



EPAM Investor Day 2026
March 13, 2026

Mike Rowshandel:

Good morning everyone. Thank you for joining us today. I'm Mike Rowshandel, head of Investor Relations. Whether you're joining us live here in Boston or dialed in through the webcast, we appreciate you joining us today. We've been planning and preparing for this day for quite some time now, and I can tell you the energy in the backstage is buzzing. Our entire leadership team is here, eager to take you inside our story, what we've built over the past three decades, and more importantly, the how and why we're positioned to be successful in the AI era. But here's the thing. Saying we're positioned to win is easy. Showing you is what today is all about.

And that brings me to our theme for the day, AI made real. Through today's presentations, you'll hear real client testimonials, providing a deep sense of the unique value we can continue to deliver each and every day.

Before we begin, I would like to remind that today's presentation contains forward-looking statements, which are subject to risk and uncertainties. Please refer to the Safe Harbor statement in our presentation materials and SEC filings for a discussion of factors that could cause actual results to differ. We'll also reference certain non GAAP financial measures. Reconciliations to the most comparable GAAP measures can be found in the appendix of today's presentation.

Now let me walk you through what to expect over the next few hours. First, some context. Our last major update was nearly four years ago. And to say a lot has changed would be an understatement. The macro, geopolitics, competitive dynamics, AI disruption, the evolution of IT services, and EPAM itself. All look very different than back then. And that's why a key objective of today's presentation is to provide important clarity on where we are today and where we're headed over the next few years.

Let me quickly walk you through the agenda. The day's organized into two parts. In the first section, we'll provide an important update on our strategy, how we're transforming our go to-market motions, and then we'll dive into a key AI section where we'll talk about AI native engineering and AI native business transformation. In the second section, we'll focus on our engineering DNA, our AI talent, global delivery engine, and then we'll have an engaging panel discussion with several of our geographic leaders. We'll then dive into our financial imperatives and then finally close with a Q&A session.

Finally, for those that have joined us live here in Boston, we invite you to please stick around after our main presentation. There's an highly interactive hour of client demonstrations and industry-led tours. These client demos and tours should give you a real sense of the AI capabilities we're delivering today.

With that said, let's get kicked off with a quick video.

And with that, I would like to warmly welcome our chief executive officer and president, Balazs Fejes, or better known as FB. Thank you.

Balazs Fejes:

Mike, thank you very much. Good morning, good afternoon, good evening, everybody. Thanks for joining us here in Boston in our 2026 investor and analyst day. My name is Balazs Fejes, but just please call me FB. I'm not going to force you to learn how to pronounce Hungarian names. Mike probably spent two hours practicing how to pronounce it. You don't have to go through that. In the next probably 20 minutes, I would like to give you a strategic overview of EPAM, the market itself, how we are positioning ourselves to win in the AI native era.

But the most important thing for me is that you have four things to take away from here. We are reinventing ourselves as a global leader in the AI transformation services space. We are using leveraging industry's best engineering talent in the industry to solve our client's hairiest, most complex business and technology problems. We are strengthening our internal and client facing AI capabilities to capitalize on the global AI transformation, and we are executing a clear strategy to drive our next phase of profitable growth.

But before we start, I think we need to really address the elephant in the room. We are reading the same headlines, same Substacks, I think watching the same Instagram, TikTok, YouTube Shorts, or YouTube videos and even just today, there are some new news popping out from everywhere. It tells a story. It tells a story that AI capabilities are growing really, really fast. And this is true, but this is only one side of the story. It only talks about AI capabilities growth, but it's not talking about the adaption rate in our societies and in our enterprise.

The two is very different. Whereas AI capabilities are growing really fast, the adaption rate, the way people are changing how work gets done is growing much, much slower. There's a gap between those. And this gap is the opportunity for EPAM. EPAM is operating on the AI frontier. We transitioned into the AI frontier, and today we're going to show you how we've done that and how we planning to stay there, and how we're going to help our clients to catch up using the learnings what we gained in the last three years. This is the opportunity of our lifetime.

Just a month ago, we are talking most of you and updated your 2025 results, delivering almost a 5.46 billion revenue. This is our sixth consecutive quarterly revenue growth on reported basis, and we are really proud of it. We are delivering across 55 countries with 62,000 EPAMers, with 56,000 delivery professionals. And it took us 30 years to get here. We have been around, we've seen a lot. And in 2025, we really delivered this growth across all the industries, across all the different geographies with a wide and very distributed presence. We don't have real concentration on this, which we feel it's very important given the current economics and grant situation.

But let me remind you who we are. EPAM, we are a build and change organization. In the last 30 years, that's what we've done. We honed our engineering heritage to actually build solutions for our clients. We are builders. We are delivering results relentlessly to our clients, helping them to navigate technology, geopolitical, and economic changes with our hybrid teams. That's who we are. And that's, I think it's very important because we just entered the age of building. We are seeing that AI enables us to build new solutions, and that's our advantage, and that's our heritage. We are serving clients and diverse and global client lists across 11 industries. 345 of our clients are part of Forbes Global 2000, 64 out of 100 are part of S&P 500 and the Global 2000 in the same time. And the top 20 clients of ours on average have 13 years of tenure. We have deep relationships. 80 plus of our top 100 clients are executing AI native projects with us. Actually, we delivering for them new transformation projects. And the same time, we're winning new opportunities, we meaning new deals, and expanding our volley share.

We are positioned to harness the value of AI internally, and also to capitalize on the growth opportunities. And I think that's very important as a key takeaway for us. Already in 2025, our results are

benefited from AI. We delivered a very strong AI native and AI foundation momentum, which was built on five different foundations or pillars. We helped our client close the adaption gap with our scaled talent capabilities. We helped them optimize their delivery using AI SDLC, which we later on launched part of AI/Run. We help them modernize their legacy system using AI where we launched MF lens, which is a modernization toolkit. And we also help our clients adapt physics, AI, and robotics. And the same time, we're helping clients globally to roll out solvent AI, which is becoming more and more important in our increasingly more complex geopolitical space where we are operating in.

Them is very, very complex. And I'm sure all of you guys came here to understand how do we see TAM. So myself, nor EPAM, we don't have a crystal ball, so we just don't have that. So we decided to borrow one from Gartner, I think. So I'm going to use Gartner's crystal ball to try to explain to you where the market is growing. Market predicts, so Gartner predicts that the total market of IT services is going to grow to \$1.8 trillion by 2029, which is a CAGR of 5%. We are operating in a subsegment traditionally delivering solutions in business, consulting, technology consulting, application, implementation. This part of the segment is expected to go at 6.5.5% CAGR. It's a growing market and continues to grow.

Same time, we took another report also from Gartner, which presented a very different picture. This really talks about the AI market itself. They are predicting that the total market of AI by 2029 will grow to \$4.7 trillion. That includes all the GPUs, all the data center investments people need to make in order to make it work and the software and the AI services too. The AI services part which we are really looking is we compromise multiple sectors and we actually took just one slice of it, what we call AI services plus AI cybersecurity. It's a very fast growing sector. It's actually growing by double digit CAGR, or sometimes strong or very strong double digit CAGR till 2029.

Now what's important to take away is that Gartner's definition, what's AI services and our definition of AI native doesn't really match because they do include some parts, which is what we call AI station revenue. But still, very important takeaway, it's a fast growing segment of the market. Now, I'm not going to able to square off what's going to be replaced by AI or how much IT services is going to be impacted by itself, because nobody can, and we don't have the data for that. But I'm just using this as an influencer to demonstrate to you that it's a vast growing market, which we're trying to tackle.

On the other hand, I would like to really focus on why we are positioned ideally to win in this \$1.3 trillion opportunity, which we call AI services. EPAM has a client zero mentality. We spend three years building our capabilities, honing our capabilities, how to harness the power of AI on ourselves, which gives us credibility. We have an engineering heritage and in the age of building and actually applying AI, it's a very difficult thing to do and you need real engineering power, real engineering capabilities to make it really work. We understand how to manage talent, how to create the next generation talent, which is so important in the next couple of years.

We have deep industry expertise because without industry expertise, you don't know what to automate. You don't know what to change and how to really take advantage of AI. And the only thing you keep talking about is how to take cost out, and that has a limit. And we have longstanding client relationships, clients who trust us, and you're going to see demonstrations of that to actually experiment with them how to use and how to roll out AI using our expertise, what we gained in the last three years internally.

We have an aspiration. Our aspiration is we want to become the goal to partner for enterprises for AI transformation, which is built on three strategic pillars. Number one, we want to position and establish EPAM as a leading software engineering services provider. We want to transform ourselves to be an AI native organization, and we want to launch new AI native offerings, which we're going to talk about. The key enable us to make this happen is talent, skills, which we talked about, strategic partnership, extending strategic partnerships, which we just very recently entered a partnership with Cursor, which is

a very important part of the puzzle, domain and vertical expertise, and continuous investments into internal products, internal IP.

We have been accelerating our internal transformation. We are true to our values of being client zero. So we spent three years implementing and changing how to run our business, how to run recruitment, project staffing, talent management, how we can do management reporting and finance and legal using AI. We got some recognition due to that in best use of AI or the best competency and skills development using AI. We have been recognized for this effort.

And using all the knowledge what we gained in the last three years, back in autumn last year, we launched a codified go to-market strategy under the brand name AI/Run, which really addresses how to do AI native software engineering and how to do business transformation, which become an AI innovation-based business transformation. This consists of playbooks, blueprints, how to manage talent, and also tools, platforms behind it. This is based on real credible evidence, based on the three-year experiments, which we're doing on ourselves. And we're going to demonstrate it to you if you are in person in Boston with all the different shows around you, and also later today, we're going to actually show you how we implemented this tooling into our internal systems.

We're creating new AI native business models and services. These are net new services, net new revenue for APEM. These services are agentic intelligent operations, AI native experiences. Just a couple of months ago, we launched Empathy Lab in North America, which is our AI native services experiences launch and brand name under this. AI native agentic operations and agentic factories and agentic security.

We are doubling down on our growth drivers, talent, skills, and capabilities. Extending on our 30 years of heritage, Sandra and Alexa, we'll be updating you how we are creating the new talent, how we creating the new roadmap to actually create the new talent, and how we're sensing who has the capability to get there. We are verticalizing and actually depending on the industry experience. We are pushed our consultancy teams into our verticalized industries. We're building, continue to build out internal platforms and IT assets, and of course, strategic partnerships where we need to strengthen and we will double down our footprint.

I think if you were following us in the last years, you heard a lot about our telescope AI. We invested decades in developing an enterprise backbone, digital backbone, which allowed us to manage our organization, manage us through crisis, manage us through different disruption, and continue allowing us to deliver with high quality. Now, we actually put an agentic backbone on top of it, which allows us for teams and agents to interact with each other and actually take real time data to drive better decision making with a higher quality output.

Our leadership team has changed. We realigned our leadership structure around industries, brought in new members. You have the chance to interact with them throughout today. Some of them is going to come on stage and present, but this is the team who is going to take us to the next level.

So why invest in EPAM? We are the best position growth leader for enterprise AI transformation. Our AI native and foundational work is expanding, driving significant growth in markets. We are the strongest solution builders in the industry with proven track record of solving our clients' most complex problems. We have a clear strategy. We are focused on accelerating organic growth while driving margin expansion.

Let's dive into the details. I would like to invite Elaina Shekhter, our chief strategy officer and transformation officer on stage, to tell us how to transform our growth to market. Thank you much.

Elaina Shekhter:

Thank you, FB. Good morning, everyone. I'm Elaina Shekhter, and as of two weeks ago, I'm the chief strategy and transformation officer. Before that, I was the chief marketing officer, but today we have our brand new chief marketing officer here, encourage everyone to meet Phil Walsh, who's going to be walking around. Today, I want to talk to you about what we're doing to transform our go to-market approach.

Over the years, EPM's been particularly interested and really honestly obsessed with building the right kind of supply and addressing our customer needs in an overwhelming demand environment. Over the last few years, we've been investing significantly in our go to-market approach and the transformation of all of our selling motions.

And so today ... Sorry for the clicker. Three key takeaways. We are transforming everything in the company. As FB just shared, our digital platforms, our talent ecosystem, how we think about delivery, everything is being built around an AI native blueprint. The same is true with our go to-market approach. We are responding to an AI-centric environment with changing everything that we do in order to more effectively meet our customers where they are. That means that we're building domain and vertical expertise into every motion. Every sales engagement, every capability is driven around deep knowledge of our customers and their domains. And we're adapting the way we go to market through our programs that address customer reach, to our engagement and commercial models, and we're doing it in sync with or sometimes ahead of emerging industry trends.

EPAM predominantly serves the enterprise. We've been doing so for years, and although we have a significant footprint in ISVs and helping high tech and software companies build, they themselves are large enterprises. And so our primary segment today are large companies, and their service needs and their landscape of service needs has changed significantly with the rise of AI, and it has never been more complex. So between market conditions that demand the addressing new competition, rising customer expectations, and all of the AI hype, all the technology trends, which are constantly shifting on a daily basis, and our demand to meet expectations for advancing the transformations with AI, and the demands of the enterprises themselves, which are shifting also on a daily basis, demanding more strategy, more growth, better optimization programs, and overall better performance, and of course, the better use of capital. We are operating in a more complex enterprise environment than ever before. And the market depends more flexibility, more capability, and more results delivered more relentlessly than ever before.

And so to address these changing conditions, we are elevating our entire game and our go to-market strategy with three key motions. Number one, we're shifting and extending our focus from building geographic capability to building full scale capabilities. Think about a full stack of capabilities that includes domain, vertical, and effectively forward deploying those capabilities to our client engagements. Secondly, we're integrating a consultative approach around the whole of the go to-market strategy. So no more is it consulting over here, engineering over there. Our goal with our go to-market transformation is to bridge strategy and execution, and in doing so, create a consulting moat, in addition to the engineering moat, which my colleagues will be talking about right after this. And we're accelerating our motions, starting with partnerships, but not only. We are changing the way we address the market in total, direct to client motions, sales and marketing transformation, and of course, the work that we do continuously with our partners.

What this means for us is that we are future-proofing an organization by creating a forward momentum that's bringing capabilities to clients to meet them where they are today. So our evolving focus areas are necessarily about value creation. FB we mentioned to our hybrid teams, we have a long-standing history of building hybrid engineering teams. Today, our job for our customers is to build high velocity

performance teams that include consultants and engineers. We are prioritizing developing critical industry specific skills. This could be vertical, this could be horizontal, and we're doing that not only around AI, we're doing it with AI, more on this to come. Finally, we're creating a global delivery value creation network that's optimized, not just across locations, and Larry will talk more about that, but also around specific services and skills and capabilities of individual people and high performing teams.

Part of this integration is not only to build consulting into everything that we do, but it's actually to open EPAM up to alternative and additional buyers in order to capture new market share. Earlier this year, we announced the expansion of Empathy Lab into North America, having had a very successful launch last year in Europe. Empathy is our AI native agency, and it offers choice to CMOs who increasingly have their own budgets for technology, and yes, also AI, to engage with an EPAM that is ready to meet them where they are in driving key transformation programs in a way that is not encumbered by traditional agency dynamics. We also continue to invest and integrate EPAM Continuum, which is our consulting brand, and the changes there are material. We are upgrading the entire consultancy workflow with and around AI. And so in doing so, we're expanding our addressable market, and we believe not only are we serving our existing clients better, but we're expanding our opportunities to attract and build new client relationships.

EPAM has always been known as a technology solutions expert. This is everything we've been doing for the past 30 years. Across all three brands and across all of our front doors, we're adopting and adapting our solutions proposition around AI. By integrating consulting, what we can deliver is end-to-end enterprise grade scaled solutions in the absolute most complex environments. And for those of you who are staying with us for the afternoon, as you walk around this space, you'll see just how complex complexity is.

And so we're driving consulting to be in lockstep with technology, and in doing so, we can guide our clients on where and how AI should be used. We're helping to determine not only the right technology platforms, but the right operating models. We're identifying critical constraints and blockers around compliance, governance, security, very material, especially these days. And we're actually starting to run AI native work streams and business models end to end, this is part of our engagement model transformation. And so we believe we are the absolute best partner to scale solutions around AI and build for the future in the most complex enterprise environments.

So what about how we sell? To reach as many clients as possible with the most relevant propositions, we are transforming our full stack of sales and marketing motions in 2026. Everything that we've been doing for the last several years has been, quote-unquote, digital. We were focused exclusively on driving optimization, modernization, and AI foundational work streams, and this continues today. In 2026, our value proposition includes the full digitization mix, but it is also driving optimization and agentic operations into both the growth agenda and the optimization agenda of our enterprise clients. How we manage sales is changing from account relationship management focus to really creating a hybrid seller, someone who is a forward deployed relationship manager who is at once a consultant, an engineer, and a relationship manager.

