI into workflows, continuously learn and record that not only in institutional knowledge that we've been talking about which are yearly, monthly, but actually daily, second of the operational knowledge and continuously training AI, learning on your business, customize according to your business and your operations and where's your knowledge, where's your proprietary data and it's eventually becoming true that nobody else can copy and that is your mode that your is your continuous competitive advantage.

Last but not least, how quickly can you implement it? That is so important, the velocity of implementing it. And if you can implement it within two months, within six months you might be able to create 18 months of advantage that nobody else can follow. And that is where you should invest. And if you identify that opportunity, put sponsorship and develop skills to execute not only for prototype but truly moving it into production to continuously keeping that gap between you and your competitor.

So with this RIPPLE framework, the opportunity filter workshop is your next step. So use that to identify these opportunities but score each of the opportunities by these four factors. Impact on business outcomes. By the way, don't forget technical feasibility and that would also impact how quickly you can leverage it and implement it. Creating (indistinct) with it. Strategic alignment and innovation velocity, we talked about the last point.

And then select the must have to move forward. So that is your map. Let's talk quickly about ROI. And we have been talking about we need to identify ROI for AI, not only just for the past three years but ever since that I started working on AI. You need to, one key learning I've learned in past many years is you need to set up observation in operation.

So not post facts that will be too late. You won't be able to separate successfully what is AI and what is not. So there are five steps you can do. You need to select the right value stream. It really have result into productivity differentiation or new revenue streams. You need to embed AI into that flow but wire the data. You need to have the entire data flow identified and separating AI and long AI and you need to actually take those usages not only just AI, no AI, how long, which part, what's the percentage and run and compare and scale.

And there are some AI ROI discovery methodologies to put my analysts hat on a little bit. So there's when you

are running pilot and you are able to isolate where you're running pilot and a place where have similar environment and similar activities but without AI. So what you can't do is test and control. So the test is where you have AI like AP testing and the control area is where you don't and you can very quickly identify how much incremental you can create.

After post-launch because you have embedded data flow, you are able to do AI attribution analysis. And when you are ongoing what you can do is actually every single touch point can be identified with the AI touch and how much AI touched to continuously optimize it. So that's number one. Business strategy to AI strategy. That's your AI north star.

Now let's look at a couple of deep dive examples. Number one is we talk about business processes. If you really keep it completely intact and only automate the separate part of it, what we can see and what we have seen from many, many interviews with customers is you don't actually see a lot of AI value creation.

What we need to do is to shift from the separate keeping things intact, completely sequential implementation of AI into reimagining the complete business process, focusing on what is the outcome you want to get and rethinking about it. Several organizations I've talked to have come back to say, "We tried that. We did a workshop, we talk about the outcome and everybody went away to think about the opportunities and redesign it and six weeks later they came back with Excel sheets."

So how do we actually do that? And what I will suggest is again ask different questions starting with the sequential that you have right now, which is using. I have a example of due diligence. Due diligence, just like any complex business process today is pretty sequential, right? So you collect the document, you have the initial review and then you flag the issues and then the issue go through legal review, financial impact review and eventually you have the risk assessment and you basically can say, okay, can I deal with it or not and is it worth it?

And do we have to go back to the previous session and handle the exceptions there? And it is time consuming and due diligence takes six to nine months in average and when you are actually should be moving forward with that business value creation. So another thing to think about is, and that's why we should not start with technology.

Very often when we are transforming business process, it's not just one piece of technology, it's not just agentic AA, it's not just generative AI, it's not just okay let's use AI assistant and make the process a little bit better. No, if you want to reimagine the process as often said, don't start with the technology and this is a key piece of it because let's identify exactly what we want to do. What is the most important thing and most time consuming thing is actually to discover issues in due diligence and then think about way of mitigating issues.

When you are thinking about that and you actually reimagine this process into a non-sequential process. So number one is I need to inject all the information I have, all the documents and then I need to do semantic analysis to identify that key clause that might be the issue. So that's number one, right? So you are already

aggregating several of the components together, not sequentially. And you can use generative AI for it.

And once the anomaly is identified and you can use analytical AI to do for example compare it with your ERP data, financial impact analysis and do the risk scoring. So that's cross-referencing and valuation. And additionally let's say the issue you discover is a vendor which actually have a clause of terminating the contract if any (indistinct) happens.

So what do I do here? You actually can parallel process and going into several agentic AI coordinated by a agentic AI coordinator. And so using legal processing agents and using financial analysis agent and actually human oversight all together. So that's your agentic routine which save time by not doing it sequentially and identify what could be the most optimum outcome and solution and then sending it to ramification.

And by the way the ramification itself can use agentic AI to draft waiver proposal for example. And human at the same time in the loop did decide the action. So all of this is to reimagine and not to think about sequentially. So we talk about the how is important, how do we do this? How do we take a sequential business process and ask different questions to say how do we completely reimagine this?

