Operator: Everyone, and a warm welcome to the Infosys Investor AI Day at our Bengaluru campus. A special hello to everyone who is joining us via the Investor Relations webcast. Today's proceedings are being recorded, and the audio transcript and the presentations will be made available soon on our website. We request you not to take pictures or record the sessions while they are going on.

Before we begin, I have some important housekeeping announcements. On your tables, you will find the agenda, an important information sheet, and the feedback forms. Please follow the timings in the agenda to help us keep the day running smoothly. The important information sheet also has the Wi-Fi details. We request you to fill in the feedback forms after every session. If you need any assistance, we have volunteers around wearing Infosys ID cards. Please reach out.

Please take note of the nearest emergency exits on both sides of this room; they are clearly marked. In case of an emergency, please follow the signage and volunteer instructions. Restrooms are also outside this hall on both sides. Kindly access charging stations on both sides of the room for charging your electronic devices. Please note we will not be taking questions at the end of each session. There will be a question and answer session at the very end of this event.

For departures in the evening, airport coaches have been arranged. The first coach for airport transfers leaves at 4:45 p.m. A second coach will leave at 5:00 p.m. Those who have flights at 8:00 p.m. are requested to take the 4:45 p.m. coach considering famous Bangalore traffic. Lastly, I request you all to put your devices on silent mode. Thank you.

With that, let's start today's program. For our first session on Tech Transitions: Why is the AI transition different? Please welcome Nandan Nilekani, Chairman of the Board, Infosys.

Nandan Nilekani – Chairman, Infosys: Thank you, and it is great to have you all here in these tumultuous times. Today I will talk about tech transitions. I have had the fortune of being in this industry for more than 40 years, and I have seen many transitions. I thought I would talk less about that and more about why this time it is different and the implications of this transition.

We have seen technology shifts for centuries, whether the printing press or telegraph. But over the last 60-70 years, we have seen much faster change: PCs, cloud, generative AI, agentic AI, and so on. The speed of change has been a constant for many decades now. Each time there is a change, the way we address it is different. We went from mainframes to mini-computers, PCs, client-server, LAN, web computing, mobile, enterprise applications, and big data. Each time, firms like Infosys had to deal with what was new.

This time, the AI transition has been much faster than earlier transitions. Internet took more than 10 years to reach 1 billion users, and smartphones took 5 years. AI is taking a couple of years. The speed of AI is largely because the infrastructure was already ubiquitous. Internet and smartphones allowed people to distribute ChatGPT, Gemini, or Claude very easily.

What has happened this time is a much more fundamental change to the way businesses operate. This is not just a layer of technology. When smartphones came, we put a front-end on existing applications. When cloud came, we could "lift and shift" to the cloud. But this time, it is different. It is an AI-native architecture change. We cannot run business the old way; customer journeys and

operating models must change.

It is also a huge challenge for talent. Writing code will not be the goal; it will be making AI work and orchestration. Our mental models must change because technology was always deterministic before. In the AI world, you give a prompt and might get a different answer every time. How do you build robustness and resilience in a non-deterministic world? This is a fundamental root and branch surgery of the way business is done.

One clear learning is that modernization of legacy systems cannot be deferred anymore. Over the last several decades, people just added to legacy systems rather than replacing them. If you look under the hood of a large enterprise, they have mainframes from 1960 coexisting with systems from 2000 in silos. That is over. If a firm wants to take advantage of AI, they have to clean this up. Large companies spend 60-80% of their IT budget on maintenance with no business value. They want to flip that to spend 60-70% on new systems, but they cannot do that without a cleanup.

Accumulated technical debt over decades must be paid. The good news is that for the first time, because of AI, we have the tools to do modernization fast and economically. The balance of advantage is also moving toward "build" rather than "buy." Building applications has become so simple that firms may replace bought software with custom-built solutions. This benefits firms like us who build these for them.