We are adapting our pricing models. Of course, much of our business continues to be very much focused around TNM, as much of the foundational work we continue to do is built around high performance teams, but we're adding output-based, ROI-based, and business outcome-based models to our engagement mix successfully. Our sales cycle is changing from a more linear sort of traditional sales cycle to one that is continuous. This is definitely a work in progress, and it will continue to evolve very quickly as we introduce agentic motions into both the top and the middle and the bottom of the funnel.

Finally, marketing is transforming, and I'm very happy about that. So from sequential brand through funnel activities, we are introducing a performance optimized marketing motion. With Phil on board,

we're going to be sharing a lot more with you on what that looks like. So beyond investing, we are transforming our sales motions and our approach to market in order to capture additional market share. And nowhere is this more evident in the acceleration of our partnering motions. So we've been making announcements over the last few months and there will be many more to come and quite quickly I might add.

But today, our ecosystem of partners includes over 160 different partners. These include the platforms, AI native players, industry partners, universities, research labs, and such. This ecosystem is constantly being built out and adjusted to suit our solutions and consulting propositions. With our partners, our motion has changed from partner-centric channel motion to one that accelerates our propositions and our value to clients. We are elevating our market sensing capabilities and helping our partners do the same through very much tailored, dedicated, and often IP-based campaigns that we're bringing to market as we speak. We're also, in some cases, working with our partners to help them build their own platforms, and in doing so, driving delivery efficiency and effectiveness for their own build out operations. These are some of the partners we work with today. FB mentioned Cursor, there's many more, obviously, and there's a number of very interesting ones that are coming up, particularly around the area of security.

Over the last months, we've announced ... These are just really a subset of the things that we've announced. The point here is our relationships with our partners go way beyond credentials. We are pushing the edge of AI innovation, and we're doing that with our partners and with our clients. And so you're seeing us show up in market with AI wins, with being named the AI Innovation Partner for some of the largest CSPs, with announcing agents into multiple marketplaces. And this work continues and will be built on as part of our evolving go-to-market strategy.

And so I want to leave you with three ideas. One, we are very serious about transforming our go-to-market approach. We understand that the environment has shifted into an AI-centric environment and we are there for it. Number two, we believe domain and vertical expertise is a critical success factor, and it is creating not only an engineering moat for us, but also a consulting moat and positioning EPAM to win in an incredibly complex market. And finally, we are innovating and amplifying our partnership motions together with over 160 of the world's leading companies. And we're using that to adapt our models, everything that we do from how we deliver to how we engage with our clients. And of course, we're EPAM, so we're starting with the software development lifecycle and the product development cycle.

And so it gives me great pleasure to welcome my colleagues, Dmitry Tovpeko, who is our VP of AI Engineering, and Adam Auerbach, who's our VP and head of AI Enablement to the stage to tell you more. Thank

Moderator:

Thank you.

Dmitry Tovpeko:

Good morning, everyone. My name is Dima Tovpeko. I lead AI engineering.

Adam Auerbach:

And my name is Adam Auerbach. I'm head of AI Enablement.

Dmitry Tovpeko:

And Adam and I are going to walk you through what is changing in how software gets built, and why does it matter for EPAM Business? Boris Cherny, the creator of Claude Code, one of the most advanced AI tools in the market, in the recent interview, famously said that coding is largely solved with AI. If that's true, why do clients still need EPAM?

We believe there are four reasons for that. First, enterprise complexity is growing and the demand for complex engineering is infinite. Second, AI demands a new type of engineering discipline that is difficult to master, and engineering depth is our moat. Third, we are agentic platform builders, not just users or adopters. We are clarifying delivery and we are scaling a new type of engineering profile to run it. And fourth, what we build for clients today becomes the foundation for autonomous enterprise AI that they will require tomorrow, and every engagement brings it closer.

So now, let's talk about the first point, the first dimension, which is enterprise complexity. Our clients operate across eight simultaneous complexity dimensions, and each one of them getting new requirements with AI.

Strategy and economics. All of the client businesses are disrupted. They are discussing what they should be doing and how they should be transforming their primary core products, in addition to technology and product transformations they need to run internally.

Data foundation. Your agents are as good as your data, and enterprise data is not ready for AI.

Vendor strategy is a good one. Everybody is talking about which tools to select, but conversation's also shifted to existing SaaS applications that are currently part of everybody's portfolio, and now clients are discussing whether they should retain them or they should rebuild these capabilities with AI. And that creates a new set of questions and a new stream of engineering work.

Every single dimension is getting new requirements. It is getting more and more complicated, and clients need a lot of help here. And this is even before we talk about the challenges that happens inside of software delivery and software engineering itself. Let's talk about it.

When AI generates the code, the hard part becomes how do you create a system that generates it right? So it all starts from design. Somebody needs to encode their specifications, what should go inside of them. All the domain knowledge, all the business workflows, all of this proprietary knowledge in undocumented systems that is sitting inside of people's heads, all of that need to go there.

Then somebody needs to architect the system. It is never a single agent that can do the work. This is always a complex, agentic ecosystem that is ever evolving and ever getting more and more complicated.

Then somebody needs to validate the output, somebody need to judge. We see that the same tools produce very different results depending on engineers who are dealing with these tools. And the gap is getting wider.

And finally, you need to connect these agentic ecosystems to your enterprise environments. With all of these established legacy ways of working, delivery pipelines, ecosystems, tools, integrations, all of that, and none of that was designed for AI, and now we need to deal with that. And that creates a huge complexity inside of software engineering, which now is getting a new AI engineering discipline that didn't exist 12 months ago, to create AI engineering layer that can run the agents that are doing the work.

And this is exactly what we have done at EPAM. We qualified the agentic system, the entire delivery pipeline with Agents. I'm not talking about agents augmentation. That was easy part. So this is going. So what we are doing, they're creating a brand new, from the ground up, AI native ways of working that we qualify it in a repeatable pipeline. And that's the blueprint.

Now, the blueprint is an easy part. The hard part is how you can actually scale it, and in order to run it, you need a new type of engineer to deal with that. Traditional, narrow specialized software engineers are actually not good in benefiting from these kind of blueprints. They can get maybe 10%, maybe 15%, but all of these stories about 2X, 3X, they require a very different profile. Somebody who can own engineering tasks end to end, across all stages of SDLC, across multiple technology stacks.

And this is where it is becoming very complicated. They need to be fluent in new AI tooling. They should be using them in very different ways. And then they should be able to judge whether the outputs at every step in the way are what they should be. And this is what we call full stack agentic engineer profile, and scaling this profile is the hard part. Aleksey and Sandra later today are going to talk in more details about it.

So the question is, can anyone do this? We believe there are two things that are required to operate this at scale, and most companies cannot assemble both. First, you've got to have strong engineering culture and depth. This is not an upskilling program. This is not a scaled certification exercise for a particular skillset. You have to start from the very high point from the very beginning. You have to have that as a part of your DNA already in order to be able to run in this race. And we set these high standards many years ago, and now, 30 more years later, we are starting from much higher point than many.

And second, you got to have delivery volume. You have to be able to run this pipeline against real enterprises over and over again, and this is where blueprint are getting better tested. This is where they are becoming real. This is where you are facing real legacy problems. And this is where they are becoming scalable and they can bring value to our clients.

So, why we think most firms cannot assemble both? Cost arbitrage firms, their model was optimized for narrow, specialized engineer. Low rates, low complexity work, and there is a required particular profile to make it a profitable business. And we believe that these firms are exposed.

EPAM hasn't been ever really playing a role there. We approach it differently. Strategy firms, they have intellectual depth, but they don't have muscle on the ground to make it real, to actually deliver on these promises. And product companies, they have engineering culture, they have great products, but they're only integrating with the enterprises. They're not working from within inside of this complexity.

And EPOM has both. We have the engineering culture and we have the volume, and that's the mode, and we believe that AI gets it wider. And this only comes from doing the work. You are as good as your delivery. In AI, the right way to build reveals itself only through doing. No one figured it out from a whiteboard. All the great founders right now for AI tools, they're all hands-on. They all have a ton of experience. And this is what we have built, and this is why we believe it is hard to replicate.

And now, Adam is going to walk you through how we are scaling this across our clients. Adam?

Adam Auerbach:

Thank you, Dima. So, what Dima is describing to you is what we call level three maturity. What we have found is that there are multiple levels to this journey, and most people start at level one. Level one is, they have access to Copilot, a code assist tool, but no one really uses that tool. So if you buy a tool, it doesn't mean that people are going to use it. People need coaching, they need training, they need support.

And what they ultimately will find out is that that tool optimizes one aspect of development, coding. So yes, does it create efficiencies for developers? Sure. But as Dima just said, there's much more to the software development lifecycle than just writing code. And that's why you also need agents, and that's what level two is. Level two is this combination of code assist tool with agents to help accelerate your current ways of working.

But that's the next challenge. So, yes, level two will create the efficiencies that people are expecting with AI, faster cycle time, higher quality, better productivity. But when Dima talks about delivery as code, he's really talking about a whole new way of working where we get to what's called spec-driven development.

But that means that your process has to change. And I have been in IT for 25 years. I know I don't look that old, but when I first started out, it was around moving from waterfall to agile. Companies are still struggling with that today. So now we're saying, hey, we're going to introduce this new way of working. We have to get over the fear and resistance from people, and then once you get to some level of accomplishment, there's yet further improvement. So this is a long journey that people are on to get all the way there.

And we do luckily have some really great case studies, like PostNL, where we are delivering agents, we are getting them to this new world, this new reality. And that is, as Dima said, that moat is the fact that we have so many of these projects right now, and we're learning from ourselves and getting this experience that we can then bring to our clients. And that really sets us apart.

As FB mentioned, we've created something called AI Run, and that's our suite of consulting services and education and tools around how to drive this transformation for our clients. We're focused here on engineering. Nir and Eli, who are going to speak next, they're going to talk about how we're doing this for the business because there's a lot of similarities here.

I'm going to double click into each one of these for a second. So, the first one is blueprints. So, Dima, you talk to a lot of CIOs. How many IT leaders can really articulate the current levels of productivity for the organization?

Dmitry Tovpeko:

Well, not very many. Definitely not on the second meeting.

Adam Auerbach:

So if the board is saying, "Hey, I want to see 20% productivity boost from AI," that's a problem, because they don't know what their productivity is today. So for the last many years, I've been at EPAM for eight years now, we have something called engineering excellence. It's what makes EPAM so special, our engineering talent, how we really raise the bar in our delivery centers and our projects. And we have a consulting offering that we've been running with our clients where we help them baseline their teams, their performance. We establish those KPIs and then build improvement plans, so that they can be more agile and leverage DevOps and get to continuous delivery.

So we're able to take that same methodology, go to an organization, understand how are they working, and then from that figure out, okay, where is the place that AI is going to have the best value for you? Instead of just saying, hey, let me be haphazard, we can be really targeted in which agents we build, the education, and then we can track the progress. And we win projects because we can really articulate, this is how fast you're moving today, this is your current levels of quality, and then here's the benefit of AI.

We have a really great case study with Edward Jones. We're working with them right now. It started a year ago with a pilot. We were able to show with our products and Copilot the efficiency gains we could deliver in a short amount of time, and now we're in the process of scaling it out to the rest of the organization. And we have many of these projects happening right now.

Dima showed you this picture, delivery as code. And so, I just wanted to go back to it for a little bit, just to articulate a couple things. In the blue boxes, which maybe are a little tough to see, these represent

different agents, or maybe it's agents calling agents. There's maybe 10 or so, maybe a little bit more in this picture.

If you're an organization, you can't just apply agents blindly to all of your teams. Every team supports a different application. A large enterprise could have thousands of applications that make up their platforms. So what that means is that every team is going to need a different set of agents tuned for them. They have different tools, different technology stacks, different ways of working. And so, the scale of this gets pretty big pretty quickly.

And so, what we have done is we have built a set of tools for ourselves under the AI Run platform umbrella. We have things like Dial, CodeMe, Alita, AgenticQA, which we can bring to a client to accelerate their adoption, as well as we can handle how you can take an agent and basically copy and paste it and tune it quickly for the next set of teams, and manage that at scale with the observability and governance that's required for a large enterprise.

Shameless plug, we have a booth, so later on, if you want to see a demo of the tools, we can definitely show you. The tools are real and they're spectacular.

We built these tools a couple years ago. It was really important for us to be able to use them to learn, and now, we definitely have projects where we don't use our tools, but what it allows us also to do is quickly understand, what are the people and process limitations that are preventing wide scale adoption of a client quickly, but also we can bake this into our own projects. So if a client's going through the transformation that takes many months, maybe years, we can come in with our tools quickly deployed with our full stack Agentic engineers and really be able to deliver the value of AI quickly.

And then lastly, before I hand it back over to Dima, when you talk about level one, level one and agile team, people are working in silos. And so, when they start to use AI, they can create some efficiencies for their tasks. In the industry, we've had this term called a T-shaped engineer, and a T-shaped engineer means that I have this one skill, maybe I'm a mobile developer, but then I can also maybe do some API development and maybe some backend work, right? So, that's T-shaped.

With AI and agents, I can really deliver on this promise because that T-shaped person can be sometimes a unicorn. But with AI, I can give people agents to really help expand what they're able to do. So I could have a front end developer who now is able to do, like Dima said, full stack engineering. They can work across all levels of the application, of the platform, and do many different things. And now with AI, they really can run the entire software factory, that delivery as code.

And so, now what we see in our new teams is this combination of this full stack agentic engineer with the combination of product and design, and this is how we deliver products in this AI native world. And Dima, I think you're going to talk to us about why the... While this might shrink some of our teams, the demand is actually far bigger.

Dmitry Tovpeko:

Thanks, Adam. And now I'm also eager to look at the tools again. So, Adam just took you under the hood, so now let's talk about the implications to the market. The common assumption is, the faster we can go, the fewer engineers you need. And for a lot of work this through, we definitely see this on the ground. But at the same time, this is actually not the case in many places where we operate with our clients.

There is a fixed pile of work at the top. Maintenance, application support, second tier applications development. There is only so much work that you can do, and this pile of work is doomed to be shrinking over and over. And all the firms that operate in there, they're all exposed. As I said previously, this has not been the place where EPAM was generally operating.

Where we operate largely is below the waterline, and this is where we see infinite backlog space. Our clients have been sitting on years worth of queued work that previously they were not able to attempt. Products modernization, technical debt, elimination, new products development, just higher velocity and productivity deliver more and more features for their own clients.

A lot of this work was put on hold or was tamed for the reason it was too complex, too expensive, took a lot of time to deliver, or simply was not possible because of our technology limitations. Now, AI makes it possible.

Edward Jones, they had a dormant mainframe modernization program that was in a slow motion mode. And now, with our AI-run platform, MFLens, we all over-competed incumbents and now we are helping them to deliver, and now it is active.

Baker Hughes, we are a strategic engineering partner for them and helping them to work on a variety of different strategic programs in the range from data products to field level AI systems across all of their operations.

NelNet, we came in and we helped them to accelerate their velocity. We'll help them to define the new ways of working. And as we increased our velocity, they wanted to do more of that. They increased their expectations, how many new features they wanted to deliver, and we scaled our footprint. More speed, more demand.

Firms built on fixed demand layer are competing to deliver the same shrinking amount of work for cheaper. Efficiency without growth is raised at the bottom, and EPAM has not been operating there. We live below the waterline. Every time we get faster, clients are attempting to do more, and that's today's picture demand.

Now, let's talk about what's coming next. Everything that we are building today, the agentic ecosystems, the agents orchestration, the enterprise hardening, that's the infrastructure that autonomous agents will require in future. Clients are paying for agentic delivery today, and tomorrow it becomes the autonomous ware.

Moreover, these autonomous agents in the first place will be attacking this top of the iceberg that I showed on the previous slide, lower complexity, lower stakes work, where we have not been operating, and this is where we can actually enter there as agents and agentic platform builders, exactly the type of complex engineering work that we've been famous for, and we can enter there as builders, not as incumbents that are protecting the margins.

So let me repeat the four key takeaways and the four points that we started from. Enterprise complexity is growing. Every layer adds another, and demand for complex engineering is infinite.

Second, AI creates the new engineering discipline that is difficult to master, and engineering depth is our moat.

Third, we are agents builders. We are agentic platforms builders. We are codifying delivery in new ways to accommodate agents-first mentality, and we are scaling a new type of engineering profile to run it.

And fourth, the investment that we are making and the work that they're actually delivering today for our clients, for agents, that's the foundation for autonomous enterprise that is coming tomorrow. And with every engagement, we are getting closer to it.

So, I started with the question, if coding is largely solved, why do clients need EPAM? Coding was never a hard part. Software engineering was. The better AI gets at writing code, the more what we do matters.

Now, I want to show you the video, the client testimonial from Larry Fitzpatrick from OneMain Financial. Thank you.