And I have a framework which is called BREAK and these are the BREAKS. Blind spot scan, reframe constraints, economic dissection, assumption audit and Kaizen, the happy path. So what is blind sport scan? When you have a business process which you have been using for years, it's very hard to just tell yourself, reimagine it.

One method is to ask why and ask why five times, and sometimes some of the whys might be why not. So as an example, this is a actual example. The step currently is manual approval for order over \$1,000. Why? Because we need to prevent fraud, it's a big amount. But why? Because big amount are risky. But why?

Because we cannot verify customer intent with that amount. Why? Because we don't have real time data. Hey, the question has already been flipped and it's no longer a compliance question, it's actually a data question. And the last why is because our systems are not interoperatable. So that five whys help you to identify again the business process problem space where you might want to dig deeper into.

That's your blind spot, which without asking the whys you would just go ahead and continues to do what you were doing before. Second, we very often have that constraints. And another idea of reframing that constraint is to challenge your process limitations. Ask yourself what if I have zero latency?

What could I achieve with that? What if there's zero touch? Then what could I have achieved that? With for example, AWS, we ask ourself what if we have zero ETL? What can we achieve that? And so the very often a new way of thinking new invention and new reimaging of business process starts with ask you a question.

This is a limit that I go to and how am I going to achieve with that? And then ask how do I achieve that? And

that's how breakthrough happens. Economic dissection. There are yes today that you might think this process is what we can do, but what we don't think about is what are actually the costs, the actual costs and value destroying elements if I don't make a change.

So for example, very often we don't think about the hidden costs, when we are thinking about "Hey this is as much as I can do in terms of delivering for example ideas of just work out". Is thinking about it's normal to stand in the queue waiting in line at supermarket but the hidden cost of time consumption and hidden cost of people being frustrated actually put the things down and never buy it.

So those are the hidden costs and value lost that you have with your customers. Think about in your process where are they and how can we solve it? We should have a assumption of it and it's about flipping the sacred cows. And sacred cows are very often again being used as there are risks, there are compliance issues, but is it truly?

Every single approval process sometimes it might be compliance issue but it might be just organizational habits. Ask yourself this approval processes, are they truly compliance needed or are we just used to it? And we have to go through the hierarchy of the organization and also manual quality checks. I only trust human to do this step. Could it be prevented? Can it be prevented?

And which sequential dependencies are actually requirements and just because historically we have always been doing that. The last BREAK, the last one is K, Kaizen, the "happy path." In process engineering we very often talk about the happy path, the path which is the easiest and which is 80% of the time you go through that and that's where you put it in.

But the new way of discovering the happy path, especially when you have agentic AI, when you actually can parallel process many, many different scenarios and find the right one is actually that thing just happen. That thing just happen for a while. The picture you're seeing right here, this happens in Ohio State University. They want to build paths for the students to get to their lecture halls and instead of building it, they actually let students just walk across the grass and identify the most useful paths and then eventually build the routes according to the the foot paths.

So that is observing the actual behavior and because agentic AI has the ability of logging everything, reflecting on it and identify patterns with memories, you can identify the happy paths and formalize it. So kaizen by the way in case you're not familiar with is Japanese philosophy. It is about, I mean some of the enterprises organizations actually use it as well in their business strategy.

So it's about continuous step by step small improvement to get to a bigger value. But that frequent changes is informed by actually what has happened in actual work and experiences. And the identical AI Kaizen at scale, as I mentioned, there's continuous observation and continuous recording and continuous reflection and learning loop with agentic AI, and there's policy refinement according to the agent's observation.

"Hey this path and with this analyst steps actually takes 20% shorter every single time and we should reroute to that." So that is the policy refinement which is continuously done and is traceable. And parallel experimentation, instead of AB testing, you have A, B, C, D, E, F, G and many, many thousands going in parallel and finding the right path and finding actually truly the experience people prefer and the feedback is always integrated including actually the human managers of the agentic AI system.

So with the BREAKs, the next one, when you want to implement, how do you do from process mapping and identify your today's process. It might be sequential, you might already have some digitized or automated part, but then use the BREAK analysis to truly challenge your limits and the the holy cows. And then of all the bottlenecks you identified, prioritizing which one makes truly a difference and what are the dependencies you have on them?

And then map it to AI solutions as I mentioned before and what you have seen in the due diligence re-imagining of the processes and what AI solution could we use. And eventually have the implementation roadmap. That way instead of just saying, let's reimagine the business process, you have a bit of structured approach to help you to get to the opportunities.

