Thank you, Cecilia.

Good morning, ladies and gentlemen. I am Shashi Menon, and I lead Digital Technology Development for SLB.

As Rakesh was looking back through our digital history, it reminded me of my own digital journey at SLB. I had the opportunity to lead the development of our first digital platform around Petrel, GPU computing with NVIDIA, cloud computing with Google, and Data platforms with Microsoft.

Now, here I am to tell you what we are doing in this exciting world of AI we're living in.

Today, I am going to show you our proprietary AI technology stack. I will tell you why it is entirely unique, how we have developed specialized domain foundation models, and the secret behind why nobody else can replicate this.

So, let's get going.

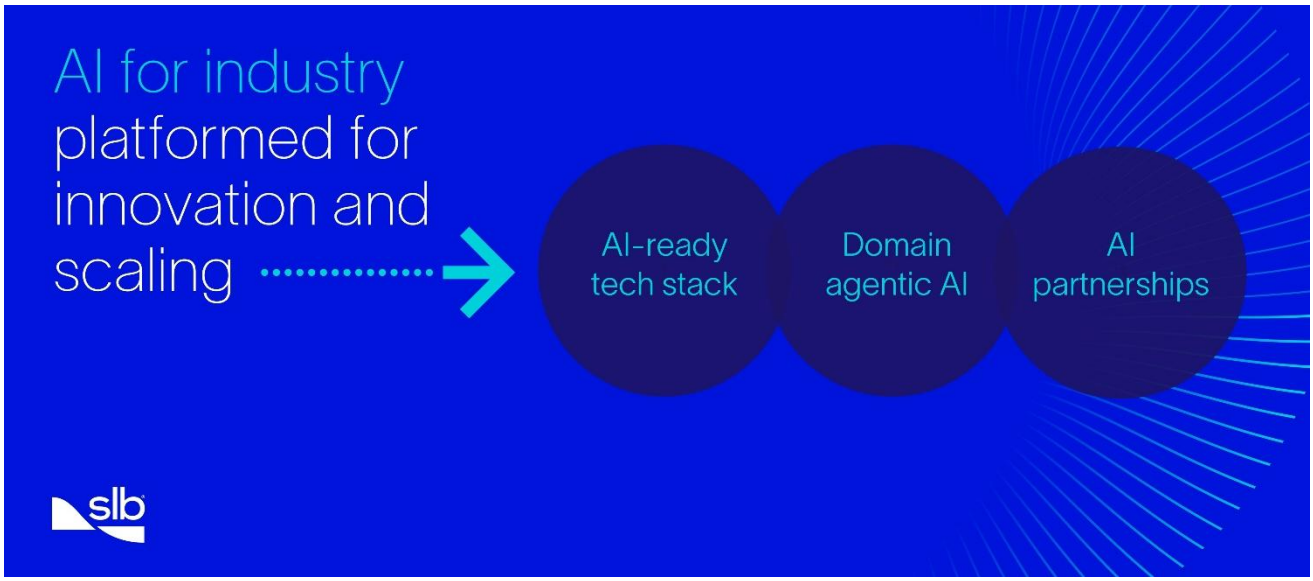

The foundation of our AI capability is an industrial platform configured specifically for the physics and data complexities of the energy sector.

These are three key elements that I want you to take away.

I can tell you that these elements are unique in our industry. No one has been successful in building these to date, not just in our industry, but across any industrial sector. And they are the ones that make-or-break AI in our industry.

So, let me explain each of these pillars and show you how they drive our AI transformation.