Larry Fitzpatrick:

Hi, I'm Mary Fitzpatrick, CTO at OneMain Financial. OneMain is the leader in offering non-prime consumers responsible access to credit. We offer hardworking Americans personalized lending solutions, including personal loans, auto loans, and credit cards. We operate across 48 states, online, and in 1,300 branch locations.

I lead our technology strategy and the teams building the digital data and core platforms behind our growth. I joined six years ago after AWS, and I've spent my career scaling technology organizations at the intersection of innovation and execution.

In 2023, we made a deliberate decision that generative AI would change our industry and we would adopt it responsibly. We started with optimizing our guardrails for the unique risks of Gen AI so teams could move with confidence. One of several strategic opportunities we are focused on is our product development and operations lifecycle.

Despite rolling out tools to teams, adoption was uneven. We met with many potential partners. Most sold slides and could not demonstrate performance. EPAM showed us how they were already working this way inside their own teams for over two years. We chose a partner who had done it, not just described it.

Late last year, we engaged EPAM to work with the organization. It spans about 100 teams across the full product development and operations lifecycle, from product strategy and design, through build, release, and run. EPAM didn't bring us a point solution. They brought an end-to-end system, a clear methodology, a working platform, and experienced practitioners who operate as one team.

We started with structure. Their SDLC maturity model gave us a simple progression, AI-enabled to AI-engaged to AI-native. On the platform side, we defined an AI Agentic ecosystem tailored to our environment. EPAM deployed their AI Run agentic platform, and we integrated it into our stack, SSO, Jira, Confluence, Git. The tools meet our teams where they already work.

We are still mid-journey, but engagement across our teams has exceeded expectations. The energy is real and it is translated into meaningful results. This is a journey, not a destination, and we've accelerated greatly partnering with EPAM.

Nir Kaldero:

Good morning, everyone, and thank you once again for joining our investor day. My name is Nir Kaldero. I'm EPAM Chief Data and AI Strategist, and I'm on stage with great friend of mine, Eli.

Eli Feldman:

Eli Feldman, CTO.

Nir Kaldero:

And together we lead our enterprise AI transformation agenda on the business side. Today, we want to show you how we help our clients accelerate their journey towards an AI-native enterprise, through robust offering portfolio, differentiated delivery playbook, and end-to-end capabilities. Our goal is simple. We want to demonstrate not just what AI can really do, but why EPAM is uniquely positioned to win in this era of AI-native business transformation.

We will walk you through four core areas around AI business transformation. The first one, how AI-native transformation is reshaping business innovation and operations, and how EPAM accelerate the journey with meaningful impact and growing book of business. Second, how our unique AI Run

Transform playbook turns strategy into measurable business outcome. Third, how strategically we expand our service mix to support and lead the next wave of AI adoption to successfully support our clients. And lastly, how are clients' biggest AI challenges driving long-term structural growth tailwind for EPAM for both business and engineering altogether.

So, let's dive in.

Eli Feldman:

Thank you, Nir. We see this space transforming approximately along the same ways that Dima and Adam just described. There is maturity levels, there's stages which organizations go through. The most foundational stage is start optimizing current operations. Easy place to start, but that requires a very meaningful foundation. Adam and Dima were talking about the foundation and engineering. That is a critical ingredient. Must be there. This is not your grandfather's business intelligence capabilities. These are foundational platforms and capabilities that need to be put in place, all the way from engineering to data platforms to business capabilities, to enable that. Once we solve that aspect of the challenge with our clients, then we can actually start transitioning to building business functions.

Now, to make it very clear, this is not about just bottom up. The bottom up is sort of the foundational and technology enablement. It is critically important, but that's not the only pathway. The other one is the top down, understanding the business case, understanding what we're actually optimizing in the business. We'll go through some of the examples.

Then once we figure out that initial optimization space, well, we can now focus on growing the capabilities, running, maybe even semi-autonomously, the capabilities that these organizations have end-to-end. And then, once we capture all of this intelligence from the business process and the capabilities and the data assets and governance that is put in place to run that capability, then we can start identifying this new business opportunities for our clients, working with them together to bring that to the market.

I'll give you two examples of work that we have done with our partners, with our clients. Critically important, each one of those started pretty much in the same place, the foundation. Cannot skip that. Have to enable foundation, have to have the right engineering in place, have to have the right data platforms in place, the governance, the observability, all of these capabilities, just to start even within a simple business process optimization. Once you have all of that data, well, all of a sudden, you can actually see how you can start optimizing, how can you build agentic around the business process and start optimizing.

In the first case with a global cosmetic manufacturer, that first business case was demand prediction. We wanted to predict what actually sells in stores. We did that. The only challenge is simple. If you know the demand, but you don't have the supply, you didn't really solve the problem. The business is not really benefiting.

So, well, the obvious projection from there was, let's try to figure the supply. The compounding problem, because you need to figure out the supply from the manufacturing process, or maybe even before that, all the way to when the product hits the store, is actually a compound data problem that is much more significant than any one of the individual elements.

Just to give you a sense, from a supply chain economic impact perspective within this organization, a weekly risk reduction, 100% risk reduction in this company from a an economic impact perspective, so sales versus costs, is about \$16 million a week. In the past, before any of this was implemented, humans looking through dashboards, again, that old fathers, grandfather BI system, dashboards and reports and

stuff like that, could solve 30%. It's meaningful. \$5 million in economic impact, they could have solved it. But there's the long tail.

So when we started introducing the capability and sort of integrating all of the data together and working with the supply chain organization to figure out their value stream business process and all of that, we realized that about 50% of what AI actually recommends, within the 30% slice still, is very much consistent with what the organization actually was doing so far. Excellent result. But it actually recommended the rest and almost closed the entire risk gap of the \$16 million a year. A week, sorry.

This opened another interesting conversation, as Dima and Adam were saying, Dima was saying about sort of SaaS platforms and packaged capabilities and stuff like that. You see, vast majority of organizations out there, manufacturing organizations, supply chain organizations, CPG and all of that, they have to rely on package supply chain tools because building a custom supply chain implementation across the board is extremely expensive. There is no, in the past, there was no ROI for that whatsoever. The largest supply chain organizations maybe, but most organizations could not.

Now, the moment we solve the supply chain from manufacturing to store, all of a sudden they say, "Well, we have another tail of that problem. How about from

... from the manufacturing and from the warehouse through the ingredients. They had another package that was solving that, but the two were not really connected, so they would manufacture one thing. Demand is something completely different. They optimize for that risk. So there is massive problem in between. Come AI. Now the implementation of that end-to-end supply chain, custom-built for that organization all of a sudden is a viable alternative to several complex integrations off the shelf tool, SaaS platforms, et cetera. So all of a sudden, we are actually capable of solving a significantly more meaningful business problem for the organization while leveraging everything that we have been talking about so far in terms of technology enablement and data enablement and governance, et cetera. Another observation that you see on the slide, this field is continuously expanded. You prove one case, not prototype, prove one case in production, organization actually seeing economic impact.

All of a sudden, well, we have this other business case, another business case. And then it's expanding pretty much exponentially in that case, even within an individual organization. And AI enables that because all of a sudden implementation is cheaper. So you can actually leverage the same budget to do significantly more work.

I would like to speak about another client. You would think that CPG, well, not regulated space, pretty easy, but the reality is risk reward. In regulated organizations, the next case is a major global pharma. In global pharma, and the foundation was exactly the same as before, build the foundation, build the data capabilities, build the engineering capabilities, solve a business use case. Once we have done that, this client actually designates as a strategic partner for the entire stream of work around AI. They said, "Okay, we have another major business problem; clinical trials." Clinical trials, 1% of defects in clinical trials, just 1% of defects in that process cost the organization \$28 million in economic impact because it delays drugs to market, all of that stuff. Most of it is a top line impact. It's not even optimization. It's not really even cost optimization. So now, if you are able to solve even into the lowest double digits in that case, you actually have a very meaningful top line impact on that organization.

Once this organization learns what actually existing setup means, then they are capable of understanding their assets, now we know what our clients need. We can actually convert that into something that is significantly more meaningful. So you have two examples here. One is multinational for consumer loan and garden products. They actually leveraging all of the foundation and all of the capabilities that we have built in business optimization. Said, well, we can go to DTC, direct to consumer. We couldn't do that before. We were selling through resellers all of our life, now we can go direct to consumer. A clear business value that was enabled by AI as well as all of the other work that was done.

Swiss Re, which you will actually hear much more details in the panel later on, but they realized being a reinsurer, they realized that they actually sell data. Again, all of that foundation actually paid off and enabled a new line of business for them. So the reality is that EPAM wins at the first stage. We help optimize because we deeply understand the technology, the bottom up enablement capabilities, the technology, the engineering, and the foundation that we can build to our clients with AI native enablement, of course. We win in growing and running the business for our clients because we can layer the rest of the pie from a business transformation perspective. We understand the people transformation, we understand the business. We understand the value streams, we understand the flows. Now we can actually layer the two together and significantly enable these organizations as well. And then we can leverage all of the deep agentic capabilities. All of our experience of the past 30 years building go to-market capabilities for our clients and actually enable them to create new set of businesses that they have that are AI native.

Nir Kaldero:

Nir, please give us the details on... Oh, sorry.

Eli Feldman:

One more slide.

Nir Kaldero:

One more slide. Sorry.

Eli Feldman:

All of that is actually quite systemic. Adam showed this slide before, and we work very closely with the technology organization, obviously, to enable these capabilities where we have the blueprints for the technology enablement, the prompts, the sequences, the workflows from the engineering perspective. We actually developed the same from a domain and industry expertise. So we come to the customer with deep understanding of the value stream of what they actually need to solve from a business perspective, enabled by technology out of the box. And we are capable to solving that. We understand that none of this is possible with individual contributors. We must build networks of experts to be able to solve these complex problems. And these are networks of experts that include, again, engineering is critically important, but people that understand people, change management, transformation, domain, industry, governance, and all of the other stuff that needs to be in place to make it work.

And then the tools and the platforms that need to be in place to enable that, the time to market needs to be accelerated. So we have to come with some accelerators, some harnesses, some productized offerings to be able to make it faster and more effective for our clients. So now, Nir, please take us through some of the details.

Nir Kaldero:

Thank you, Eli.

Eli Feldman:

Thank you.

Nir Kaldero:

All right. So now let's talk how we are expanding our strategic capabilities to lead the next wave of AI adoption. Successful AI adoption comes down to three kind of main pillar. Think about it as a three-legged stool. The first one is the data, which is the fuel and really the foundation for AI. The second one is technology, which is the environment and the infrastructure to really deploy AI and use it at scale. And the third one is people, culture and process, which is probably the most important pillar here, where you really want to make sure that what you built is adopted and then delivering the business value following the investment. Across these three, we are expanding our AI strategic service capabilities to help our clients transform at scale, to ensure we stay ahead of the market and help them. Let me walk you through these kinds of four areas of expansion.

The first one, we are reshaping our consulting model into something entirely new, AI native, verticalized consulting built with and for AI. This isn't just traditional advisory. We use AI to conduct consulting itself. Instead of slide decks, we deliver prompts. Instead of static artifacts, we enable small language models across evolving processes. Instead of isolated recommendations, we co-design simulations with agentic tools, and we really aspire to help business leaders run scenario planning with AI agents in days and not weeks or months even. We deliver consulting also for AI, the practical building blocks that make AI successful in productions all the way from operating models, governance, responsible AI, cybersecurity, adoption programs, and value tracking. Our consulting proposition is built for one purpose; moving AI from experimentation to production at scale. The second one is we are building the future of business operations, as FB mentioned. We are experimenting with and plan to disrupt the market through an agentic led business operations offering, where we design, build, and run high-end processes powered by agentic AI.

This lets us expand our share of wallet, evolve our service mix, and grow our total addressable market through next gen managed services. The third one is where technology and domain expertise truly converge. We are building deep industry knowledge with strong AI capabilities and acumen altogether. Through our proximity to clients, we are developing industry-specific data models, co-creating vertical ontologies with strategic partners and assembling pre-built agentic workflows tailored to how industry run. The payer for our client is simple. Faster AI deployment in their specific context with less risk and greater precision towards the ROI. The fourth one, we are evolving our accelerators. We have been already expanding migVisor into an agentic-led migration platform. We also extending, for quite long time, dial AI as an agentic orchestration platform.

Think about it building agents with prompts. You have the ability to deploy mixed frontier models and ensure that AI is really deployed at scale all the way with governance, security, and FinOps from the get go and from the start. Together, these kinds of four capabilities position us to be ready and ahead of the market to deliver real measurable value for our clients.

Let me close with why we believe AI is a long-term structural tailwind for EPAM. Real AI business transformation isn't really just about deploying models or tools. It demands business model reinvention. Think about it as the culture and the mindset shift that enables completely new ways of working. Process reimagination, targeting the right workflows and designing AI enhanced experience. Data monetization and modernization, really breaking the silos, capturing new datas, and building the architecture and the semantic layer for reusable real-time intelligence across the enterprise. And obviously other critical services and elements across the end-to-end AI innovation life cycle, all the way from AI strategy to MLOps and AIOps. But think about it. This complex business transformation work stream, also generating significant downstream investment in core technology and engineering demand to enable the foundation to run, deploy, and use AI at scale, which altogether, if you think, creating a significant opportunity for EPAM to lead in the market.

The business transformation work and the technology work also deeply interconnected, and we see both of them are growing. We are uniquely positioned to deliver strategy and implementation simultaneously to enable the full deployment and full scale reinforced by our AI native talent and unique playbook. And our end-to-end capabilities is really and truly our competitive advantage. And this is why we believe EPAM will continue to capture market share as AI accelerates globally.

With that, let me conclude and have some kind of key takeaway to leave you with. The first one is we are driving our client's AI native business transformation at scale. A few great examples that Eli showed on stage. We are leveraging our unique and proven AI run transform playbook on the business side to turn AI strategy into measurable business outcome.

Eli Feldman:

We're expanding our service mix to unlock new opportunities while staying ahead of the market to support our client's AI adoption journeys, and our client's biggest AI challenges create long-term structural growth tailwind for EPAM within both engineering and consulting strategy simultaneously. With that, it is my pleasure to introduce our next client testimonial.

Ahmet Tezel:

My name is Ahmet Tezel, and I'm the chief innovation officer at LivaNova. My role is to lead end-to-end innovation in the company. It was clear to me that we needed an external partner to help us out in creating a cloud platform and products that go with it. And I had experience with EPAM from a previous company, and it was a good experience. One of the challenges if you're an epilepsy patient is that you have to go to a physician's office about eight to 10 times in your first year post-implant. And the reason is that you go there to get your device adjusted with respect to its parameters. Now, this is not easy for epilepsy patients because they're pediatric patients or if they're adult patients, they usually don't have a driver's license. So it's a complicated task. On average, you travel more than 30 miles for each adjustment.

So doing this in house in a hospital setting is difficult. Now, there is a huge unmet need here where you can do this adjustment in a remote setting where the physician can connect to the device remotely and talk to the patient and do the necessary adjustments. So that's the program that we developed with EPAM, where EPAM was able to create us and work with us a secure private cloud connected care system that enables physicians to connect to our products remotely and adjust the parameters of the patient's device remotely. I envision that we will continue to work with EPAM.

We now have the first FDA approval for our first franchise, our epilepsy franchise, through the product that we developed together, but I envision that we will continue to work together as we expand the partnership into other business units that we have in the company. We have a broad neuromodulation franchise with different disease states that could benefit from cloud-connected care, and we also have a cardiopulmonary franchise that can certainly benefit from having a connected ecosystem for their devices. So I envision that we will continue to work together with EPAM as we roll out our digital ambitions to our broad business units.

Moderator:

So we're now going to turn to our question and answer session. I'm joined by FB and Elaina here for about the next 20 minutes, call it. Just a quick couple of points for those in the room, just please raise your hand, wait for the mic to come to you, state your name and firm, and we will get to as many questions as we can. We also, of course, covered the overall strategic overview, our transformation, and then our AI native pieces of the business. So we kindly ask to keep your questions tailored to those

sections as we have much more coming up later in the afternoon, including our financial imperatives and multi-year outlook. So with that, we'll go ahead and open it up. We'll take one here in the front, Mr. Bergen.

Brian Bergen:

All right. Thank you. Brian Bergen from TD Cowan. So appreciate all the color you've given so far. I wanted to ask on the go to-market transformation. So trying to understand really how material this changes for you. You talked about a consulting-led approach in the past. So what are you going to be doing differently now? And I think you also mentioned maybe potentially some client-facing personnel changing. Just talk about how you're going to manage execution risk around that.