Foundational systems will increasingly become systems of record, but the interface will be agentic. Agentic interfaces allow for a pro-consumer design that hides complexity. Enterprises will want to put agentic layers on top of all applications. This requires orchestration.

The technology is currently far ahead of its deployment. Frontier models moved from 100 billion parameters to 1 trillion in a year. However, implementing this is hard. It involves organizational change, retraining people, and breaking data silos. We call this the "deployment gap." This concept is similar to what Professor Clayton Christensen called "technology overshoot" 25 years ago. The technical capability will keep getting better as billions are poured into it, but enterprise deployment is slower. Addressing this deployment gap is where we can help.

Talent transformation is huge. We will need AI engineers, forward deployment engineers, forensic analysts, and data scientists. Productivity gains in "greenfield" projects are easy, but the real world involves "brownfield" systems with trillions of dollars invested, technical debt, and undocumented dependencies. Some old systems are only understood by 70 or 75-year-old specialists. Modernizing these is much more complicated than greenfield development.

Implementing AI requires focus to avoid generating "slop" or fake productivity. We need usage guidelines and quality gates. First-principles thinking is more important than ever. We must train people to solve problems without tools so they understand the "black box" once they use the AI. Understanding enterprise context is also essential; every company has different legacies and silos. This context cannot be captured by a tool alone.

There is no opportunity gap; if anything, the opportunity is bigger than ever. The question is how a firm transformation its talent, products, and services for this new world. This is an execution risk, not an opportunity risk. I hope you will hear today that we are on the right track. Thank you very

much.

Operator: Thank you, Nandan. For our next session on the AI services opportunity, please welcome Salil Parekh, Chief Executive Officer and Managing Director, Infosys.

Salil Parekh – CEO & MD, Infosys: Good morning and welcome. I want to share where we see the opportunity today in AI services and how we are going after it. Our clients trust Infosys in driving and delivering AI work. We are currently doing AI work for 90% of our top 200 clients. This is not just pilots; these are parts of large programs.

We have introduced an "AI-first" value framework, represented by a hexagon with six large areas of growth.

First is AI Strategy and Engineering. This involves strategic work with CEOs and boards to build and orchestrate agents. Second is Data for AI. Every enterprise is protecting its data, and there is massive work to be done to make enterprise data ready for AI. Third is Process, where business processes are being driven by agents. Customer service is a major driver here.

Fourth is Legacy Modernization. As Nandan mentioned, this is a massive opportunity to move away from legacy landscapes. Fifth is Physical AI, where AI software is embedded into devices. Sixth is AI Trust, covering cyber security and responsible AI.

This hexagon is not theoretical; it is happening on the ground. This work represents 5.5% of our revenue in Q3 and is growing at a robust pace. We break this down further into 30 offerings and 100 sub-offerings, all enabled by our engineers and agents built on Infosys Topaz. Our entire go-to-market team is focused on these to drive revenue growth, cost optimization, and innovation for clients.

We have always talked about "Navigate Your Next." This "next" is about AI. Our strengths lie in our understanding of the client landscape, our domain knowledge, our engineering talent, and the platform and IP we have built.

We see an external market opportunity between 300 billion and 400 billion by 2030. While AI productivity leads to some compression in traditional IT services, we believe the expansion opportunity is larger and will be the driving force of our growth.

Our playbook involves three pillars: AI-first services, AI-augmented services (where AI is infused into traditional work), and foundational components like Topaz Fabric. We are also going through a huge reskilling process. We have recruited 20,000 college graduates this year and plan to recruit another 20,000 in the next financial year.

Increasingly, we are connecting with the CEO client base, which is necessary for success in the AI world. I will pass it to our delivery leadership now and return at the end of the day for Q&A.; Thank you.

Operator: Thank you, Salil. For our next session on the AI Services playbook, I would like to welcome Satish HC, Chief Delivery Officer.