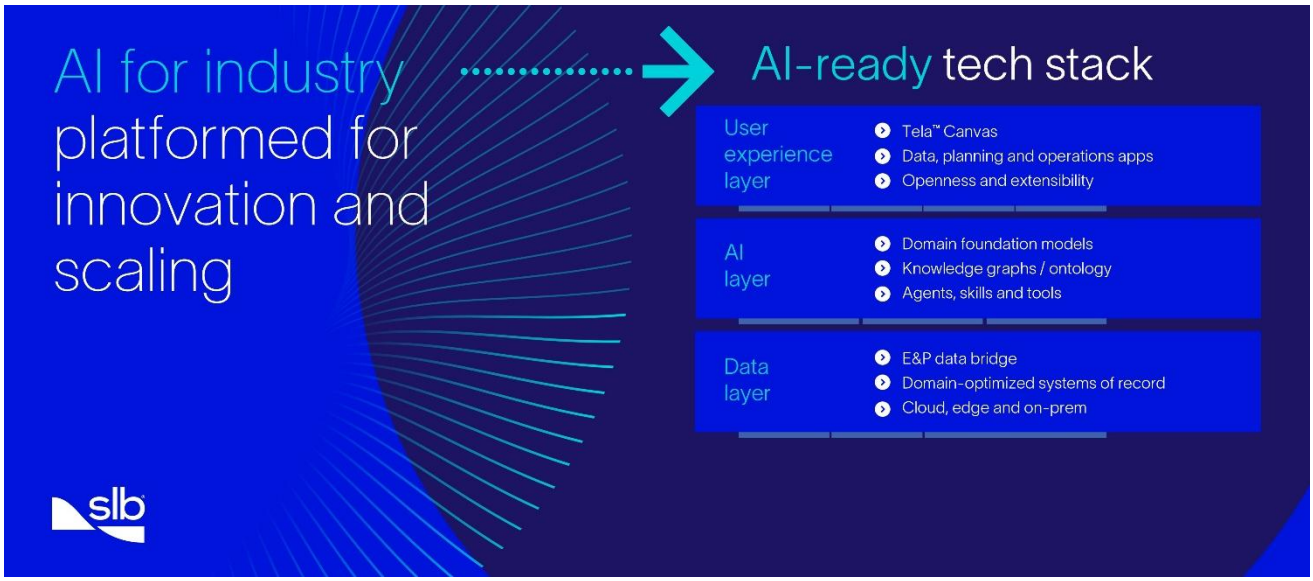

Number one, the AI-ready technology stack. Let's take a look under the hood.

The data layer. I think we will all agree that there is no AI without data. In our industry, more so than in any other, data is all over the place; public versus private, on-prem, on the edge, on the cloud, pretty much everywhere.

We work with mission critical technical data like seismic, reservoir, and real-time operations data.

These are stored on customer systems of records, that are often proprietary, and many outside the industry do not even know they exist. And our customers don't want to move or duplicate data from their systems of record.

In fact, they cannot. As data is core to their business processes – and any missteps can cause serious issues.

We have implemented a unique exploration and production data bridge from the ground up to honor these constraints. Our data bridge is the fabric that allows us to connect to customer data sources without moving or replicating data.

It is designed to work with the many, many data sources and cloud setups of our customers. It is what allows us to search, discover, access, and consume data in our AI workflows.

The AI layer. It is arguably the most important. It is the "magic" dust on how SLB's AI differentiates. And it just works for our industry.

Everyone here is familiar with Large Language Models (LLMs) and the weekly developments from

frontier AI companies. We have seen many of our peers and customers adopt and force-fit these generic, big and unwieldy models into their AI implementations.

I can tell you this. That is the uphill battle.

Trying to decide which LLM to use and worse, using these generic LLMs for technical workflows is like trying to boil the ocean.

To uniquely solve this, we have built proprietary domain foundation models (DFMs). Not one model, but models for several of our petrotechnical domains.

And they are special in three ways.

One. They don't replace these generic LLMs. They work in tandem with any customer preferred LLM. We solve for domain specifics while the LLMs solve for the generic.

Two. Our domain foundation models are purpose-built. Purpose-built because we know exactly how domain data is structured, what to look for in the data, and how to use it to deliver on the user's intent.

And three. Because we know which parameters are important, we know which data to feed it.

We have augmented our curated publicly available data sets with proprietary SLB data. And we have absorbed the IP and knowledge that comes from decades of oilfield services.

So, for our customers, it is simple. Our domain foundation models give them a powerful base model that is continuously updated. They can even refine these models within our platform with their own data to customize them for their own assets.

The third Layer is user experience.

The key here is the Tela Canvas. Think of it as the ultimate industry Copilot. You know how ingrained ChatGPT, Claude, or Gemini have become in our day-to-day activities. Our users will soon find Tela indispensable because it is just as easy to use and is integrated into the products that they use daily.

Even more so because Tela understands their technical context, has strong guardrails that ensure it never goes off-road, and uses domain benchmarks to ensure it stays strictly within the bounds of domain science.