Balazs Fejes:

Let me go back a little bit, Brian. Good to see you and thank you for that question, but let me go back a little bit about EPAM. EPAM historically was very much operating in a seller's market, right? If we created the capabilities because of the resource shortages, people were coming to us, and it was very much us showcasing our capabilities. I think in the last years we learned that it's much more of a buyer's market, which means that we need to be more proactively marketing our services to them and actually start creating a more targeted go to-market motion backed by marketing. At the same time, the way how we're managing our client relationships are also changing, and we started to make those changes probably in the last one or two years, very much focusing and becoming more client-centric and very much highlighting the way how we're solutioning with our clients.

Also clients right now increasingly more transforming how they're delivering their businesses, as Eli and Nir were talking about, and we need to provide help. Elaina, could you add something to this?

Elaina Shekhter:

Yeah. So thanks, Brian. Good to see you. For sure, there's a couple of things going on. So as FB said, we have to go get more of the business than we've ever had to do before. And we're actually changing that go get motion, not just to sell AI, but changing it with AI. So there's a fair amount of training that's already happened. There's more in terms of sales enablement and sales training to come. And, yes, I think that there will be some rotations in the field. I think that's natural and expected and in fact welcomed. But one of the biggest changes that we've made this past year is really integrating the industry consulting groups, which were historically for us more of a standalone service line into our IBUs, into our industry business units. And what that's creating is these high velocity teams that I spoke about and that you heard about now. Is it a risk? Probably. Is it absolutely critical? Definitely.

Moderator:

One here in the front. Jason, please.

Jason Kupperberg:

Thank you. Jason Kupperberg from Wells Fargo. Really appreciate all the detail. I wanted to ask about these full stack agentic engineers. Interesting new role it sounds like. Tell us a little bit about the profile of these individuals. How many of them do you have today? How many of them do you think you'll have in two or three years?

Balazs Fejes:

Jason, good to see you. I think it's a really good question. Clearly this is something which we are growing rapidly right now. We have very much focused on this space. You will hear probably in one hour or so from now from Sandra and Aleksey how we're actually creating, how we finding them in the organization and what training program we are putting through that. Actually, this capability is growing really, really fast because that's the real focus area. What we're doing is we're identifying them. We are actually putting through them with a rapid pace of understanding it. And probably in the last just three months, we just doubled the capacity of that capability or that headcount. This is something which is going to be our standard motion going forward. And in every discipline, every line of business, we are basically pushing our engineering teams, not just on engineering teams, but even account managers and delivery managers or the sales team at how to adopt and how to use AI.

And just two weeks ago, we launched quite aggressive and pushy programs to making sure that our salespeople, account managers are actually using agentic tools to not just deliver their account plans and solutions, but actually understand fully how to deliver these applications. So this is ongoing effort, that's where our investments are going, and we believe this is what's going to differentiate us and going to allow us to really scale in the years to come.

Moderator:

Thank you. We have one here in the front.

James Floss:

Thank you. James Floss from Morgan Stanley. Thanks for putting this on today. I wanted to ask a little bit as you change the engagement approach and sounds like some of the development approach, is that going to necessitate also a change in the way that things are architected from the beginning, and how does that impact things like sales cycles and project scaling and that kind of thing? Thank you.

Balazs Fejes:

Good to see you. Absolutely it's changing. And actually if you will, just a shameless plug as you are in the audience, go after the session. We have a whole video actually explaining to you how we are using what we call AI factories in the sales process, how it's actually integrated in our RFP creation, RFP responses, which is really going to change the way we are going to market and actually sells our efforts. But it is changing not just how we're selling. It's not just the way we are contracting, it is changing how we are architecting the solution, how we're putting together the solution itself. We will be talking about how we quality assure all the proposals and all the estimates and using AI. So this is very much ingrained into our go to market motion, the way we're delivering, the way we are go to market and actually how we build a solution and how we're using AI in every possible step where it's possible. Not just possible where we are able to figure out how to plug it in today, and we're finding new and new ways every day.

Moderator:

Two over here.

Brian Keen:

Hi, it's Brian Keen at Citi. Can you talk a little bit about going after that fixed demand, some of that work that you guys didn't do traditionally that was more labor arbitrage, how you guys can get into that market through AI, and how fast can you disrupt that market at coming in at different prices?

Balazs Fejes:

I think it's a good question. I think we had early indications that we had success in this space in the last months and weeks. We made many, many proposals in this area. It's probably too early to call a full success, but we see real promises in this area. You're going to, again, shameless plug, you're going to see some amazing videos and demonstrations behind you around how we're going after the manual testing space and how we're going after the intelligent operation space with AI, how we're helping that in this area. Also, we're going to start seeing capabilities, how we actually doing BPU automations for some of our clients. And actually we have public case studies around that where our clients are starting to see real ROI, us replacing more traditional call center agents with AI based solutions. So how fast is going to scale? It's probably early to tell, but we are seeing demand, interest from our clients.

And because we're coming in with a very fresh point of view, we're coming in with new ways how we're approaching it, with a new price point, a new way of delivering it, new way of taking advantage of AI to do knowledge transfer, this creates quite a buzz in our community. Yes, absolutely.

Elaina Shekhter:

So just to maybe put a fine point on it, for us, it's a transformation pitch. It is not a labor arbitrage optimization pitch and all of the attenuating things that go with it, including organizational design, platform architecture, et cetera. So it's much more than just a labor arbitrage market capture opportunity.

Moderator:

I want to hear in the front and next.

Punit:

Thank you. This is Punit from JP Morgan. So as you pursue AI native SDLC, bring AI into SDLC, which changes the way you engage with your customers, talk to us about change management aspects from clients perspective. Are they ready? Or more importantly, are their employees ready for these changes, and will all the recent news flow around Anthropic and the development there in Cloud and everything? Has that changed their behavior in any way?

Balazs Fejes:

Punit, great question. I think if I want to summarize it is a change management process. We're going engage by engagement, project by projects, and we're talking about thousands of engagements, which we are migrating, which we are elevating in terms of maturity. Is our clients employees all ready? No, but it's an opportunity for us. We are giving them education. We are giving them advice how to change organization, how to introduce new tools, how to actually go through this whole education coaching process. Most organizations just went out as Demo and Adam talked about, went out and bought the tools, and they said, "Here you go. Now we expect you to be 15, 20% more efficient." And then a couple of months later, they found out that it's actually a J-curve, and their productivity dropped. So they said, "Okay, why don't I use some online resources?"

This is where you can read about it, and there are some forums, and nothing happened. This is the point where we are entering to the picture, where we really start advising them and coaching them how to actually mature the engagement model, but they are not ready. I think all the changes you are referencing, which is Anthropic or OpenAI, launch of cloud code or Codex, this is only for the really mature clients and mature engineering teams. If you just launch in a legacy code business brownfield, any of these tools, these tools go wide, and they're actually not going to create any productivity because

you need the specs, because you need to describe the brownfield itself, the expectations, and you need to have the right tooling in place. So it takes a while to adopt. It's a change process, and we see a multi-year adoption for the enterprises.

Moderator:

I think we had one here, and then we'll go to the one over there.

Nate Svenson:

Thanks very much. This is Nate Svensson from Deutsche Bank. I'm going to build on Punit's question here. I really like the slide with the three levels of AI adoption. I thought that was a useful heuristic. It sounds like most companies are on that first inconsistent and ad hoc usage stage of AI adoption. Your differentiation in moat is going from the second to third stage. I guess the question is, if most companies are in stage one today, how do you help them get to stage two to ultimately get to where you have the most competitive differentiation? Why are they going to choose EPAM to go from stage one to stage two versus a different system integrator, other sort of competitor, and how do you maintain that client relationship as we continue to progress?

Balazs Fejes:

Very good question. I think why are they going to choose EPAM? Because we will go in and show you not just slide decks. So this is the case where we're showing slide decks to you, but in most cases we are coming with real examples, real blueprints, real proof points. How are you going to get there? Very practical. How can we actually go in there? It's a very hands-on experience. Our clients are seeing that the leadership team who we have, the people on the field really understand how to make this happen. So this is the experience. When they talk to me, they actually see on my computer, I'm learning a cloud code, and it's very different discussions when the CEO really starts talking to them about the best way how to use in the enterprise for all the different purposes agentic tooling itself. So it brings a level of credibility.

Most organizations are actually not even at level one. Most organizations are still level zero. They haven't purchased the tools yet because they've never done the investments. It's just in the last six months when people really started to understand that this was really happening. Previously, based on all the different data points, people were skeptical. Now that skepticism is gone, they start investing, but the only thing what they're able to do is go out and make those purchases. That's why probably the revenues of these companies are skyrocketing right now, but the adaption is very, very difficult. So we are going out with the blueprints, with the run books of how to make the transformation with the educational materials, understanding how to actually go through step by step the change process, understanding how to mature engage by engagement, because it's not a top down, I would say, big bank.

It is happening. You have to do it project by project, going step by step. And as you are maturing these engagements, you can go to the next level. We have examples, and we can actually show how you are able to execute that in an organization such as EPAM at 60,000 people's scale. And that's very unique.

Elaina Shekhter:

And that's why they're called foundational services for us.

Moderator:

Let's go here and then here.

Jamie Freeman:

Thanks. It's Jamie Freeman from Susquehanna. I was revisiting my notes from Dmitry's talk about the four reasons to need EPAM; enterprise complexity, engineering mode, agent building, autonomous enterprise. If I mess those up, I apologize. My question is, if those are the reasons to need EPAM currently, I'm wondering, does it change the relevance of the global delivery footprint, and does it potentially argue for a bigger onsite onshore presence?

Balazs Fejes:

That's a great question. I think what we are seeing right now is our clients and enterprises, the same time they're trying to mature AI SDLC, mature the engagement model, mature the maturity of what they're doing. Same time they are executing in parallel other strategies, such as moving to GCCs in India or other locations. And they're coming to us, how can they upscale

They're existing so called legacy GCC with new skills. How can we help them to increase their internal efficiency? Just the other week I was talking to our client when making this pitch, they are actually expressing their need that can you engage with EPAM, with the EPAM skill, globally to tackle their own internal legacy? And their own legacy is not on site, their own legacy is it's a global footprint with different GCCs in different countries starting from India to Spain to, in this case, it was Portugal and Slovakia. That's where engineering is happening today and you need to meet your clients where their engineers are. So for us, we don't foresee that. And actually later on you are going to hear on the panel how we seeing all these things play out in each and every different geographies where we are.

Surinder Thind:

Surinder Thind with Jefferies. Following up on earlier question about the client journey and going from level one to two to three, and I think FB, you mentioned that maybe a lot of them are even at level zero. Can you maybe talk about the propensity of clients to move away from level one in the sense that, if the models continue to get better, we look at the journey over the last couple of years, would a client not want to continue to try and do more themselves, especially if the models continue to scale at the current pace? And are we in a situation where we have to wait until maybe there's a more maturing of the technology before clients move to level two and three? Or what gets them across that line? Because it just seems like industry demand remains relatively tepid.

Balazs Fejes:

Thank you much. It's a good question. I don't... Okay. So I think the models are maturing very, very rapidly. We all know that the capabilities. Also, the price point of the certain level of capability maturity is continuously dropping. For different business scenarios, business cases, you need different level of maturity. Depending on your price point of engineering, depending on the business case you would like to use AI for, there is different entry points. It might be possible that due to tokenomics, today for one company, this is affordable and/or actually economical to deploy AI today, or some decides to wait a little bit later while let's say the models mature or the cost drops, because there's two things happening at the same time. Newer models going to come and enter at the same level of price points where they are today. Old models continue to become cheaper as the token price, execution price, inference cost for all those models are dropping.

So some people are start deploying and actually actioning on this as they reach a certain entry point, and some people are waiting for newer models, as you were saying. But maturing, going through in a maturity model, it's not really, I would say optional. In order to get access to the capabilities of the model, you have to go through this maturity. So, one way or the other, if you want to tap into the power of the models, you will have to go from one to three. You're not going to get the benefits at level one. Actually, probably you're going to, as the models continue to evolve, you will be continuously even more disadvantaged by staying on level one. I don't know if it makes sense, but that's probably the right answer to this.

Moderator:

We have time for one more question here in the front, please.

Jonathan Lee:

Jonathan Lee from Guggenheim. Thanks for hosting. FB, you mentioned different price points as it relates to models, but can you expand on EPAM's pricing strategy overall as it relates to how your new go to market and your AI native approach impacts your pricing strategy going forward, especially as you balance agents versus perhaps higher cost team structures given talent scarcity?

Balazs Fejes:

Jonathan, thank you much, it's a great question. I think as you saw from our results, we continues to be predominantly in a timely material model and we actually also communicated to you that we were, in Q4, we saw very successful... We were successful getting rate increases from our clients, which actually indicates to us that the clients are receiving benefits of the more value which we deliver to them in the TNM model. But also I have to tell you that most of the times, the tokens are paid by our clients because we are operating in the client's infrastructure due to security reason, due to data confidentiality. And in that infrastructure, the clients are the ones who are deploying the models and they paying for the tokens. Going forward basis, as we are migrating away or transitioning away from time and materials to more advanced capabilities or more advanced contracting models, we will be seeing that it's going to be part of our commercial model.

We're going to factor in the price of the tokens into our model itself on top of it or maybe on a transparent way, it's a work in progress, how are we going to charge our clients the tokens because all the tokens are... The price is very volatile, so it's very difficult to figure out how to price it in at this point of time, but we expect that once we are more in the fixed price or more advanced models, the cost of compute will be included in our price. Last but not least, I think one takeaway that our AI native projects and revenues are operating at higher profit levels compared to EPAM average, they're more profitable.

Moderator:

And that wraps the first Q&A session of the day. We're going to take a break and reconvene here at the bottom of the hour, so 10:30 for those that are attended in virtually. For those in the room, please enjoy some refreshments and drinks, and then we'll get back to our seats here. When we come up next, Arkadiy Dobkin, our executive chairman will kick us off getting into our engineering DNA. Thank you very much.

Chris Michaud:

People are probably familiar with the MIT report that boldly states 95% of companies are getting zero return on their AI investments.

Eli Feldman:

You don't have 10 years. You have two, three, maybe.

Angela Masella:

I think AI will be transformational for the clinical experience in surgery. I think it's going to improve patient outcomes, it's going to reduce burnout and burden on surgeons and nurses and respective teams in the hospitals.

Shan Jegatheeswaran:

At least 90% of the AI projects that are rolling out are failing within companies, and that's because it's an organization and a people adoption problem with AI.

Chris Michaud:

We appear to be the anomalies, I think is really cool. We have a client base that is beating the trends right at the forefront of it.

Elaina Shekhter:

Taking something as nebulous and as confusing and sometimes scary as artificial intelligence and all the hype around it and turning that into examples of really meaningful programs that EPAM is either in the middle of or fully executed is the vision.

Paul McCormick:

We really are starting to unlock useful, tangible results for our clients.

David Billings:

These all go far, far beyond POC. These are scalable deployments in AI that are really delivering tangible business value for our clients today.

Alex Stassen:

We engrain it in all of our projects, it's basically nature and fundamental to what we do. Trying to improve things and make things better.

Claudia Fenske:

EPAM is a fantastic partner for us actually on the sustainability journey and also building our global innovation strategy.

Michael Callahan:

We've leveraged the EPAM partnership with their expertise. Put our first foot forward, has been a huge benefit.

Azhar Kholwadia:

EPAM have been a great delivery partner for us, both in terms of challenging us to making sure that we push the boundaries and making sure we're getting the basics right as well.

Kevin Labick:

What I think people get wrong about AI, is that it is there to automate tasks and remove humans. It's going to be much more of an exoskeleton, so it's going to enhance people's capability. It's going to make them faster, it's going to make them smarter, it's going to improve decision making.

Elaina Shekhter:

Artificial intelligence in many ways is a compliment to human intelligence and something that we should be looking for to actually propel enterprises and to propel the enterprise of humanity forward.

Arkadiy Dobkin:

Hello, everybody. Good to see many familiar faces here. So I'm Arkadiy Dobkin, executive chairman and founder of the company. Been here for a long time and past CEO position to be in September of last year as you know. I think being here for a very long time and hearing the previous conversations and Q&A sessions where we actually try to answer very, very difficult questions and present the picture which conflicts, not in very simple terms. We engineer our presentations well, and I probably, based on the years, have a little bit more holistic and casual conversation today, but there are three key messages which I think important and I would like to concentrate on this, that engineering excellence is still very... Sorry.

Why Kevin was there? Engineering excellence is critical differentiator and in AI age, it's even more important to cover through entire implementation cycle. I think history matters and similar in previous ways, I don't think it's going to be a revolution, it's going to be evolution for multiple reasons, and I think history is important to remember. And I think similar like in the past, the human talent will be the critical differentiator. Everything else will become eventually equalized and become more commodity, but actually the people who deliver in the last mile will be critical. So with this I would like to, for a couple of minutes, go back to the history and explain, at least for some new people in the audience, that from the very beginning, EPAM was slightly different than other major players on IT services market.