Satish HC – Chief Delivery Officer: Every tech shift leads to a rewiring of enterprise workflows, and AI is the next rewrite. A typical enterprise landscape is far more complex than a simple diagram because of fragmentation, technical debt, and regulation. Integrating AI is not just a software plugin. It is about the reimagination of processes and rewiring legacy power structures.

The enterprise stack consists of systems of record (deterministic/compliance), systems of intelligence (human engagement with data), and a layer of non-deterministic work (unstructured/unique problem solving). This top layer is underserved and ripe for AI transformation.

The "co-intelligence" of humans and machines will reinforce every layer. Systems of record will accelerate, and systems of intelligence will integrate structured and unstructured data. Approximately 60% of effort in an AI project goes into this integration before models are even applied. This leads to a new layer we call the "Systems of Cognitive Work." Humans will shift from acting on data directly to an oversight role.

Our playbook drives value at the intersection of intelligence, engineering, and domain. Sustainable advantages for an enterprise are created by deep integration into specialized workflows. Our engineering approach codifies enterprise context, which helps scale AI while protecting the client's competitive differentiation.

We are seeing new deal archetypes: legacy modernization with reduced risk, large integrated operations deals, and organizational transformations. We take accountability from strategy to execution.

For example, we worked with a CPG client that wanted to grow revenue to 7 billion. We used Infosys IP to build their unified data foundation and an enterprise agent AI platform. They now have 10 agentic AI products. In their R&D; function, they have a line of sight to 50 million in new revenue through an agent developed for product formulation. They have also unlocked 25 million in cost savings through operational optimization. Thank you.

Operator: Thank you, Satish. For the next part of the session, I would like to invite Dinesh Rao, Chief Delivery Officer.

Dinesh Rao – Chief Delivery Officer: Good morning. Our priority has been moving customers from experimentation to scalable industry-level AI. We have codified our services across six strategic pillars.

Pillar 1: AI Strategy and Engineering. We help create an AI blueprint and transformation office. For example, we worked with the board and CEO of Danske Bank to enable them to be "AI-first," identifying core processes like KYC and fraud detection.

Pillar 2: Data for AI. Data is everything AI needs. Most enterprise data is unstructured (videos, speech, etc.). Our frameworks help customers transform data into a uniform data fabric using semantic ontologies. For one industrial manufacturer, we harnessed 10 petabytes of data to drive supply chain optimization by 20-30%.

Pillar 3: Reimagining Business Processes. Processes must be reimagined based on how human intuition works with agents. At Toyota Motors Europe, we used a supply chain transformation agent

to automate drop-ship processes, increasing inventory visibility.

Pillar 4: Modernization. Tech debt is the obstacle for many organizations. AI models powered by Topaz Fabric allow us to transition large legacy codebases into modern cloud microservices with a viable ROI.

Pillar 5: Physical AI. Intelligence is moving from the cloud to the actual physical edge. This accelerates decision-making for vehicles and industrial operations. Our acquisitions of InSemi and In-Tech fit directly into this pillar.

Pillar 6: Trust and Governance. Hallucinations and model breaches are major concerns. Trust must permeate all layers.

I want to highlight NOVA Chemicals. They are a large petrochemical manufacturer. We brought in data from machinery and OEM manuals to help a planner use a simple natural language interface to guide maintenance. This orchestration between OT and IT systems created efficient work order processes in their SAP environment.

Finally, consider Hertz. They are modernizing their reservation and pricing systems, which consist of 3 million lines of code on tandem computers. We used our iLEAD platform to document these programs, showing database and table dependencies that were previously unknown. What would have taken 4 years will now take 18 months. Thank you.

Operator: Thank you, Dinesh. For the final part of this session, I would like to invite Balakrishna DR, Head of Global Services.

Balakrishna DR - Head of Global Services: Satish and Dinesh talked about AI-first services. I will talk about AI-augmented services—taking traditional services like testing, development, and infrastructure and inducing AI into them.