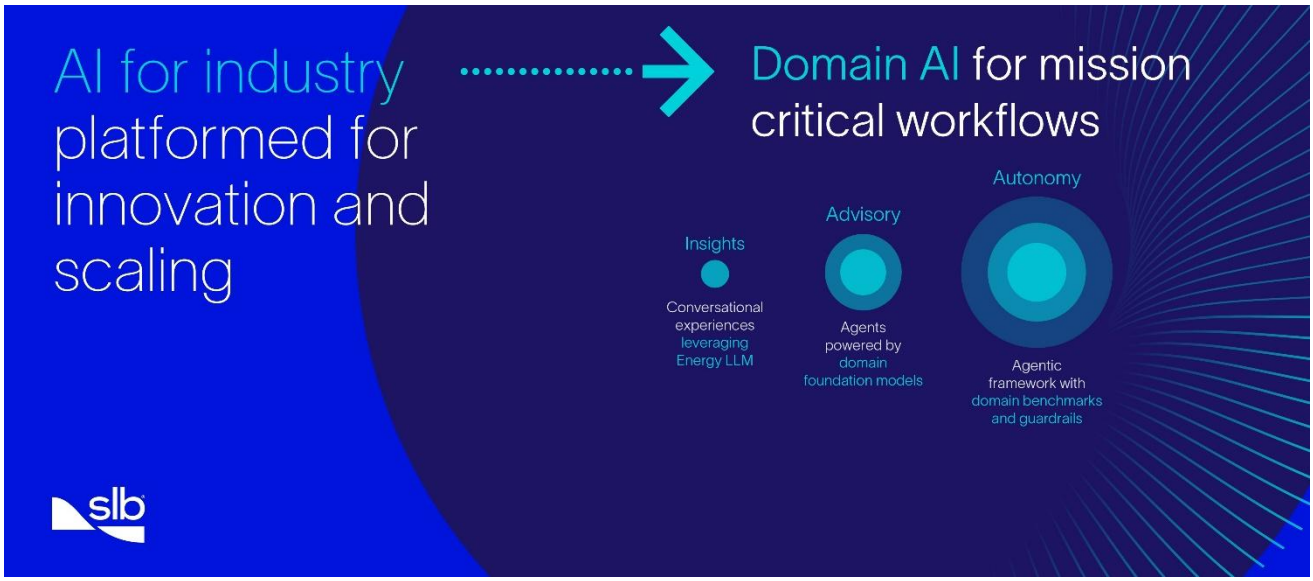

The second key element is our domain AI.

Our AI implementations are structured to meet customers wherever they are in their AI readiness.

For organizations beginning their AI transition, we offer conversational interfaces that extract insights from their data, project histories, and operational context. This is powered by an energy-specific LLM infrastructure that is trained on SLB's technical documentation, intellectual property, and a century of oilfield services experience.

The barrier of adoption is low, and the value for our customers is immediate.

For customers who are further along in their journey, our technical software embeds dedicated AI agents powered by our domain foundation models, skills, and tools. These provide high-value engineering recommendations while keeping the human in control of the decision.

For advanced operators, our agentic framework runs fully autonomous workflows – the observe, plan, generate, act and learn loop – but they are constrained at every step by domain science and physical guardrails.

Because these agents are trained on validated domain data, they offer a level of accuracy and assurance that is unique to SLB.

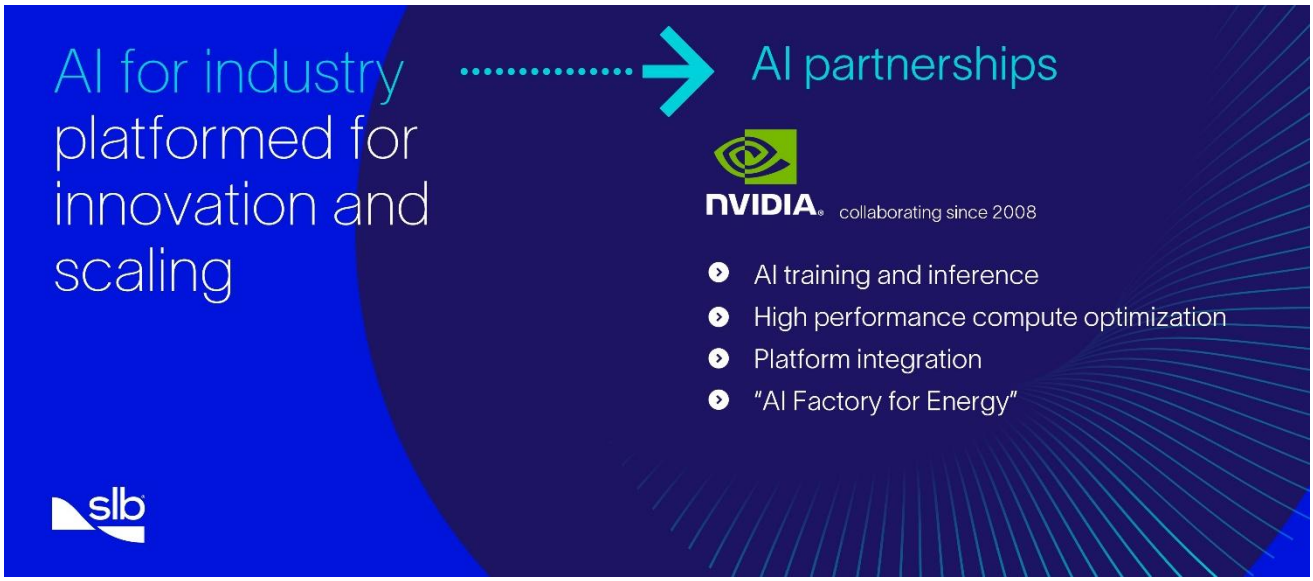

The third key element is our technology partnerships.

You heard from Trygve earlier about the breadth of our partner ecosystem. I want to go deeper into one in particular.

Our relationship with NVIDIA is special.