Our first clients were software companies, and for the first 10 years, 100% of our services were focusing on building products for software companies, very, very different business. The second 10 years we started to work with digital natives, Google's experience, Epics games of the world, and actually helping them to scale. At the same time, you understand that this 20 years of our first years of existence actually established very different DNA, very different processes, very different talent selection than majority of the industry.

It's important and it's become important after our IPO when we grew very, very fast when we are able to address the demand of completely different skills. So, I'm using the same slide which is already in FB presentation because the question which you asking and we are asking ourself, is it still important? Is this engineering DNA still going to be differentiator with all this noise and rumors and credible people talking around us, how code is over and maybe code is over, but what about engineering, but maybe engineering is over and what the next model will bring in all of this. With this, I would like to add opinion on one more expert and I'm not going to read the slide, but please read it.

Or even in short, the author said this, "Programmer is about to share the fate of the dodo bird. By the end of this decade I foresee massive unemployment among the ranks of programmers, system analyst and software engineers." And it was published in this book in 1992. It wasn't published by somebody, it was published by Ed Yourdon, who was the father of structuring programmer and critical person in creating object-oriented programming, and he was a visionary and one of the top 10 computer scientists of his generation. Why I'm saying visionary, because this book was published in 1992. His status was that offshoring and new programming methodologies will kill American programmers.

Think about 1992, the whole offshore in the market was \$100 million from about \$100, \$200 billion global IT. He was brilliant, but three years later, four years later, he published another book called Rise & Resurrection of the American Programmer, because he admitted that he hugely underestimated entrepreneurial drive, Silicon Valley innovation, growth of economy thanks to internet, and one more point, complexity of the enterprise. He hugely underestimated that, and he was wrong in his first book. So this bring us to actually the EPAM lifecycle, the history, the ways from foundation to going through the crisis. We started in 1993 actually at some level inspired by his book about offshoring and we ran into internet era and you know what? That is this point, the skills which have to deliver this new type of applications didn't exist. You cannot go to the market and buy, each of this internet, including actually created the hive, the programmers, C++, C or C++, like real people don't need anymore, HTML coders will do it.

And then it was disappearing because each time complexity was underestimated. Then we came to era of cloud, mobile and data and same stuff. And we as an engineering firm, we starting to build our own platforms. It's been mentioned as TelescopeAI, we will talk a little bit more about it, but we also engineer not only digital platform, we engineer our educational learning platform as well, because we cannot find these people. We have to find the right candidates and develop them. And that's what we did during this second era better than anybody else. That's why we were growing. This is where it was impossible to predict what type of new applications going to happen. Think about it, we're talking about AI impact on existing type of applications, and that's what underestimating coming from, because the main change going to be in the future, we don't know yet what it is.

And now in AI era, and that's what we were covering before me. The pattern again across all of this was that every productivity went from 4GL to, object-oriented to, open source, low-code promised to reduce builders' demand and in practice each time the lower cost went, the more market expanded, more opportunity, more cheaper way than new and this new were growing like a snowball. So, that's why if you think about, in addition to everything else, what's happening with regular productivity, which we kind of focus in the first part, entrepreneurial drive of people, innovation levels of something which you have no idea about it today, potential economy growth with AI and making everything cheaper in intelligence and enterprise complexity, which I don't think I need to explain.

Even with the comment before that some companies even didn't buy the tools, the silos of knowledge, so huge in corporations, you work in there, you know. AI not going to bring any benefit unless it's all uncovered together. So with this, in the AI era, it's going to be actually growing demand. I'm pretty sure about it, not theoretically, for very real. We enter the market when traditional software could never afford to serve before.

The last mile become very critical built differentiation. Everything else will be equalized, and we're going to address level of complexity. We have no idea about it today. Similar like think each time 10 years back, think from AWS to Amazon Bookstore. Can you imagine all this happening? So, I think the shift of the bottleneck going to be up and up. And the last 20%, which usually taking 80% of the big engagement because of complexity, they will move even to higher average and the 80% was relatively easy. Yes, it would be much more easier to do.

So, I think at this situation, the people who delivering this last mile, leading this last mile, which would be very, very scalable, it's a key differentiation and these people who has to work in ambiguity, in unknown, think very quickly, because AI making everything all that very, very fast. And I think bringing another current authority, Boris Cherny, probably you saw his podcast, he was talking about exactly importance of engineers, and this is what Dima was talking about it, this full-stack agentic engineer who can coordinate, people who can orchestrate. And if you think that it's very new thing, that it's a mistake. I think that's exactly EPAM was benefiting from this type of people in complexity during the previous

decade. This is how we differentiate ourselves in the past. The point was that with talent we built, sometimes these type of people were not even in enough demand.

We were putting them on some coding positions, but we understand with our insight to the systems and to our educational learning process, which we going to talk in a minute, how to identify them, how to develop them, and how to scale them historically for the last decades. So, key takeaways, we're probably really underestimating the scale of AI driven market expansion and the complexity of enterprise. The second engineering matters and Anthropic people saying this as well, coding simple engineering becoming much more sophisticated. Think about it like new terms which come in almost each couple months. Prompt engineering, okay, this is legacy. Context engineering, intent engineering, I don't know what will be tomorrow. And right talent and this was 30 years of our focus. By the way, Dima was presenting here, he was graduating from computer science, but he went through our educational six months bootcamp before he started to work at EPAM.

That was happening 20 years ago, and this is what happened today. Thank you. And I would like to invite Sandra, our Chief Learning Scientist, and Aleksey, who is the head of Engineering Excellence, actually to bring much more details on what I was sharing with you.

Aleksey Didik:

All right. Good morning, my name is Aleksey Didik. I'm a head of engineering excellence AI.

Sandra Loughlin:

And I'm Sandra Loughlin. I am EPAM's Chief Learning Scientist.

Aleksey Didik:

Arkadiy just showed us that we need the right talent. Dmitry and Adam gave us a glance of a full-stack agentic engineer, and it was even a question from the audience, who are those people, so let's take a look. Full-stack agentic engineer is not just a new role which build from scratch for AI era. It's a evolution, it's built on a foundation of narrow specialists available in the industry. In EPAM narrow specialist, there was already better because of our engineering DNA culture and excellence. Now we need to extend this foundation with a full-stack development ownership of application layers across all technologies.

We need to deepen it with the understanding of AI tooling and also understanding of AI native workflows, capability to orchestrate agent fleets across all stages of development. But how we can even approach this new talent profile? How we can build it? And we are doing it by breaking it into skills. Skills which are becoming less prominent and important. Skills which are still need to stay because they still might have, and the skills which are emerging and rising because they're becoming a new must have.

How we build those talents? To build those talents, we have our educational program with universal coverage. We build this program using our own proprietary courses. We do not want just to use materials from the market, because we believe that external knowledge need to be processed and pass through the lenses of EPAM experience, our experience to deliver AI native work. We combine it with a formal education and informal education, running a global AI conference last year, thousands of people, 45 countries, because it's important to build horizontal connection with people, between people, so they can exchange knowledge, learn from someone next door. We run masterclasses together with our partners from Amazon and Microsoft. So, it's a very good program, but is it enough?

Sandra Loughlin:

If that seemed common to you, it is. Percent of employees who've gone through courses, who's clicked through what? How many classes do you have? Those metrics are table stakes. Worse, they're illusions of competence. Training people is not a strategy, and it's certainly not a differentiation. Leading in the AI services market requires going far beyond those basics. Building the AI native talent that you've been hearing about today is actually a three-pronged challenge. It starts with identifying the skills that are in demand today and critically the ones that will be needed tomorrow.

Exactly the kind of skills that you heard about this morning from Dima and Adam and Ark. But development really isn't about training, people can train and learn nothing. And most people learn from informal things like reflection and practice and getting feedback. For development, there are two key things to learn about. One is motivation. Can organizations drive their people to learn even when it's hard or not fun? And the second is validating the skills. If you can't use those skills in production, it doesn't matter. But the most critical metric for professional services organization is actually deployment.

Can you put the right skills and the right combinations on the right client projects to create value? This three-pronged challenge fundamentally shifts the metrics that matter for talent development. Instead of focusing on number of people trained, the companies that grow people and those that invest in them should be thinking about different metrics altogether. How quickly can you sense the right skills? How fast and how thoroughly are people upskilling and demonstrating that they're using those skills in practice? And critically, how quickly are you staffing the right people to the right client projects? In this era, the future will be had by those people who focus on those metrics.

You're not going to be surprised to hear that is who we are. For years, you have heard about TelescopeAI, EPAM's proprietary 30-year homegrown, in the making system that is focused on people and the backbone of our business. Today, you're learning why we keep talking about it, and that's because Telescope was purpose built to do exactly these three things. In a world where organizations know more about the chairs in their buildings than the skills of the people who sit in them, EPAM has built our business to know exactly what we need, who we have, and where best to put them. And for a company whose business is people, that knowledge is competitive advantage. Before Aleksey shows you the metrics that we track, please know that some of these numbers are operational and proprietary. It's why you're not going to see hard numbers for everything.

But most importantly, you can't interpret these numbers without a context and the industry is just not there yet. They're not tracking the same numbers that we are. We believe that they will get there, we think that it is inevitable and we are welcome to... We're excited about that the day when they do. But until then, we're going to offer you a glimpse into how we treat talent as a business asset.

Aleksey Didik:

So we have three functions, sense, develop and deploy. Sensing starts from market and industry. Industry first, our practice leads carefully process all information coming from the industry on what is going to happen in the next months and years. But we do not just listen, we process and converge this information into skills, skills which are retiring, retaining, or rising, because it drives the development of our learning programs. And the same skills, I use it to understand demand on the market and predict demand on the market, because we know how much new positions our clients need with AI ready skills and I should say that this demand is quickly accelerating. But it's not enough just to sense industry in the market. We need to sense our people to understand the why, how we can provide them to our clients.

And this sensing is definitely not only about how many training modules they completed. This sensing is about the way how they converted this knowledge in their real work experience and build real skill. So we combine evidence from different sources, from the complexity of delivery,

... of real work they done, from reviews and assessments, how quickly they learn. Endorsements from the peers and it's all together creates a universal standard applicable across all our global workforce, across all our countries. And I should say that we are sensing that we have enough AI ready engineers to cover all our client needs. And now, we need to deploy, but we don't want to deploy people just based on availability, we want to deploy people based on the verified mastery, based on our ability to provide feedback purpose engineers to our client. That's why our Terrascope AI and proprietary AI driven matching model uses 25 different attributes to find the right people with the right skill for the right project of our clients.

And results are evident, roughly 80% of positions with AI skills at this moment are staffed within seven days, and the rest doesn't take much longer. But it's about speed because you can staff quickly, but this is the quality and quality is here. Our NPS in comparison from 2024 to 2025 grew on plus 4%, and take into account that our NPS is already above industry average.

Sandra Loughlin:

The fact that we have hundreds of university partners is good, but the way that we use them is actually what matters. Instead of relying on faculty to keep pace with AI or hope that they listen to us and change their courses, we learned long ago to engage directly with students like Aleksey, like DEMA, using our own instructors and our own proprietary coursework, the same coursework that we use with our people inside. This means that students in our pipeline are trained on our evolving definition of AI talent and they're tuned for local client demand. And because we've invested in them and because we have built relationships, EPAM gets to snap up the best talent before anyone else. This model is not new for AI. It's how we've operated forever and it's not something special that happens in one geography.

We built this model in Eastern Europe and then scaled it to all of our major delivery centers around the world. And that's why, as you will hear from Larry and Vic, we can have a standard very high for engineering talent anywhere we go in the world. Four years ago, I stood here and said young EPAMers can't be hired. They can only be built. That has not changed, but the value to our business has. In a world where AI native juniors aren't available anywhere on the market, EPAM has a global pipeline prepared for local client demand on day one.

Aleksey Didik:

Results are evident. We have engine and it's running. We are sensing market in an industry, which allows us predict what's going to happen next and how much people we need. It helps us to build the supply depths through our learning programs, combining formal and informal education, and then verify those skills. Be sure through the real delivery through the production. We are able to deploy our people, feed for purpose, right skills, right people for the right project. We can do it quickly and keeping quality. And deploy function goes back to the sense, and that's the way how the feedback loop completes. That's why how the whole engine is working. And we can not only today create several full stack agentic engineers and many of them. We can do it tomorrow and the day after tomorrow, just because this engine is what drives this success.

Sandra Loughlin:

The market commonly conflates EPAM's success with our historic footprint in Eastern Europe. That has never been correct. Our roots in Eastern Europe set the highest expectations, but this talent engine that

we've been telling you about all morning is what has scaled those expectations to EPAMers worldwide. In other words, our ability to provide clients with the best engineering talent is and has always been due to what we're showing you today. A platform and business model specifically designed to sense, develop, and deploy cutting edge talent. As you have heard through the years, what defines cutting edge has changed, but the success of our model has not. We have maintained world-class talent in every era and in every area of the world. From this perspective, full-stack agentic engineers are not a new challenge for us. They're just the next frontier. Competitors are scrambling right now to recreate telescope AI and our skills-based organization as they should.

Meanwhile, we will continue to refine our engine and use it to help our clients get ahead. As you've heard all morning, AI is changing and expanding, not diminishing the need for expert engineering talent. In fact, AI has only made the need for that foundation stronger. Value follows constraints. And in a world of AI, one of the biggest constraints is human skills. To meet the moment, IT professional services organizations must sense what those skills are, motivate employees to develop them and deploy the right combination of skills to the market. That's it. That's what it takes to lead in the IT services market. And in this, EPAM has a 30-year structural built in advantage. Thank you.

Aleksey Didik:

All right. And I want to welcome to the stage Chief People Officer, Larry Solomon. Thank you.

Larry Solomon:

Hello, everyone. It's great to be here. I'm Larry Solomon, as you just heard, EPAM's Chief People Officer. And I've been in that role for coming up on 10 years now. Now, earlier in the session, you heard Adam Auerbach comment that he's been in and around the IT industry for 25 years. I've been in and around the IT industry for 40 years, approximately. Now, I know what you're all thinking. I know what you're all thinking, especially those in the front row. There's no way that that guy up there has 40 years of work experience under his belt, right? All right. I see a few... Okay. All right. Thank you. Thank you for that. But to get more serious, I first want to thank you for coming today. It's much appreciated. I'm going to quickly take you through our global talent and delivery model that has evolved over the past few years, and why that evolution has made us stronger, more resilient, and better positioned than we've ever been in the history of the company to support our clients all over the world.

Our delivery model today is not only stable, it's optimized. It's optimized. We've built a model that's more balanced, global, and flexible than ever before. And that foundation has been what's let us scale rapidly. Move talent where we need it, move talent when we need it, and deliver for our clients no matter what in the heck is going on in the world around us. And we've had a lot going on in the world around us, as you all know. Now, I like threes so there are three key takeaways that I'd like you all to take away today. First, we've successfully rebalanced our delivery base. The 2022 invasion of Ukraine was a catalyst, an unbelievable, almost unreal, incredible catalyst that accelerated our move into nearshore and offshore hubs without sacrificing client continuity and the quality of our delivery to our clients. And let me assure you, you can't learn that from the fine educational institutions that we have within a few miles from where we are today. You can only learn that by experiencing it, by living it, and that's what we did.

Second, it wasn't just about moving people. It was about de-risking our entire delivery execution model. And we've built a rock solid culture of resiliency, a rock solid culture of resiliency. Resiliency first. Finally, we're now truly distributed around the world, harnessing the lessons that we've learned from crisis. And fortunately or unfortunately, crises have become a core competency of ours, but it's helped us create a global engine that provides better access to top talent. And you've heard about the importance of top

talent, and you'll hear about it today after me. And this is now a durable competitive advantage for our enterprise.

Now to understand where we are today, you need to look at where we came from, where we started. Back a few years ago, in late 2021, we were already in the process of diversifying, but as many of you that have followed us know, our footprint was still quite heavily concentrated. And at that time, 59% of our delivery professionals, 52,000 strong at that time, were based in three countries, and you probably know them, Belarus, Ukraine, and Russia. And while this served us well for many, many years, it represented geographic concentration risk that we knew, we knew we absolutely had to address and deal with.

Let's fast-forward now. A few months ago, the end of 2025, look at the shift in the circles on the map here. Our delivery force has grown to almost 57,000 production professionals, but the distribution of where they are around the world is like night and day. We've reduced our concentration in Ukraine and Belarus by 38%. And at the same time, we aggressively ramped up other parts of the world like India, Latin America, and Western and Central Asia. This is what optimized and balanced looks like. We're no longer dependent on any single geography, on any single region. We're much more regionally balanced and diversified today.

Now, I'm extremely proud, extremely proud of how fast our teams pivoted. We have a saying that we use quite often around the place, speed kills when you don't have it. Speed kills when you don't have it, and we had it, and we still do today more than ever. We relocated people, we opened up brand new locations, we expanded our mobility programs, and we built talent pipelines in new markets, all at a pace that no other company could match. This speed and agility is absolutely part of the core of what and who EPAM is today. Today, the model is a real strategic advantage for us. It helps us deploy the right skills to the right clients in the right places at the right times. It improves our cost positioning, it expands our access to top talent, and gives us the geographic flexibility that is extremely, extremely difficult to replace.