An example that I already been working with a few customers and is clinical trial process. Clinical trial is when we bring a drug to the market, it takes roughly 10 years and billions and billions of dollars normally because you have to go through phase one safety, phase two efficacy and phase three in a bigger population. And many, many people believe that all of this needs to happen sequentially, because of course safety first and then small population tasks for the efficacy and bigger populations.

What we don't challenge ourself is can some of these processes happen in parallel. When you are preparing documents, when you are referring to what is in the market already, when you're referring to the literatures, many of this actually can happen at the same time and through trial and error as well. So using similar mindset we can rethink about how do we redesign the experience with AI because it is often very, very similar problem, which is the silo problem.

When you're rethinking about it, the way of doing it is put a human at the center instead of to say marketing needs to do communication and sales lead generation. This is the person that you need to actually serve. What is her need and what does she want during the entire process? Agentic AI can help you with that because it has reasoning, planning, action and orchestration all done at the background similarly.

So let's use a actual example. Returning products can be quite frustrating for customers. So when you are doing that, I want to return this and at the behind the scenes AI can do it in parallel, assessing the possibility of, I mean by the way the purchasing history is being leveraged to identify hey, does this customer always always return the products she bought?

Is it valid? And what has been done before and what are the warranty? Can she return it? Is it within 30 days after she purchase it? And reason it at the same time and getting all of the related information needed and acts.

So if it's valid, if it's still within her right and let's process it, let's update also the inventory immediately and let's arrange the shipping and let's notify the customers.

All of this, imagine if it's our original process of processing the return is going through several departments and all of this actually are done without going through that silo and with actually that person with this particular needs at the center. And you need to think about how are the possibilities of dealing with it more proactively.

Okay, we solve a problem with returning the product but eventually if you put the human at the center, these possibilities of thinking about proactively of designing her experiences, can I propose a new product if it's a customer? Can I actually deal with the project guideline before the person notice it and let her know, hey, it is approaching in three days, what are you doing can be accelerated?

And if you are dealing with partners, can you actually identify the challenge from all the arrow locks and the repeated problems that you might have or the partner might have about this issue? So to implement it, an idea is again during a workshop and mapping that event you want to solve and identify the silos and also the entire orchestration.

Think about how can I use it to completely transform the experience and without thinking about sequential steps that I need to go through if it's human team handling them. Example is actually again in Amazon, the entire Amazon organization moved into an new agentic AI solution called A to Z.

So instead of when I need a vacation, I need to go through HR policies, I need to ask from my manager and all this information I provided, I just type in with the interface, I need a vacation of five days and all of this needed information and are truly send it to me. Okay, so this is a system according to the policy, you are in Switzerland and this is what you can do and you have remaining 15 days.

Yes you can take five days of vacation right now. So if we put for example in this case employee as the center there is HR manager, payroll, IT, all of this departmental silos are put aside and truly think about what needed to be achieved with this human in the center and every processes and tools can be leveraged.

Now let's talk about how do we actually work together with AI for all of your employees. First of all, not everything should be given to AI, whether it's agent or generative AI. There are agentic AI big problem spaces that truly deserve to use a launch agentic AI. For example, you have a clear objective because that's what AI needs to achieve.

You have sufficient data truly available and you have many decisions but it is possible to reverse. In Amazon we talk about one way door, two-way doors and two way door decisions can be reversed back and those you can give to AI. And you can retain human control and you should retain human control in more strategic decisions, relationship sensitive decisions and novel situations.

And let's not forget and high stake decisions and one way door and irreversible decisions. And we need to think about also we need to think about human AI manager. It is no longer a task definer. So what is a task definer? Run this report. You define what needs to be done, but truly now when we're working with AI, especially agentic AI, we need to give AI the objective.

I want to achieve this, please do the payment on time, please help this customer and according to the policies, return this product with the right way within the time limit. So that paradigm shift of task dedication to objective setting is our truly true change of the mindset. So human's role as AI manager requires three competency, objective setting, performance monitoring and strategic intervention.

So the details of objective setting for example, what is the primary goal you want AI or agentic AI to achieve? And you set the goal but you also need to set what is the success criteria. For example, CSAT scores, for example, resolution time for example, fast contact resolution percentage. So all of this is potentially success matrix. What are the constraints? What are escalation triggers?

When do you actually need to keep human in the loop? And how often do we review the performance? Weekly analysis, daily, monthly? And leveraging AI's memory. And very importantly, AI manager skillsets are different from human manager skillsets. And one true important skill is actually learning agility.

Human managers. When you're working with AI agents and we always say that it's actually whoever knows how to best leverage AI is going to become the high performer in the future, the best people you have. What he needs to do is actually be very, very good as at pattern recognition, as in how exactly is the best way of leveraging AI and combining with my team members to achieve the objective the best way, the best route identification we talked about.