It is nearly a twenty-year joint engineering program. We have direct access to their top engineers, and they choose to work with us for one specific reason: we bring the unique domain physics needed to push the boundaries of digital in energy.

Together, we are building the "Tela AI Factory for Energy" to bring the most powerful set of agentic AI implementations in the industry to our customers.

Let me be clear about what this deep integration means: guaranteed peak performance.

Every domain foundation model we produce is optimized to be the absolute highest-performing model in the industry.

NVIDIA engineers are actively taking our source code and tuning it to run optimally on today's Blackwell chips, and they are already future proofing it for tomorrow's Vera Rubin architecture.

We are fusing the world's leading AI computing architecture directly with our unparalleled domain data and science.

No one else does this at scale today. It creates a competitive moat that simply cannot be replicated.

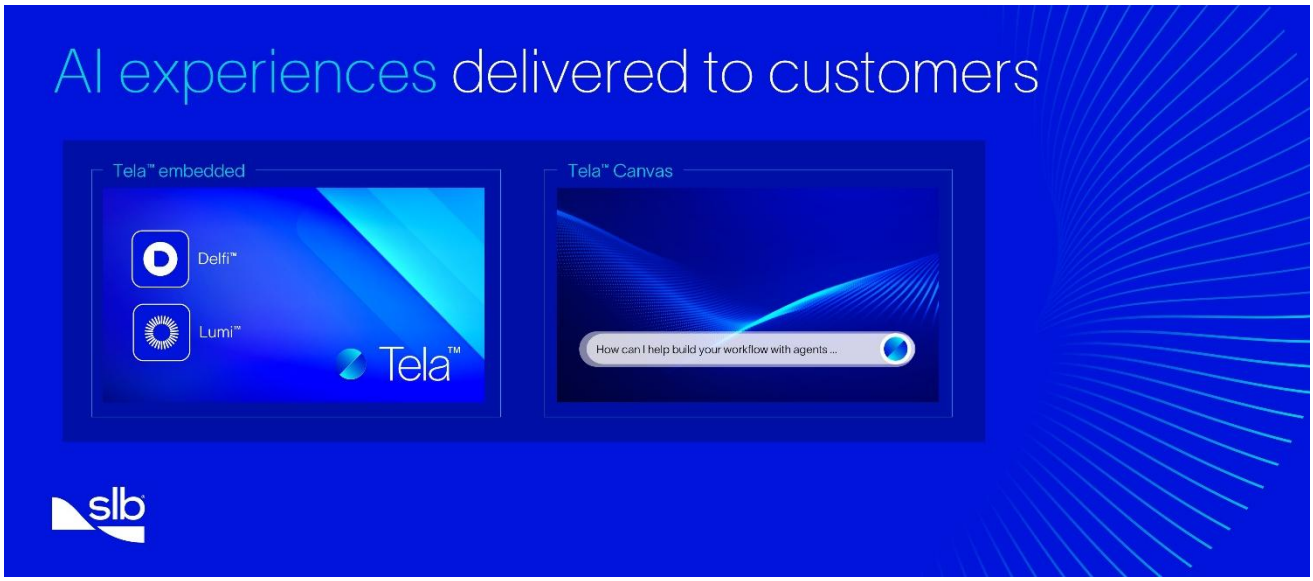

We deliver these AI capabilities through two distinct user experiences; each one designed for a different mode of working.

Tela embedded integrates agentic AI directly into our existing and widely deployed software: Petrel, Techlog, DrillOps, OptiFlow, etc.

Users access conversational and agentic tools natively within the applications they already work in every day. This drives immediate productivity gains and reinforces the value of our software. This provides the indispensability I talked about earlier - it is simply there for them to use.

Tela Canvas is a standalone experience designed for broader, cross-functional use.

It works across traditional application boundaries and data silos, orchestrating complex end-to-end technical workflows by combining AI with the physics-based domain science that forms the core of our software portfolio.

Where Tela embedded enhances individual application workflows, Tela Canvas connects them.

Customers can choose between rapid, transaction-focused interactions via Canvas or deep, immersive engineering workflows within our core applications. Both run on the identical agentic AI backbone; the same domain foundation models, the same physical guardrails, and the same data infrastructure.

Now, let me show you what this looks like in practice, starting with planning.