We're poised to capitalize on this more balanced footprint. We've de-risked our execution with stronger business continuity. We have better and faster talent access from a much wider pool of specialized and unique skills that our clients are demanding from us every day. And as you'll hear from others, we're integrating AI-enabled optimizations across our company to improve our cost profile and utilization across the regions. Ultimately, the model that we're talking about here supports an important 24/7 or follow the sun delivery cycle, and that creates faster iterations and turns for our clients. Today, we find ourselves even faster, safer, and more globally diversified than at any point in the company's history.

Now, some of you may recall the concluding comment that I made at these events in prior years, and I'm going to say it again today, so it's okay if you don't recall. I'm going to say it again today because I believe it's more true now than it ever has been. We hold the cards that we've been dealt and I wouldn't trade in our hand for anything. I wouldn't trade in our hand for anything. With that, I am excited and delighted to hand over to my good friend, my colleague, and frankly, one of the most talented and smartest leaders that I've had the privilege of working with in my 10 years at EPAM, Victor Dvorkin. Thank you very much.

Victor Dvorkin:

Thank you, Larry. It'll be hard for me to prove it. Good morning, everyone. I am with the company for 28 years, in the role for 10, and I will try to prove as a scientist that what we built is actually one of the best delivery engines in the industry, and that it will be actually rewarded by a native wave. Let's start. First, clearly enterprises got access to really powerful models right now. There is no doubt. And what it means that the demand for AI work will increase because they will understand better, they will want better,

and they will ask us to do more. We spent the whole morning talking about enterprise complexity, legacy systems, complex platforms, integrations, hallucinations. We didn't talk about that. And real operating pressure, which they have.

This has been our environment for 30 years, large transformations, regulated industries, complex platform engineering, and also Google scale product engineering at speed. This is actually how we won the digital wave in the past. Great models for us, in our opinion, is an opportunity because this is what makes our delivery engine actually unique and that what will make AI run the enterprise. I will show you what's probably one of the most complex slides, so bear with me. Larry spoke about allocation strategy. Allocation strategy is a serious advantage. I'm showing you an example of a large client. They have a headquarters in U.S., a headquarters in Europe, a local GCC, and a Latin American subsidiary. Think about the complexity.

We have nearly unlimited flexibility how to configure this type of engagements, meaning the most strategic, regulatory, and pricing needs of our clients. And we, as Sandra and Aleksey said, we sense, develop, and deploy our talent. But I will add, we also continuously assess our talent globally and unify our skills globally through global unified assessments. Every engineer in order to get promoted need to be assessed from an engineer actually to NSVP. We just finished the cycle right now. And this consistency is a key.

I will complicate the slide more. Data practice. As you see, it's global, it's in every location. And this is our major strengths as well as cloud, digital and product engineering, SAP, Salesforce, and other practices. This horizontal capability is massive with thousands of professionals, leadership, competency centers, partnership with cloud and platform providers, methods, tooling, training, certifications, and operations. I will complicate the picture more. I will add verticals. By the way, talking about hallucinations. In healthcare, they produce unsafe outputs. In financial services, non-compliant responses. In supply chain, the output looks great, but think about the world. It will not work today. And that's why vertical is super important. That's why T-shaped talent matters the most for AI adoption. And that's why organically, with our clients together and through acquisitions, we are continue to develop vertical capability. Elaina and Eli spoke about consultants, I will talk about engineers. Think about healthcare, life science, financial services, media, gaming more. T-shaped talent makes system work with AI.

Think about now 3D pictures we just covered, right? It is absolutely unrealistic to run this manually. That's why for so many years we developed our digital ecosystem covering talent, skills, knowledge, technology, and I will show you delivery. This is a delivery view of a delivery engine. See, it looks fine. Green, but I actually would say it's a bit too much green. And that's why we will drill down on a specific account. What we can see through AI powered systems, we can now instantly understand what the risk we are doing, what type of analysis we should have on how to remediate the problem. And this also accumulates our reusable delivery knowledge, which helps us with estimates and with many other things.

The same view through the agent. You can see that agents and teams can query it from cloud code or from other environment, or actually through the agentic factory. And this is coming the VIKbot. Yeah, it's me. I forgot the glasses. How to explain VIKbot? That's very easy. If you have a project, VIKbot will come to you. That's how I was explaining it somewhere in the kitchen. Yeah. But most interesting, we both also work on our newly baked AI factory, which we can demonstrate today. It helps to validate our proposals. It helps to check the estimates and the value if we bring it to our clients. With that, I will leave you with a few things as [inaudible 02:30:36] we have really advanced capabilities, global scale, consistent standards, T-shape expertise, and AI run platform, which runs our delivery engine, which runs enterprise, and which will be ready to win the AI wave.

And with that, actually, I have one more thing. To feel the organization heartbeat, we prepared a panel with leaders building and running teams around EPAM. And I would like to welcome Amit Singhal, Head of European Delivery to introduce the panel. Thank you.

Amit Singhal:

Thank you, Vic. By the way, that VIKbot is real. It's calling me every day. It's much nicer him calling than Vic calling me. My name is Amit Singhal, SVP and Head of Delivery for EPAM in Europe. Joined EPAM roughly 10 years back, but who's counting? As Vic said, I'm going to host a panel for you so you can hear from some of our regional leaders. Please join me in welcoming them. Okay. How are you all doing?

Stepan Mitish:

Great.

Martin Mendez:

Very well.

Enver Amdiy:

Yeah.

Srinivas Reddy:

Amazing.

Amit Singhal:

Okay. Sitting at the far end is Enver, my partner in crime in Europe.

Enver Amdiy:

Hello.

Amit Singhal:

Maybe we should have sat together, but it's okay. Enver heads our business in Europe. By the way, congratulations on getting to the top place in Whitelane survey in Europe.

Enver Amdiy:

Thank you.

Amit Singhal:

You and I both seeing an interesting trend in Europe where business is growing much rapidly than we hope. We like it, but we had hope.

Enver Amdiy:

We do.

Amit Singhal:

Yeah. Across sectors and industries, hope that's not an accident and there is a strategy behind it. Keen to hear from you what's going on there. Next to Enver is Srinivas, Srin, as we fondly call him. You and I join roughly the same time in EPAM.

Srinivas Reddy:

That's right.

Amit Singhal:

Your mission was to build a different kind of India for EPAM in the region and scale it. And it's one of the largest location now, so you must have done something right, Srin. Congratulations and welcome.

Srinivas Reddy:

Thank you.

Amit Singhal:

You took a long flight to get to Boston.

Srinivas Reddy:

I did.

Amit Singhal:

Through a narrow air corridor.

Srinivas Reddy:

That's right, over Iran.

Amit Singhal:

Good. Martin.

Martin Mendez:

Hello.

Amit Singhal:

New kid on the block.

Martin Mendez:

Hi.

Amit Singhal:

Very new to EPAM leadership team. Martin, you were born in?

Martin Mendez:

Argentina.

Amit Singhal:

Argentina. You lived in Brazil and Mexico, so you know the region a little bit?

Martin Mendez:

A little.

Amit Singhal:

And you're enjoying your journey so far with EPAM?

Martin Mendez:

Very much.

Amit Singhal:

Okay. Much like Srimi, Martin joined us to consolidate our investment in the region and create one team, which can be plugged into a global delivery model that Vic talked about. Welcome.

Martin Mendez:

Thank you.

Amit Singhal:

And last but not the least, Stefan. Stefan leads our teams in Ukraine and all of us know the board is still ongoing and everything that throws at Stefan and his team and continue to do good work. On behalf of entire EPAM family and our clients, Stefan, can I say thank you. Thank you.

Stepan Mitish:

Thank you. Thanks for having me.

Amit Singhal:

Thank you. That's amazing. Thank you all. Let's get into it. There's a lot to talk about, but let's try and focus on few things. I'm very keen to talk about resilience of EPAM delivery, how GenAI adoption is going across complex enterprises. How do we balance our global mindset, but equally Srimi, for example, in your case, working locally with GCC. Enver, if I could start with you. You and I know Europe is a melting pot of cultures and languages and it's fragmented. There are complex cities. How do you lean on Big EPAM to deliver best in class services for them?

Enver Amdiy:

Thank you, Amit. Indeed, Europe is a great mix of cultures and languages, the place from where global companies, number of global companies are rooted from, but also significant market for a number of localized businesses. As Vic and Larry said, we as a company invested heavily into reinforcing our global delivery engines so we can serve clients from everywhere in the world. At the same time, working for number of years with our clients, shoulder to shoulder, we accumulate a significant amount of industry

knowledge. Today, I believe our winning combination is in market, Western European talents for client proximity, senior leadership and regulatory alignment, Eastern European teams or nearshore teams for in depth engineering talent and for business knowledge, and offshore teams for technical talent for scale and cost-efficient execution. If you add on top of this mature governance and now AI powered productivity gains, then you get an engine which is both resilient and highly efficient.

Amit Singhal:

Right teams supplied in the right proportions at the right time to the right situation is the winning combination.

Enver Amdiy:

Absolutely.

Amit Singhal:

Martin, if I could come to you next. Pretty similar to what Enver said, but we know Latin America is not Europe right, so you know it better than rest of us. What's your sort of winning combination in the region, both for your local customers, which you brought with you from NEORIS and plus EPAM global customers?

Martin Mendez:

Thank you, Amit. Pleasure to be here. It have been almost a year and a half since NEORIS became part of EPAM. And I think that due to that, we have a much stronger EPAM in Latin America and I'm glad to see that. And the reason is that because we now have great engineers plus all the AI platforms that EPAM built. Plus now we have a strong leadership team based in the region that knows the regions for a while. And we also have a strong install base of customers that were customers that were born in Latin America or they are playing in Latin America. As Larry, I like the threes, but I need to tell you four things. Four avenues that we are pursuing in Latin America.

The first one is that the one that EPAM was pursuing since the beginning is how to supply or how to do nearshore from Latin America to the U.S. And that's something that we continue growing and developing more capabilities and those employees or those consultants that are working with us are also going to be able to serve our local customers in Latin America. The second avenue is how to bring those new technologies that we are building on a global basis to our install base in Latin America. And we are very proud now to have, I say, the Navy SEALs that will help us to expand our presence in the region. We have the platforms.

Amit Singhal:

We surely need more of them.

Martin Mendez:

But this is something that in the past story of NEORIS we are not having and so I'm very proud now. I think that we are going to be very successful on bringing those things to Latin America. The third avenue is that we also have global customers that we are having operations in Latin America, but we were not able to serve in the past. Now we are working with them, we are helping them to deploy those technologies in Latin America. You know that Brazil, for example, is a very complex country and we do have an operation there and we are getting to know them and expand that relation.

And then last is we also have, or EPAM has a very installed and mature relationship with a lot of the large technological partners and they were demanding and kind of how EPAM was able to go with them to the region. And now we are partnering with all of them. And I think that that's going to be another fourth avenue that we will explode in order to grow in the region. And I'm very happy to see this combination as a winning one.

Amit Singhal:

Yeah. You said something very interesting that if you want to be resilient in the global world, one of the critical item is to have strong leadership based locally. Okay, great. Stefan?

Stepan Mitish:

Yes.

Amit Singhal:

I've got so many things I want to ask you, but time is limited. I can see it there. First of all, there is hardly a week goes by where I don't come across a customer who's been working with your teams in Ukraine. And all I hear is great words, and I know it's not just sympathy. On the other hand, when I talk to your people, I see motivation, I see high degree of engagement. What's the secret sauce? What's going on in the middle?

Stepan Mitish:

Thank you, Amit. I think that's

That's the most common question I get asked. And first of all, all the credits should go to Ukrainian team. They are awesome, brave, resilient, and I'm really proud to be part of it. Now answering your question, I think there are several components that help us to be successful. First and foremost, I believe that we secured the foundation. So basically companies stood with us from the day one. They created a 100 million dedicated fund to help our people, their family. And there is an expression, you want your employee to take care of your clients, take care of your employees. And that's exactly what we did. Second, I think we focused on a purpose, not the pity. And that might be not obvious for people outside Ukraine. But for Ukrainian, the biggest motivation is feeling yourself useful. So you can protect the country in trenches or you can protect the country on economical front.

And I remember a conversation with a client who was kind of reluctant to open work for Ukrainian teams just out of sympathy. And he told me, "Stephan, I cannot push your people toward during wartime." And I told him, "While I appreciate your heart, but the truth is that our people have a tremendous motivation to work because it brings revenue, it increases taxes, it helps to protect work in places, it's industry, et cetera. And that purpose gives our people control while kind of everything else is volatile." And he was like, "Whoa, I didn't think about that from that angle." And that was an eye-opener for him. He opened work for our teams and they've been delivering for him ever since. And last but not least, I believe it's our results relentlessly attitude. So London Stock Exchange, our huge client. So they have a massive program of migrating hundreds of applications from on-prem to Azure.

And by the way, several previous attempts failed with other vendors. So just recently we completed a first migration of the application that was done by a small Ukrainian team with a little bit of sleepless nights, of course, with the usage of AI. And it was done on budget, on time and client's feedback was that it was the most seamless migration ever in his career. So at the end of the day, while the environment may be volatile, we prove that our delivery is a constant thing. So we don't just meet the standard. I hope that we set the new one that could be called resilient partnership.

Amit Singhal:

No, it's absolutely right. So I see this every time we are having client conversation about Ukraine. And as you said, if you want to help Ukraine, work with us. So great. Thank you, Stefan. Quick follow up. One other interesting thing we saw in Ukraine was very early adoption of GenAI. And in fact, some of the EPAM IP like CodeMie and ELITEA was born in Ukraine, which became part of EPAM AI platform. Again, what was sort of the driving force behind it?

Stepan Mitish:

Well, great question, Amit. So in order to understand how it happened, you have to understand Ukrainian engineering DNA. And to be honest, we've always been very fast adapters of everything new. So look at our ministry of digital transformation, our DM mobile app with the government in the mobile, with all the documents, services, defense tech, nearly cashless society. So for us, AI is not a hype. It's a kind of skin in the game. If you don't disrupt yourself, we're not going to win. We're not going to be successful in front of the clients. And yes, indeed, both tools, CodeMie and ELITEA, which are part of AI run platforms that Adam and Dima was talking about was born in Ukraine and by Ukrainians, which basically proves that we not just deliver despite all the adversities, but we also innovate. And yes, we see a shift in the engineering role from kind of how, which is code generation to a certain extent, to what and why, which is focusing on the complex client's challenges.

And here is an example from real life. Vadym who is a product manager of the CodeMie, it's an AI native agentic platform, he was doing a demo. And during the demo the app crashed. And instead of just panicking, he just went to CodeMie agent and described the bug and asked agent to fix it, run the test and deploy to production. We went for a 10 minutes coffee break, returned back, and boom, it's already fixed and in production.

Amit Singhal:

So like Boris from Anthropic said, "Cloud is coding Cloud."

Stepan Mitish:

Exactly.

Amit Singhal:

CodeMie is coding CodeMie.

Stepan Mitish:

That is exactly the level of maturity we bring to our clients. So it's not just about code snippets generation. It's about automation, the full cycle of software development. So at the end of the day, we believe that AI is a multiplier for human brilliance. And given all the companies we have and our strong engineering DNA, we believe that we're going to remain steady, innovative partner that our client trusts to continue solving the complex challenges.

Amit Singhal:

That sounds amazing, Stefan. And as you said, if you want to win, you have to disrupt yourself first.

Stepan Mitish:

Exactly.

Amit Singhal:

It's great.

Stepan Mitish:

Thank you.

Amit Singhal:

Srini

Srinivas Reddy:

Yes, Amit.

Amit Singhal:

We should talk about India. Not about the pollution and traffic and the population, but EPAM India. So it became the largest location in EPAM and in span of what, 10 years or so roughly?

Srinivas Reddy:

That's right.

Amit Singhal:

And clients tell us, and we see it ourselves, but client tell us which is probably a bigger proof point, that it's very different when they work with EPAM India versus the rest of the competition. What's behind the scenes' and story? How did you go about doing it? What's really different?

Srinivas Reddy:

Yeah, thank you, Amit. So I think we differentiated ourself in the India market by building a modern engineering company. And it was built on our EPAM's global engineering culture and hiring quality talent. We did this over 10 years and we did this very differently. Today I have very senior leadership teams in India that manage mature practices in cloud, in data, in data science, and now in AI and gen AI.

Amit Singhal:

So again, local is strong leadership. It's critical.