And also knows intuitively if after working with AI for a while, what are actually the contextual interventions human needs to very often do and do not let AI handle it. So that human AI augmentation happens when human intuitively know when and how best to leverage AI. And those are the four key skillsets they need to develop and eventually is truly on the job by working together with AI system.

Now let's talk about AI organization, the organization which will enable you to launch AI. We have the leader, we have the organizational structure, we have the innovation ecosystem, and we have the talent system. Number one, the leader. Anyone here who has the title of CAIO?

It is truly becoming a trend that organization are hiring CAIOs, people who they believe should be leading the AI strategy and telling the organization that how should we best leverage AI. That's great, but sometimes we do encounter one leader cannot scale because a lot of the organization, especially the organization at the top, the board, the C-suites are very often also new in this.

They don't know about what AI can or cannot do and they don't know what is the right way of making their

organization truly know how to leverage AI. So one key thing about CAIO is actually to educate, to orchestrate and distributed AI fluency and ownership across the entire C-suite. And we see that is the most effective way of CAIO leading the organization.

One person cannot scale AI strategy and cannot make it successful for the entire company. And then the AI organization itself, I have done and being tasked and done AI or data and AI transformation for four global organizations. And the model I often use is hub and spoke. Yes, we do need many things in the hub in order to scale quickly for the entire organization, but we also do need spokes into the other business functions and into the other departments to help people learn and to help actually people who are facing that business problems and opportunities daily to identify the opportunities.

But how big is the hub and how big are the spokes and how flexible or permanent are the spokes are the differences. So I mean I have led AI research heavy models where the organization is truly creating something and doing research on that, doing cutting edge research on that, doing something that which are the industry's holy grail problem. If you were to establish an AI organization, which is research heavy, by all means, AI center for expertise needs to be big and very often the spokes into the other functions are temporary.

Going there solving one particular problem and moving back into the center which is called squats. And then another model of hub and spoke is small hub and bigger and more permanent spokes. And those are very often with AI heavy model. When I was working for global life science companies where that commercial organization's hub was actually at the center, but many, many of the markets have spoke which are permanently there.

But most important thing is how are you going to discover and scale that AI innovation and how do you actually multiply that. And that is what I call AI incubator model. Whether you are hub heavy or spoke heavy and having that ring, it's really important. Because very often successful ideation and pilot of AI use cases happen in the spoke.

But when people discover that and even it's successful, they have no obligation to scale it. And in order to actually mobilize that and motivate people to actually moving it and finding opportunities of reapplying it in other organizations and then feeding back to the hub to actually create permanent products to enable that fast scale is that ring and that's your innovation ecosystem and that is the incubator.

Whether it belong to hub or belong to spokes, it doesn't matter. We very often say this is the true two piece team with multifunctional team, business subject matter experts, but importantly also AI experts. Having the right talents, having a talents which have that business acumen on the right hand side and problem solving skills, but also AI functional skills, the left side and balanced left and right side is very, very important.

Whichever organization hub or spoke they are in. But what you can adjust are the balance. So some needs to be left brain heavy and you need a lot of in depth AI engineering or scientist skills or some needs to be balanced. Like for example, a very important role of AI translator is about problem space discovery and AI opportunity

identification.

So that left brain and right brain needs to be balanced. So from strategy to implementation takes 90 days and to think about that whether hub and spoke and which type of hub and spoke where your people today are and doing the assessment and identify that fluency program for your team and whether it is increasing their left brain part or right brain part is can be done following this 30, 60, 90 days roadmap.

Last but not least, data. We have been fixing the base for years. Who hasn't encountered some of these, data warehouses, data lakes, data lakehouses, data mesh and data fabric. What we need to think about, what we need to identify now is how can I actually build a living foundation?

Because your AI foundation should not be a cathedral and should not be keeping your head down two years and building it, but really having it as living city and lay the pipe and continuously develop because agentic AI unlike all of the others from the past do not need perfect data. It can take the data, it can improve on the data and it can actually continuously use data to both enrich the use case but also further develop your models.

And AI technology evolution is accelerating and as we know this is just some of the examples in the past three years. So what we need to do in terms of architecture is also keep it simple and modularized. And that's where, for example, solutions and platforms like AWS AgentCore can help you and keep the choices and do not log in and give you production grade of observability.

And we need to not only keep responsible AI and risk management in principle, but also operationalize it. Okay, so for the implementation, again for the 90 days, if you want to actually think about Agile foundation, these are some of the implementation roadmaps you can follow. Start with one use case, create that template and expand.

So these are my six chapters of AI strategy playbook and this are the summary for you to take back. So if you have, our time is up, but if you have any questions and feel free to approach me after the session or throughout the entire re:Invent and enjoy the rest of your re:Invent. You've been a great audience. Thank you very much.

(crowd clapping)