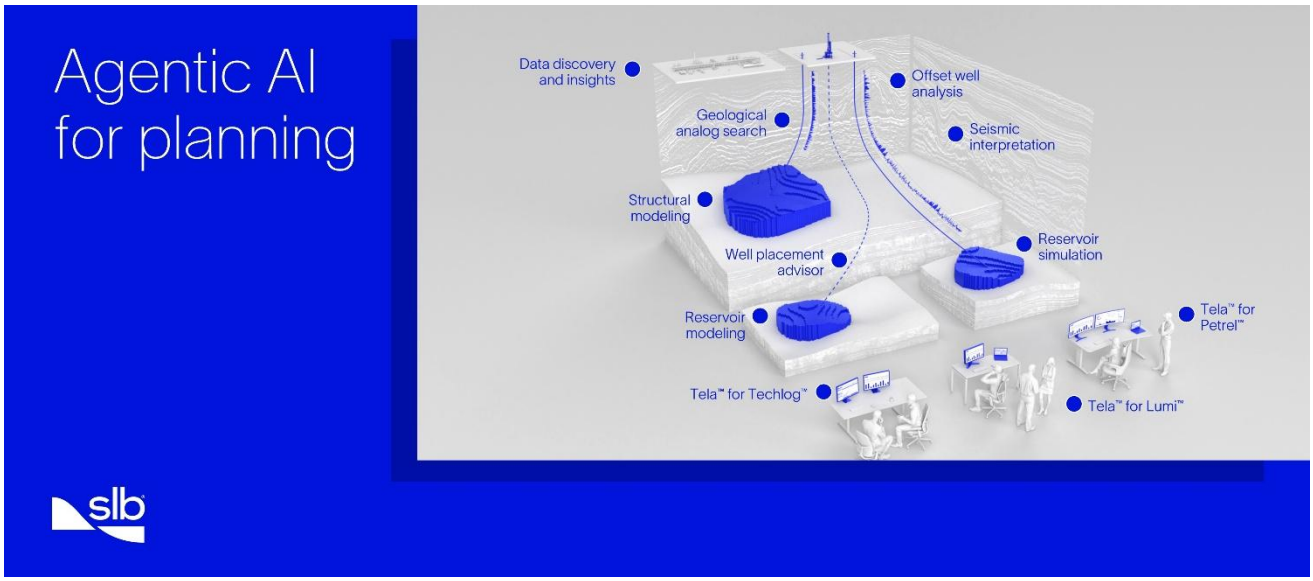

Our planning solutions and workflows are driven by an extensive portfolio of subsurface agents, models, and specialized tools covering the full spectrum of technical workflows used in exploration and field development.

They provide comprehensive coverage of the key domains across geophysics, petrophysics, geology, and reservoir engineering that Rakesh talked about earlier.

Let's see an example of these planning agents in action within the Tela Canvas and connecting to Petrel, our leading subsurface platform.

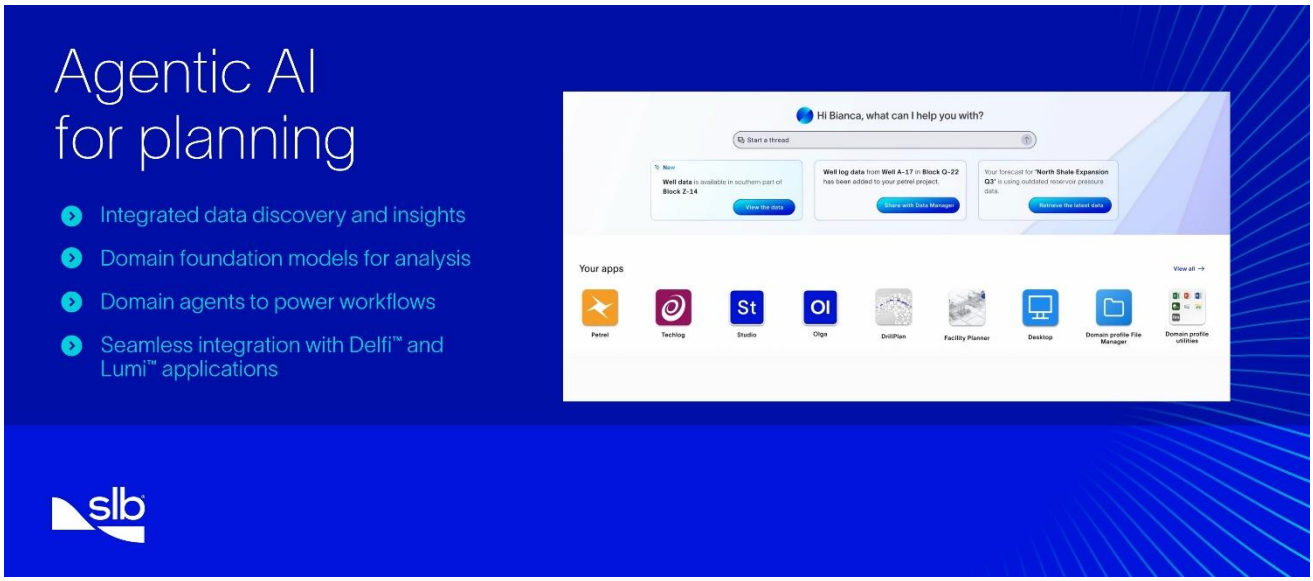

[Tela video demonstration plays]

You just saw a glimpse of Tela in action both as a Canvas as well as Tela embedded in Petrel.