Srinivas Reddy:

Yeah. And I remember the first global AI workshop was conducted in Hyderabad. This was for a week, more than two and a half years ago. And we have all our senior AI leaders in Hyderabad, and most of them are actually in this room today. And when we were done, one of the OKRs that we came out was to make EPAM India the first AI native location in EPAM. And that for us really started with training our engineers. Today, our AI literacy in EPAM India is 90% plus. More than 70% of our projects leverage AI tooling, either our AI run or some agentic ID that is provided by our client. In addition to that, the teams in India have also contributed to our AI initiatives. We built the AI Ops platform that we leverage on all our managed services engagements.

Amit Singhal:

Which is now part of EPAM AI-run bigger platform umbrella.

Srinivas Reddy:

That's right. We also built an AI reverse engineering tool and agentic swarms for use on our modernization projects. And if you think about it, some of EPAM's largest implementation on CodeMie, on ELITEA, on cloud code are being run out of programs in India today.

Amit Singhal:

Yeah. Now we see that. We see AI adoption at scale with large enterprise in India. So well done. Thank you, Srini.

Srinivas Reddy:

Thank you.

Amit Singhal:

Quick follow up. Again, being the largest location, you play a very big role in EPAM. JIRA does diversification for US customer, European customers all over the world. But then you have other side of the coin as well, which is GCCs, which are rapidly coming up and building in India. So could you talk a little bit about it? And do you think GCC is a big opportunity for EPAM?

Srinivas Reddy:

Yeah, good question, Amit. So as you're aware, we today work with 150 plus clients in India, and we work with them in various different delivery models. The traditional outsourcing where the teams are exclusively based in India and hybrid. And we do quite a bit of work today in the hybrid model where we have teams in India, but we also have teams in one or multiple of the other locations. And like you said, in the local market we work with those global capability centers or GCCs. Today, we work with somewhere between 50 and 60 GCCs in India, and we've been working with them for more than 10 years. And over those 10 years, we've built strong local relationships. And today, I think in most of them, we are their trusted partner. And I think that's happened mostly because of our advanced engineering skills, our AI capabilities, which really compliments what the GCCs themselves are looking to build in India.

So the answer to your question is yes, Amit. I think their rapid growth over the last few years in India is actually an opportunity for us.

Amit Singhal:

So premium skills and proximity to GCC is what they're looking for. And sounds like we're winning there. No pressure.

Srinivas Reddy:

Thank you.

Amit Singhal:

Okay. So before we wrap up, there is one big question we have to sort of talk about, which is we see, and industry is talking a lot about it, that a lot of large complex enterprise clients are stuck in this R&D and POC phase and not really able to scale GenAI into their environment. We have seen some early success with these organizations. So can I ask both you and Martin to share some examples where you believe that we managed to unlock the key?

Enver Amdiy:

With pleasure. Martin, would you like me to start?

Martin Mendez:

Okay.

Enver Amdiy:

Okay. Yeah, indeed clients are quite excited and to a certain extent are under pressure. As you rightly said, they run multiple POCs in the last couple of years. And now they really get to see real implementations with real returns of the investments. So I believe opportunity is big and EPAM is very well positioned to capture it. The key arguments for this are, as you all heard, our top-notch technology excellence. Second is industry, knowledge accumulated over the years. And third is our early and very practical investment into AI. All of this helped us to form very concrete industry aligned points of view on how AI and modern platforms can transform our clients' businesses. And importantly, we didn't stay on PowerPoint levels. We went all the way and implemented industry-specific accelerators and our clients use it nowadays. Just to give you a couple of examples. First one was Swiss Re.

They used to produce Sigma Report, very well known in the industry. And we helped Swiss Re to ID, validate, and deliver Sigma Explorer, something that connects all resources, all publications, all data sets, and helps end users to talk to the data, do the data analytics using the natural language. So we developed the system using two EPAM accelerators, DIAL and QuantHub, and it went live in a very short period of time. And another important example is going to be one and one, or how they call it in Germany, Einstein's, major German telco. And they wanted to reimagine the way how they interact with the users. So we developed for them an agentic AI platform that today handles hundreds of thousands of end user calls. Not only helped to cut operational costs, but it improved the client satisfaction level. And again, we used our AI run transform platform to develop it. And the first agents went live just within several months.

Amit Singhal:

Sounds like we need to say good luck to our friends who are running traditional BPO industry.

Enver Amdiy:

I'll do this.

Amit Singhal:

Okay. Martin?

Martin Mendez:

I have three examples of Latin America, and this is Latin America. One is there's a large utility company in Brazil that is doing a big migration from a legacy system into SAP on the cloud. And there's a big need to migrate tons of data, but at the same time they are going through an M&A strategy towards acquiring companies in the region. So they were thinking about how to do it. We were competing with some of the local competitors and they were going from other traditional approach of migrating data. So we came with Run:ai and Miser as one of the platforms that we have. We've been able to prove them that by using this platform, we are going to be not just able to do it faster in the first time, but also have a repeatable agent that will help later in the future acquisition.

So that's one of the first cases, and it's very interesting. The second one is like, we have a large manufacturing company in Latin America where it's having cameras for the surveillance of the plants.

Amit Singhal:

Not for spying.

Martin Mendez:

Huh?

Amit Singhal:

Not for spying?

Martin Mendez:

Not for spying, but at the end, we transform those cameras into a living agent that is serving what's going on. So now we are able to track all the tractors that are getting into the plant and foresee what they are doing and optimize their roads. We're also monitoring inventories and we are helping them with health and safety in terms of seeing if the people are in the right places, they're using the helmets and the like. And that was a physical platform that was there without taking the volume and we explored that. And last but not least, in terms of Agentic, we also implemented an Agentic platform in one of our customers that is in the bakery industry, to help them to better serve their suppliers and give them information of where their payments were, if there was something that was blocking, when to expect that those payments to happen.

And that allows them to reduce like 30% of their physical agents, and in this kind of case, of the BPO that they were mentioning.

Amit Singhal:

So I think it sounds like based on examples both of you said, it's a combination of industry depth, good old EPAM engineering, but applied in a forward deployed capacity to work with clients closely. That's great. Thank you. Thank you.

Enver Amdiy:

Thank you.

Martin Mendez:

Thank you, Amit.

Amit Singhal:

Martin and Enver, for sharing cool examples. And thank you, Sridhar and Stefan, for what you do, how you do. Between two of you, we have good part of EPAM, so no pressure again. Sounds like resilience by design. We need more of it. So thank you all, and that's a wrap.

Next, you're going to hear from our CFO, Jason Peterson. But before Jason comes on the stage, again, let's hear from one of the EPAM clients. They're called LDC, one of the largest global soft commodity trading company and logistics company. So enjoy the video and then you'll hear from Jason. Thank you.

Guy-Laurent Arpino:

It's great to be here with you today. I'm Guy-Laurent Arpino, and I serve as chief information officer at Louis Dreyfus Company or LDC, which is one of the world's leading global merchants and processors of agricultural goods since 1851. I oversee our global digital strategy and technology initiatives, spanning trading, supply chains, and corporate functions. I've been part of LDC for 10 years, following 15 years at Procter & Gamble and five years at Bacardi. We began working with EPAM in 2019 on our analytics and AI transformation, quickly scaling teams across Hungary and Belarus. The partnership expanded naturally into software engineering with projects such as MYLDC, our customer portal, and our largest front office program.

Throughout COVID and various geopolitical disruptions, EPAM has ensured delivery continuity and helped us scale effectively. They also executed the award-winning migration of our five global data centers to Microsoft Azure in just 24 months. More recently, we've partnered on strategic AI initiatives, including next generation pricing engines. What truly differentiates the relationship is our shared focus, not only on what we build, but how we built it, ensuring scalable, sustainable, and modern engineering practices in a fast evolving landscape.

Today, we jointly embark on fully leveraging the potential of agentic AI software development lifecycle, transforming the way IT solutions are being designed, implemented, and delivered at scale. This will allow us to achieve our ambitious roadmap to become an AI-powered company. Over the years, EPAM has played a significant role in modernizing the application landscape, enhancing our enterprise architecture, and reducing technical debt by approximately 40%. In addition, we've made substantial progress in our data estate, developing a data platform on Microsoft Azure from the ground up. This platform serves as the cornerstone for both our data and science initiatives and AI-based products and services.

We anticipate further opportunities for collaboration and innovation and AI-enabled software development. I can only advise you to keep the high degree of ownership and accountability in the delivery of our projects, and I look forward to benefiting from the broader perspective across industries and the technological landscape, particularly around AI.

Jason Peterson:

In this final presentation of the day before FP's closing remarks, it's probably going to be no surprise that I'm going to talk about the business from more of a financial perspective. I'm also going to lay out our expectations for the coming three years, 2026, 2027, and 2028. I'm going to focus on our accelerating revenue growth. I'm also going to talk about our improving profitability, and then I'm going to talk about our ability to continue to generate strong free cash flows.

First, I want to explain what's in this slide. And so off to the left, clearly it's 2022 through 2025 are actuals. For 2026, what I'm showing you is just the midpoint of the guided range from our most recent February earnings call. I think the point that I want to make, and I think Larry did a really good job of

kind of reminding us what we've been through over the last four to five years, is that we had to deal with increasing difficulty in our operations in Belarus. We had the invasion of Ukraine. We exited Russia, that was both a delivery location, and it was also a revenue generating market for the company. We were able to maintain steady revenues throughout this time period, returning to growth at the end of 2024, further improving our growth rate in 2025, where we recently discussed our organic constant currency growth rate of approximately 5%. And more recently discussed our expectations for 2026 with a 3% to 6% organic constant currency growth rate.

If I look over to the profitability side, from a non GAAP operating income standpoint, we're nothing if not adaptable. Again, what we've been through with having to move our populations, support our employees, make certain that we've maintained our customer commitments, continue to invest significantly in capabilities, particularly all the AI capabilities we've been talking about. We're able to maintain sort of steady non GAAP operating income throughout the time period. Again, returning to growth in 2024, further accelerating that growth in 2025. And we're talking about solid growth as we move from 2025 to 2026. Off to the far right with the non GAAP diluted EPS, again, you've had growth throughout the last three years. For 2026, including the share repurchases, we actually returned to double-digit growth and non GAAP EPS between 2025 and 2026.

So I think this is really interesting, and you've seen this kind of throughout the day. And why it's interesting to me is that not only is our expanded geographic footprint a source of revenue growth for the company, but it's also an opportunity for us to continue to expand profitability. I think we've talked about the fact that we really have delivery excellence regardless of geography. So we've got AI capabilities globally in all of the regions in which we operate. And so what we'd understand then is that, of course, now instead of just delivering from Belarus, Ukraine, and Russia, we now have the opportunity to deliver around the globe. We can meet client expectations for different time zones, different price points, and when clients have specific sort of preferences in terms of geography.

Okay, now on top of that let's talk about profitability. And so, I think what everyone would understand is when the invasion of Ukraine happened, the exit from Russia, we had to move quickly. We had to move into new countries. We had to grow rapidly. The net result was that we were taking care of our employees. We're meeting our client expectations for delivery. At the same time, we were growing and then obviously focused on cost efficiency, but it was a lower priority. Today we've been much more focused on cost efficiency in some of the newer geographies and the geographies that scaled quickly. I think I've been saying for the last couple of years that even if you're worried about bill rates in India, we can still maintain high levels of profitability there. And that India actually generates higher profitability than the company average. If you add to that, the fact that we've been focused on the cost efficiency, India continues to improve profitability and the gap between India profitability and EPAM average profitability continues to grow. We've done similar things in Western Central Asia where we've continued to grow our profitability.

LATAM has been interesting because one of the advantages of the NEORIS acquisition, is we've learned a lot more about how to operate efficiently in Columbia and more recently in Argentina. So in all these cases, this gives us a further opportunity to sort of improve our profitability. It's one of the reasons why we're guiding towards profitable revenue growth in 2026 with an expansion in gross margin. So this is effectively just a reiteration of guidance, so it's 1385 to 14 billion for Q1. For the full year, 4.5% to 7.5%, which digests down to three to 6% organic constant currency growth. What you'll note for the non GAAP income from operations measured as a percent of revenue is that the 13,500 to 14.5 for Q1, at the midpoint is higher than what we generated in Q1 of 2025.

The same thing's true for the full year 2026, the midpoint of the 15 to 16% range, again, higher than what we actually produced in 2025. So we're seeing not only revenue growth, but improving

profitability. Again, you go to the bottom portion of this page here and you add the addition of the share repurchases. And you've got double-digit growth in non GAAP diluted EPS, approximately 14% in Q1 and at the midpoint of the range, approximately 11% for the full year.

From a long-term financial algorithm, what you're really looking at is a focus on continuing to grow and to accelerate that growth through success in the market for AI native and AI foundational services. We're looking to continue to expand profitability. Again, that'll be with a focus on sort of cost efficiency. And then we've all talked about AI productivity and the opportunity to share in those benefits with clients, give them some cost efficiency, retain some for ourselves, which improves gross margin. We've always had strong operating cash flows, modest capital expenditures, so that produces strong free cash flow. And then from a capital allocation standpoint, we continue to invest in our business. We've done strategic M&A, and more recently we've introduced share repurchases, including the \$300 million ASR that was announced in March of this year.

So from a growth standpoint, I think what I first do is take you off to the right side of the page. So we are participating in an immense 1.8 trillion IT services market. Again, we're quoting the Gartner statistics that market is growing. Underneath or within that market, there's the much higher growth opportunities associated with AI native, AI foundational. And then we've talked off and on throughout the day about kind of the more greenfield opportunities for EPAM, Agentic BPO, AI enabled managed services. There'll probably be some contribution from M&A over time. So what we are looking to do is continue to accelerate our revenue growth by participating and more importantly, succeeding in the high growth markets with the goal of eventually returning to 10% or a double-digit organic constant currency revenue growth.

From a profitability standpoint, I've talked about the fact that we were from an actual standpoint in 2025, 15.2% adjusted IFO. As you move to 2026, you've got the guided range of 15 to 16%. And then what we're looking to do is continue to improve profitability over the next couple of years, returning to a 16% plus in 2028. And again, what we'd be focused on is both improving gross margins, and then in 2027 and 2028, also gaining some additional benefit from SG&A. I think I've talked over the last couple of quarters about our focus right now is to continue to invest in business development and sort of sales focused marketing. I don't expect us to see a lot of leverage in SG&A in 2026, instead which you'll see gross margin expansion.

Then over time, you'll see a little bit more efficiency benefit from SG&A. At the same time, we're focused on generating between 50 and 70 basis points, gross margin improvement. After 2026, that would come from the cost efficiencies, that would come from the pyramid or seniority index we've talked about. Nothing heroic, just kind of returning back in the direction of what we might have generated historically, utilization improvement, and then the AI associated benefits, again, sharing those with our clients.

From a free cashflow generation standpoint, we've always had strong free cash flows or generated strong free cash flows over 500 million in 2023, over 500 million in 2024. More recently, we actually generated over 600 million in 2026. As I look forward, we'd be committed to continue to maintain the 80 to 90% conversion rate that we have historically targeted. As I look at our financials over time, that means we would generate over 1.8 billion plus in free cash flows over the time period, 2026 through 2028. So I think most of us are aware of the fact that the company's got a very strong balance sheet. At the end of 2025, we had \$1.3 billion in cash. We have modest debt. We have an untapped credit facility. On top of that, we've got the ability to generate the 1.8 billion in free cash flows that I talked about.

As we look over the last couple of years, our historic use of cash, clearly we invest in our business. I think we've talked about this throughout today, in terms of the skill development, the education, the platform technologies, the AI capabilities, the IP and the assets. And so we're spending hundreds of millions of

dollars on that. That keeps us at the cutting edge and gives us the opportunity to continue to grow faster than the rest of the market. So we'll continue to make those investments. We'll continue to do strategic acquisitions. And then over the last couple of years, we also introduced share repurchases, we would continue to do those. So off to the right here, so we'll continue to reinvest in the business. You'll have capital returns in the form of share repurchases. And of the billion that was authorized most recently, we still have \$450 million left in that.

Finally, you'd continue to see some level of M&A, probably more in the tuck-in category in 2026. So if I then just sort of close here on the M&A objectives, we clearly would look to sort of expand our in market capabilities, particularly industry vertical capabilities. And clearly that augments our AI capabilities. We also might use M&A to help us effectively be an entry point for select geographies. There's the type of idea where you sort of create a beachhead, which then you can grow behind. Then finally, we would use M&A to sort of deepen the scale of certain high growth capabilities. I've always thought of our M&A strategy as one that is not necessarily designed to buy revenue, but it's really designed to sort of help shift the company. To create an opportunity for the company to address different opportunities and then further our organic constant currency growth rate.

Most of the companies we acquired do have our services businesses, so there's strong free cash flows. And then finally, our historic focus is to make certain that we're able to sort of maintain the 16% plus profitability. Over the last couple of years, we got away from that. In the future what you'd see is we'd make certain that we did acquisitions that allowed us to either achieve the 16% or actually sort of facilitated the achievement of the 16% profitability range.