We have created detailed playbooks for 20-plus traditional services. We are working with leading models from Anthropic, OpenAI, Amazon Nova, and open-source models like DeepSeek and Llama. Infosys has even created its own coding model. In the Hertz example, we used an OpenAI model to "critique" code generated by another model, improving accuracy.

LLMs need enterprise context. We bring this through mcp registries and knowledge graphs of the enterprise standards. We have also created 100 specific agents for the application development life cycle. Engineering is changing, but great engineers are still a requirement. We have trained 90% of our developers on AI.

We have a 360-degree partnership with Microsoft. They are our partners, our customers, and we are their customers. They are currently moving to a "Master Customer Agreement" (MCA) to eliminate paperwork and evergreen licenses. Infosys built the IT system to manage this transformation, achieving 2x developer velocity and 35% improvement in time-to-market. For their intelligent cloud support, AI agents now monitor logs and predict issues before they happen, giving a 10x improvement in root cause analysis turnaround.

At Danske Bank, we have 20 use cases live. We created an internal chatbot called DanGPT with over 16,000 users. We used agentic AI to generate more than 2 million lines of code, which were then validated by our engineers. Thank you.

Operator: Thank you, Bali. For our next session on Infosys Topaz Fabric, please welcome Ravi Tarafdar, our Chief Technology Officer.

(Note: The transcript continues with Salil Parekh and Jayesh Sanghrajka for the summary and Q&A;)

Management: Good afternoon. Hopefully, the day has been interesting. I will spend a minute on a summary, and then Jayesh and I can answer questions.

First, we have a comprehensive set of AI offerings across our hexagon framework. Second, the opportunity is huge; we see an external market expansion of 300–400 billion. Third, enterprise clients trust Infosys, as seen in the testimonials from CEOs and CIOs. We have deep engineering talent and a platform with Topaz Fabric and AI Next. We are ready with the training and the balance sheet to make acquisitions as needed.

Jayesh, please join me for the Q&A.;

Analyst: Thank you. You highlighted that the new opportunity from AI will be 300 to 400 billion. Could you elaborate on what the net opportunity is after adjusting for the compression in traditional services? Examples today showed everything being done faster with fewer resources.

Salil Parekh – CEO & MD, Infosys: We have not quantified the compression number for external use yet. However, we believe the expansion numbered provided by external sources looks larger than the compression.

Analyst: You spoke about the advantage of understanding client data and context. Many peers would claim the same. What exactly is the delta or incremental capability that gives Infosys a better right to win?

Salil Parekh – CEO & MD, Infosys: The focus is on the platform and Topaz Fabric capability where we built our own agents and can integrate others. It is also how we identified the six specific growth areas and executed by reskilling our people. Continuous execution is the real difference.

Analyst: Has our pipeline to TCV conversion timeline improved because we can now build prototypes and working models much faster?

Jayesh Sanghrajka – CFO, Infosys: Large deals signed in the last few quarters have not seen a significant shrinking in timelines yet. The timelines remain relatively similar.

Analyst: One of your competitors recently said ERP migration programs are seeing compression from years to weeks. Do you see that practically happening?

Salil Parekh – CEO & MD, Infosys: On ERP migration specifically, we have not seen that. However, where you look at overall modernization from an old legacy landscape to a current landscape, the timeline is much more compressed and cost is reduced.

Analyst: Comparatively, in the previous digital cycle, we saw initial compression followed by opportunity. Where are we in the cycle now? Does the expansion from generative AI outweigh the cannibalization yet?

Salil Parekh – CEO & MD, Infosys: It is intertwined with the macro environment. We see the macro improving in the US with tax and interest rate expectations. Some areas like customer service and data are moving very quickly and are incremental. We do not see an acceleration of compression at this stage. Financial services and energy are showing visibility for strong growth next year.

Jonathan Lee – Guggenheim: Can you help us understand the level of investment necessary to reskill and hire laterally, and the impact on margins?