It showed how a combination of domain foundation models and other AI models integrated into a project workflow can transform them.

I want you to think about why this is so differentiating, and why this demo couldn't have been done just six months ago.

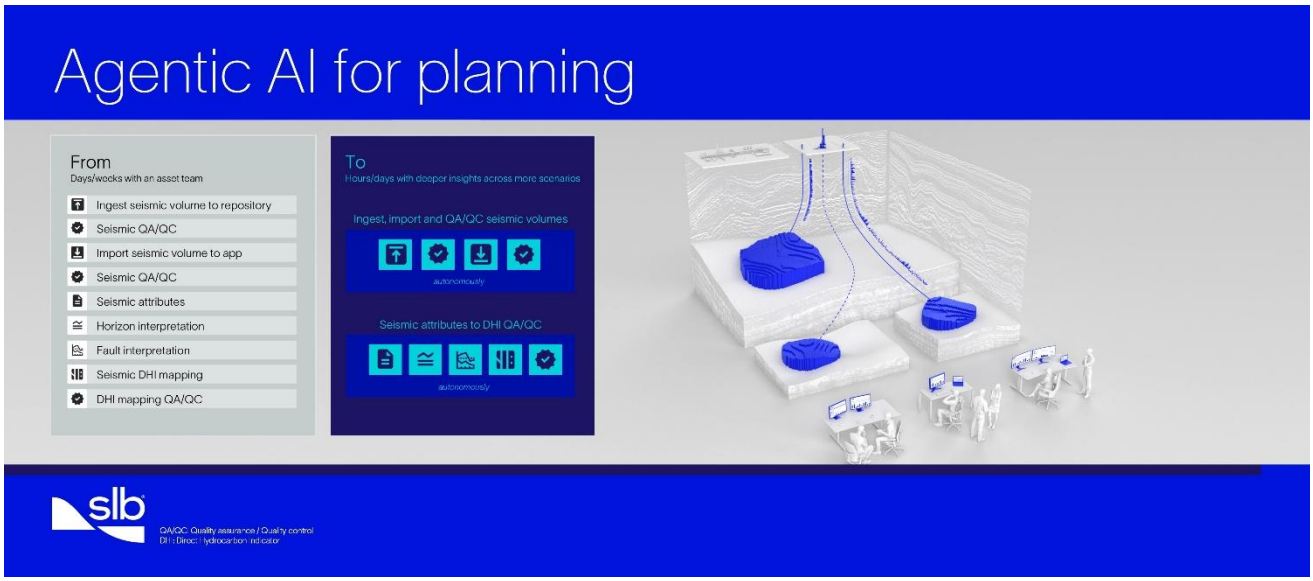

A typical exploration workflow to directly detect hydrocarbon rich areas in the subsurface takes a team of geoscientists weeks to execute.

We are transforming these nine complex steps into two simple clicks – almost magically.

We are compressing weeks of complex analysis into a few hours, while covering a broader range of scenarios than was previously not practical.

Let me tell you the way this works.

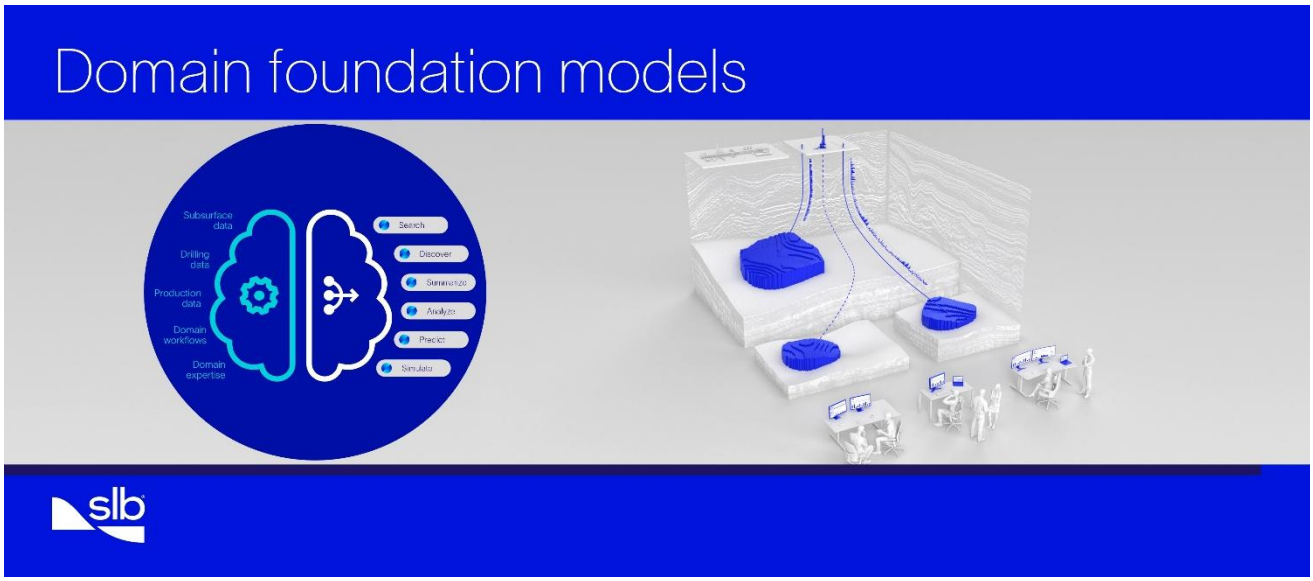

You have heard me say domain foundation models several times.