So in conclusion, we're focused on ongoing acceleration and revenue growth. We intend to continue to improve profitability, returning to the 16% plus adjusted IFO level by 2028. We're going to continue to generate strong free cash flows, the 1.8 billion plus that I've been talking about. We'll continue to make disciplined capital allocations, including share repurchases. Again, with that, that's the end of my presentation. We are

... we're going to have Q&A after this, but right now we've got one more customer video. I think it's Bank of Ireland and thank you very much.

Myles O'Grady:

Hi, I'm Myles O'Grady, Chief Executive of Bank of Ireland Group. I want to share the story about our new app, which EPAM has strongly supported us on, launching in the coming months. As we all know, customer expectations have changed fast and continue to evolve at pace. Keeping up isn't enough. We need to ensure our technology is market-leading. So, we've implemented change and made improvements with increasing momentum. We've upgraded systems across the bank, delivering greater stability and resilience and better customer service. We've invested in technology that our customers use most, facilitating payments across Europe in seconds, and upgrading our digital banking offering. And AI has helped us protect customers more and do things faster and better.

The core focus has also been the reinvention of our mobile banking app, a hugely important part of customer service offering. We've been working closely with EPAM on this, and I know they understand our vision and ambition for what we want to deliver here. The result is a banking app in pilot right now and launching soon that we can be proud of that will deliver great outcomes every day for our customers and will help us go further, faster over the time ahead. 2025 was a transformational year for Bank of Ireland in our tech delivery, and this year promised us to be even more exciting. I'd like to thank FB, Ark, and all the team at EPAM for their strong support in the progress we are making and in the delivery of our ambitious plans for the future. We're looking forward to the journey ahead and to what we can achieve together.

Moderator:

Okay. A lot of content today. We've got our final Q&A session with several of our leaders. Same rules go as the first session, so please raise your hand if you have a question. Allow the mic to come to you. Raise your hand and excuse me, state your name and firm. And then for those online, if you submit a question, we are looking, we'll field those as well. Let's go with the first question right here. The glasses.

David Grossman:

Thank you. Yeah. Thanks. It's David Grossman from Stifel. You did a great job of laying out the structural tailwinds from AI and why EPAM is well positioned to benefit from those tailwinds. And I think what's notable is historically at least these massive changes in technology have been accompanied by accelerating growth for the industry. On the other hand, industry growth has been relatively low, call it the last 18 to 24 months. So in your opinion, what is so different structurally about this cycle and what needs to happen for growth to not only reaccelerate for the industry, but obviously for EPAM as well?

Balazs Fejes:

David, good to see you and let me try to address that. I think what's really different at this time is the rate of change of these fundamental technologies are so much faster. Our clients, ourselves, and everybody who is participating in this is just watching the race, what we are seeing. And somebody in the audience, we discussed it already this morning, that some people are just waiting til things gets cheaper. They are waiting for ... If you wait one more month, maybe the model will get better. Maybe if you wait one more month, maybe the model not just gets better, it gets cheaper.

And just probably six months ago when we reached the point where the model's results are good enough and they're cheap enough that you actually can launch these transformational programs. So I think we got so accustomed to that people rush ahead and allocate capital and start making these investments that we underestimate the resistance and the time and cautiousness people having with these new AI models because it's no longer just an IT change. Cloud was the internal affairs of the IT department. This is a business change. This requires business leader committing into a massive change program, how to change the whole business processes, and also addressing the technical debt which they carry today.

Because without addressing the foundational element of AI, which we keep talking about it, the cloud migration, the data and data product creation, the legacy modernization, you're not going to get the benefits. And also you need to upscale your teams. Everybody's reluctant to get locked in to a vendor, logged into an engagement model, and everybody's hoping that they can do this without massive changes to their organization. So this is created kind of, I would say, wait, wait and see period. But now what we are feeling that organizations are no longer able to wait longer. They are just ready to launch into it. And we have these active discussions, which makes us very optimistic that going forward, we're going to see the demand bouncing back.

Now, when do we see for the whole industry to start doing that? When one player, one client of ours or maybe a client of our competitors actually succeed with the transformation. Maybe they don't even do the full-blown transformation of a whole company, but if you transform just one line of a business and achieve some level of efficiency gains or speed to the market, which we never seen before, that will force everybody else in that vertical, in that geography to do the same and do transformation in masses.

So I think this is where we are and we are in the stepping point. This is the first time when you start hearing from the frontier players, the Anthropic, the OpenAIs that they actually done the engineering and they actually now seeing that [inaudible 03:18:41] is actually seeing that the first time he really saw real efficiency gains from software engineering using AI, something which really surprising. The first time

he actually trusted the AI to do the coding was late last year. I think we just underestimated how much time it will take to get to this stage, and we are there. Now it's going to start happening.

Moderator:

Maybe in the back. Great.

Maggie Nolan:

Hi. [inaudible 03:19:13]. Yeah. Thank you so much. Hi, Maggie Nolan with William Blair. Why is vertical expertise more important now? And EPAM has a broad set of verticals that they address. So do you need to narrow that focus or are there ones that you're going to start with first to maximize success?

Balazs Fejes:

Maggie, good to see you. And I think it's a good question. Why now? In the digital transformation space, horizontal skillsets were much more important. People were actually applying horizontal capabilities into variety of industries. During the AI transformation, on the other hand, they have to solve business problems. They have to now tackle the business challenges. That requires real industrial knowledge. That requires you to discover how to automate that piece of functionality. That requires you to really understand deeply what to do. You can walk around and you will see how we are tackling it in energy, how are we doing in healthcare, life science.

But in order to do that, you really need to understand what you're doing because as Vic mentioned it, the models will hallucinate. If you're not grounding it into vertical expertise, what you're going to produce is not going to be safe. If you look at LivaNova case study, which is an AI engineering case study, the fact that they got their software FDA approved, and by the way, they used a ton of agentic solutions, it's underlines what you need to bring to the table in order to make it safe, order to make it really productive in this environment. Vic, Elaina, do you want to add something to it?

Elaina Shekhter:

Just I think that the cycle change that we see is broadly a move from digitizing businesses, which is what we've been doing for the past 10 years with cloud and modernization to transforming them. And I know we've been using digital transformation kind of as an industry term. I think we're really on the forefront of actual business transformation. So it's not just digitizing or creating data platforms. It's actually re-imagining new business models with an AI centric point of view. A, that's hard, technically hard, but it's even more difficult from an industry point of view, because particularly in regulated industries, you don't just get to try it and see if it sticks.

Speaker 1:

Maybe one more thing. Am I with mic? Yes. We have actually ... We are not starting from scratch. Like from 2014 or 2015, we are building healthcare and life science at scale of the whole organization end to end. And you'll see Gregg today, he will see it, show it, and it's engineering, it's consulting, it's advisory strategy and all the levels. So it's a good question about focusing, but we are doing it and so it's improvement. It's not necessarily you need to be locked in somewhere. The same happening with energy from 2016, from large workloads there. And so time after time ... Gaming. So you see it with us actually over time.

Moderator:

Let's go here.

Surrinder Thin:

Surrinder Thin with Jefferies. As you think about the build part of the equation and you start to build bigger, more complex platforms and all of the orchestration that's going to be on top of that in terms of the agents, who at the end of the day owns the IP and do you have the ability to maybe manage or run those platforms and monetize the agents or the capabilities? Or does EPAM still remain within the build parts and then you just kind of let your clients run the platforms at that point? And then you, I don't want to say you walk away, but then you work on the next project.

Balazs Fejes:

Thank you very much. I think it's hard to see where this leads to. I think we're already in the position that the build versus buy equation is flipping. We are starting to build more and the clients are choosing to build more instead of buying off the shelf packages or SAS solutions or we are actually rebuilding some of these SAS solutions for the clients who take it in house. In the current wave, what we're seeing is that people want to own their own IP, especially when you are talking about agents. Remember, this is a workforce transformation. You want to own your own workforce. And even if your workforce is no longer humans, they could be agents, but you as an organization, you have to risk mitigate, you are dependent on your workforce.

So we are seeing clients really want to own the agents themselves. They want to have flexibility. This is at the next level of risk management. You cannot be logged into a vendor like EPAM or even to a model or even to a hyperscalers because now your whole business starts dependent on that workforce. I think we have to start seeing this transformation itself in a very different eye. We have to see it as a workforce or actually whole business transformation and less of an IT transformation.

And once you start seeing it from that angle, and I had one of these discussions just the other week with an insurer, they are seeing this from a risk management point of view, who owns the agent itself because they are going to be so much ... Because in the end, in the end of the transformation, they're letting go their original workforce and now they are relying on a new agentic workforce. And in that equation, your business dependency is switching and they want to own that dependency.

Elaina Shekhter:

Can I just-

Balazs Fejes:

Yes.

Elaina Shekhter:

So there's not just the IP over the build versus buy, the critical sort of rights issue is to the data. And it actually is probably one of the reasons why the frontier companies will likely not end up owning the full end to end because there is a critical mission type of not just rules and logic, but actually the ownership over their own data. And yes, we can maybe operate some of the platforms that we design and build, but it would be a rare thing where we would be the owners and controllers of the underlying enterprise data set.

Moderator:

Let's go here.

Fanikan Muri:

Hi, Fanikan Muri from HSBC. As you said, the technology is evolving very rapidly. So how do you ensure that your employees are upscale in such a fast changing technology and how do you price in this technology as AI native services likely requires a long contract period? So how do you price in these kind of changes?

Balazs Fejes:

That's a very good interesting question. I think we already saw in this AI transformation space, probably two or three technology shifts by this point, right? I do remember when Andreessen Horowitz very proudly launched the next generation AI architecture, which is probably six or eight months later, was out of date. I don't know if you remember the reg architectures and all the different vector databases, which everybody was crazy about probably 18 months ago, 12 months ago. And today nobody talks about it anymore because the context window grows so much.

I think staying on the frontier requires you to continuously do R&D, continuously have people who are on the frontier and actually on the edge of the model. So we have team members work evaluating, building new types of capabilities in house. And from that, designing new types of educational programs, the educational programs which Sandra and Aleksey was talking about, how we pushing it out.

And I think Vic was also indicating that we codified the environment which we are operating. So in reality, we are running the internal systems with agents, with MCP connections to the internal applications. So you can actually run operations as a code in the organization. So this is a very different shift, but how you educating them continuously, how you are actually making them work continuously be running programs. Maybe Larry or Vic, you want to talk about this.

Larry Solomon:

Yeah. Just the one thing I was going to add is to some degree, a lot of this starts from the beginning, from the selection, and we have very, very rigorous technical requirements and expertise that comes from EPAM that starts even before the individual joins the company. So we try to only include in that pipeline of candidates that we select, those that we believe have the strongest technical chops.

Speaker 1:

And just to continue on this, this is very important comment because one of the key differentiators of the AI agent engineers is judgment. So we actually continuously in this work, Ark started to say, we were selecting people with judgment all the time, maybe subconsciously, maybe consciously, but now we understand exactly how to select them, how to separate those who have it, those who do not have it.

Dmitry had a great message on talent density. Those are important building blocks. Now about the speed and desire and everything, it's all about also top down by example. If [inaudible 03:29:19] is coding, I'm coding, we can demonstrate it, we can show it, and the next managers will be coding with Claude Code, with something else, and this will stimulate newer technologies. We also have mandatory requirements for education. So if you are not educating, it's a bit of a problem.

Balazs Fejes:

Now going back to the pricing, which was I think was your underlining question, seeing where you're sitting and what your role is. I think we are clearly, for this type of skills, we are able to charge higher rates, which actually we talked about it before that our AI native portfolio drives higher margin, but I think it's a continuous moving target because some of these skillsets becoming commodity and as technology rapidly changes.

Moderator:

Let's go here and then we'll go there.

Brian Keen:

Hi, it's Brian Keen at Citi. Maybe for Larry and Victor, how do you imagine the delivery model in terms of people? Do you need more or less as this evolves over the next three years? And then just one for Jason, accelerating revenue growth, given the industry dynamics right now, the question obviously is on visibility. What kind of comfort can you give us in visibility, either contract bookings or when you look out two years, what percentage do you see and how do you get to that number of accelerating revenue growth? Thanks.

Larry Solomon:

Yeah, I can start by saying on the more or less, I think the answer is yes, it's both. I think not only from a geographic perspective, but from a skillset perspective, more of some of these, less of some of those. And if you go back to recalling some of the things that Sandra was talking about, the level of data and the signals and the metrics that we track in order to figure out as best we can, where we need those people, what skills, helps us get a little bit ahead of that curve.

I think the other thing that I would say is, we work really hard to put a plan in place every year, and my view is, it's only valid for one day, January 1st, because on January 2nd, something has already changed, especially in the market that we're in today. And the companies that are going to win are the ones that can figure that out and pivot the fastest.

Jason Peterson:

Yeah. And so for the accelerating revenue, so first I'd just start with this year, 2026. So with the guide at 3% to 6% organic constant currency, our focus is on making certain we can at least hit the midpoint of the range and clearly we're all driving to achieve something to the higher end of that range. There's already kind of line of sight to larger opportunities that if we can close and we're trying to close, I think sort of drives us above that midpoint of the range.

As I look further ahead, it's the ongoing success with clients. It's all the things we've talked about that clients can't do this themselves, their increasing dependency on partners like EPAM. But hopefully what we've convinced you of today that this isn't easy and EPAM is extremely well positioned to participate in these high growth market opportunities. And if you can be successful in a market that's growing rapidly, that drives higher revenue growth. And that's kind of how I would think about it over the next couple of years.

Moderator:

We have time for one more question. Sorry, him.

Punit:

Thanks for the presentation. Punit from JPMorgan. It was interesting to see all those regional heads coming here on the same table in the panel. So talk to us, how does EPAM operate across different regions? Because the individual regions might have different cultures, like the policies. So is it the same culture, same EPAM across everywhere, same type of people in terms of profile, type of people you hire across all regions or are there differences based on that region's policies or culture?

Balazs Fejes:

So let me try to start with it. Vic is running a global delivery platform. Basically, he runs the factory itself, right? It's an engine. In this engine, we are enforcing certain level of uniformity, right? What Vic highlighted is the assessment, and basically that's our requirements, how you're going to get promoted. That's the way you are actually being assessed. That's the way you are actually reaching the next level. Is it the same culture? No, because we are coming from a different parts of the world, but there are unifying elements. There are volumes which we are sharing. There are ways how we're communicating.

And we have to collaborate. We have to work together. We are working together to deliver to one client. And throughout this delivery, we're actually kind of syncing up. We have the same volumes what we're pushing out. We are being assessing people in the same way. We're hiring for the same goals and for same profiles with the same criteria. We're running the same process globally, how we run compensation, how we run assessments, how we're going to provide feedbacks and performance management. It creates one level of sync. And overall, we're hiring engineers, and engineers kind of understand each other and kind of sync on it in a weird way, right? In a geeky way. But I think that's who we are.

Larry Solomon:

If I could just add.

Balazs Fejes:

Absolutely.

Larry Solomon:

Sorry.

Balazs Fejes:

Go ahead.

Larry Solomon:

Were you done?

Balazs Fejes:

No, exactly. Go ahead.

Larry Solomon:

No, go ahead. Go ahead.

Balazs Fejes:

Go ahead, Larry.

Larry Solomon:

I think a short way to look at it is globally consistent, locally relevant, and at the end of the day, it's what's best for the client. Client-centric decisions that are locally relevant with the global consistency in processes, culture, core values, but locally relevant is extremely important.

Balazs Fejes:

Excellent. Thank you.

Moderator:

That wraps our Q&A session. I'm going to hand it back over to FB for closing.

Balazs Fejes:

And if we figure out where the clicker left the building. Who has the clicker? All right. I hope by this point with the team, we made clear our positioning and why we have the right to win in the AI native era. And I really would like to thank the team itself to make such a great presentation and actually present this message. Our people have navigated technology, social, and geopolitical challenges and changes. We have emerged stronger out of it. We learned a lot, and I think we are the most resilient organization out there, not just in terms of against geopolitics, but any type of technology and social change. Why invest in EPAM?

We will be the winners in AI era. We are best positioned to be a leader for enterprise AI transformation. We have the strongest engineering talent or engineering DNA in the industry with a track record of solving our client's most complex hairiest problems. We are delivering already AI foundational and AI native work, and it's expanding and it's growing significantly. We have a clear strategy, focused on accelerating and driving profitable growth with margin expansion. Our 2028 goals are accelerated revenue growth, 16% plus non-gap operating income margin, and delivering \$1.8 billion cumulative free cash flow throughout 2028.

Thank you very much. Okay, doesn't work. As I said to you, something has to break. Thank you very much for coming. And for the audience online, we would like to thank you for attending and see you next time. Thank you.