Salil Parekh – CEO & MD, Infosys: We will maintain our margin guidance. We have a strong program to make our costs more efficient, and we will take those savings and invest them into scaling AI faster. We are also ready to use the balance sheet for appropriate acquisitions.

Jayesh Sanghrajka – CFO, Infosys: We expanded margins by 50 basis points in FY25 and have kept them stable this year while absorbing all the tech, training, and sales investments.

Pankaj Murarka – Renaissance: You mentioned 5.5% of revenue comes from AI. In the context of Fortune 2000 clients, that still seems small, with average deal sizes maybe around 4–5 million. When do we see 500 million deals and accelerated adoption?

Salil Parekh – CEO & MD, Infosys: It is becoming part of almost every discussion. When we started calling out digital numbers, they were around 20–25%, and they grew to 70% over 5 years. This transformation might take 18 months or 7 years; we are ready for either.

Aditya – UBS: How should we think about pricing models? Some projects that needed huge teams now use lean teams and platforms. Is the model evolving toward outcome-based or platform-based pricing? And what happens to the utilization of the fresher workforce as we move toward more specialized projects?

Jayesh Sanghrajka – CFO, Infosys: Pricing is evolving toward a combination of outcome-based, agent-based, and traditional models. It depends on the client context.

Salil Parekh – CEO & MD, Infosys: Utilization will remain an important metric. We are still recruiting freshers and teaching them to operate with and without AI tools. We retrained our base during the digital shift from 25% to 65% of revenue, and we will do the same here.

Gaurav – Morgan Stanley: Is the line between software and services getting blurred? Nandan mentioned "build versus buy." What does that mean for the addressable market?

Salil Parekh – CEO & MD, Infosys: It expands the amount of work we can do. If building at the "edge" becomes easier and more effective for the client, there will be different custom builds for different clients, which is a larger opportunity for us.

Surendra – Citi: Is there a way for Infosys to capture the value better? Historically, value savings usually go back to the customer.

Salil Parekh - CEO & MD, Infosys: Today, enterprises expect model-type benefits in pricing before the enterprise reality has fully caught up. In the future, as they become more aligned, outcome-based sharing might become more equitable.

Jayesh Sanghrajka - CFO, Infosys: There are also indirect ways. Being the AI strategic partner for 15 of the top 25 clients gives us a larger share of the client's overall landscape.

Kunal - Bank of America: Is there a bigger opportunity in mid-market or emerging market customers if the model is no longer labor-intensive?

Salil Parekh - CEO & MD, Infosys: Small language models are very useful for limited data assets in large clients. We have not looked at the mid-market deeply because the cost of sales is different. We see growth markets like the Middle East as very strong for these models.

Sandeep - Aquarius: In a "post-AI adoption" era, when everything is modernized and automated, will the terminal growth rates of the industry be negligible?

Salil Parekh - CEO & MD, Infosys: The amount of software demand is becoming 100x the current size. Even with 10x productivity benefits, there is a massive amount to be built. Economic growth always drives a need for new features and technology. It is not a static world.

Kalaljeet - Kotak: Are the AI agents you deploy homegrown or from companies like OpenAI? And what is the service intensity if you are using a third-party model—how much service revenue do you get for every dollar spent on a model?

Salil Parekh - CEO & MD, Infosys: We use a mix of our own agents, foundation model agents, and third-party agents. We don't have an exact ratio of "model spend to service revenue" yet, as it varies by integration and security needs. We will build a larger data set on this over the coming quarters.

Analyst - Capital Group: For several quarters, net headcount has been flattish. With AI services picking up, are you gearing up for net headcount to pick up again?

Jayesh Sanghrajka - CFO, Infosys: We added 13,000 net headcount in the first three quarters.

Salil Parekh - CEO & MD, Infosys: As the macro improves and AI opportunities expand, we expect to continue with headcount increases in the coming quarters.

Operator: Thank you. That brings us to the end of the Q&A.; Thank you, Salil and Jayesh. Thank you, everyone.