Let me explain what they are and the key role they play in our AI implementations using the seismic foundation model as an example.

Seismic data is core to most exploration and field development workflows. However, seismic data modalities – its structure and format – are unique to our industry and LLMs are unable to work with them.

We started with a base vision transformer model. We adapted it to handle seismic data modalities. We incorporated seismic and geoscience domain priors and context. And then we trained it with public and SLB multiclient seismic data sets.

The outcome is a rich and capable seismic foundation model with less than 1 billion parameters.

For reference, leading frontier models are already well beyond a trillion parameters.

The smaller parameter count means our models are cost effective to develop and to operationalize. It will also allow our customers to fine tune the model with their data for continued use within Delfi and Lumi platforms.

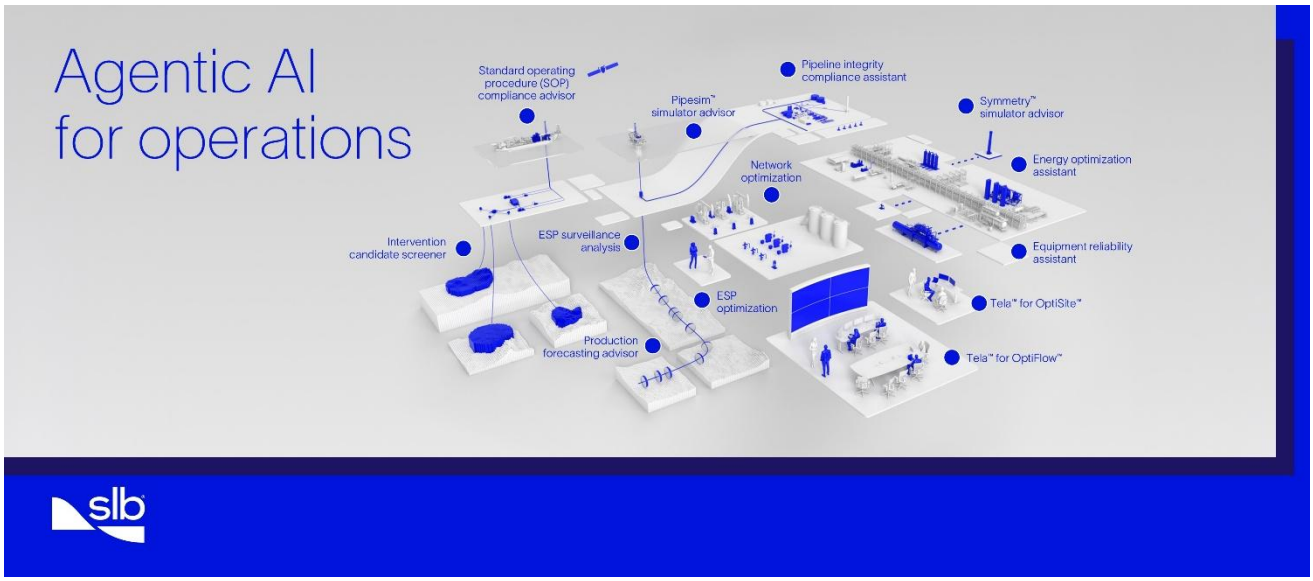

As in planning, our operational execution relies on a dedicated, scalable portfolio of operational agents, models, and tools.

These cover drilling operations in our DrillOps family of products. They enable the increasing autonomy of complex operations such as real time geosteering that Cecilia talked about.

On the production side, agentic workflows form the core of our Opti suite of production offerings.

These range from lift operations in wells, to pipelines and networks, and to complex facility operations like floating production, storage and offloading vessels (FPSOs).

Let me exemplify this from a deployment for a customer in the Middle East.

In traditional operating models, engineers must manually trigger, review, and process data to manage facility equipment.

This is incredibly frustrating and monotonous for a production engineer while for the company this means operational adjustments can only be made when there is someone at the desk.

Tela converts these manual actions into autonomous, evergreen loops, reserving human intervention for high-value capital decisions.

You might say “so what”?

When we talk about a production asset, we talk about a lot of equipment. These come from different vendors, are of different vintages, each with different sensors, working with different parameters and data formats. All being used differently in different conditions.

Leveraging agentic AI is the only way to reach operational autonomy at scale while retaining human oversight.

We can do this because our equipment digital twins are trained on real-world physics and validated through continuous iterations.

We can do this because we are OEM agnostic and can support equipment from multiple providers. We can do this because we can model and simulate at the asset level. These enable round-the-clock facility optimization, directly lowering operating risk, minimizing unplanned downtime, and maximizing barrels produced.

You will soon hear from Stephane about our broader monetization strategy but let me be bombastic for a moment.

Every single agent, every model, every tool that a customer uses within our ecosystem — we have implemented the platform so that we can monitor their consumption every single time, at scale.

That is exactly why and how we can monetize AI.

We do this through four distinct channels.

We've already talked about Tela embedded into our existing, widely deployed software platforms. We monetize through up-selling Tela subscriptions and then consumption of agentic workflows.

Tela Canvas opens a new channel for monetization where a baseline subscription paired with consumption-based pricing allows us to cross-sell to the existing customer base. It also creates new sales opportunities with those customers who find adopting platforms like Petrel to be a heavy lift.

Going beyond these two channels, we also develop fit-for-purpose AI workflows for our customers through our Innovation Factori model, a network of seven AI centers of excellence around the world.

These solutions are then deployed within the customer's Lumi and Delfi environments, creating long-term platform stickiness and ongoing consumption.

And finally, the SLB Digital Marketplace that we announced on Monday.

It enables a platform business model for SLB and verified third-party developers to offer specialized agents, models, and applications that can be deployed in our digital ecosystem. We monetize through revenue sharing driven by consumption; a high-margin, scalable channel that grows with the ecosystem itself.

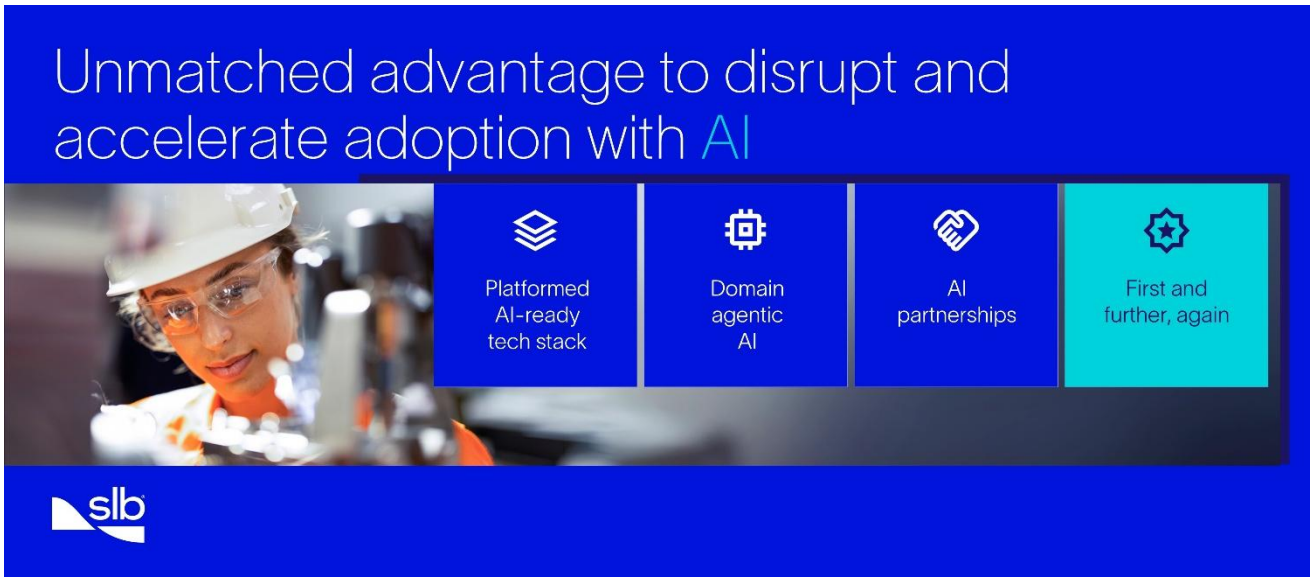

As I conclude, I want to emphasize one critical truth.

Today, no other company in our industry can do what SLB has accomplished.

We did not just ride the wave of generic AI models. We built proprietary domain foundation models from the ground up.

We did not ask the industry to completely rewrite their operations. We embedded intelligence directly into the applications that tens of thousands of geoscientists and engineers trust every day.

And we have created an agentic framework that is open for our customers to extend, yet rigorous to operate autonomously within physical constraints.

This combination of proprietary models, trusted platforms and deep domain science, augmented by strong digital partnerships is what positions SLB to lead the commercialization of AI in our industry.

And it is a position that will allow us to be first and go further once again.

Next, let me welcome Stephane to the stage to discuss the financial impact of our Digital business.