Standard Chartered's Investor Event
Tuesday, 19 May 2026
(Amended in places to improve accuracy and readability)

Delivered by:
Tanuj Kapilashrami, Group Chief Operating Officer, Standard Chartered
Noelle Eder, Group Head, Technology & Operations, Standard Chartered

Tanuj Kapilashrami: Bill started the session today by saying that Standard Chartered has a very clear ambition. We want to be the world's super-connector, not just a global bank but a global financial network, solving for cross-border needs of our clients, connecting capital, trade and payments across the network in a world that is becoming increasingly fragmented.

Delivering on this ambition requires more than geographic reach. It requires an operating model encompassing people, processes, and technology that is interconnected and purpose-built. An operating model that can be scalable, reusable, and that is standardised. That's what I'm hoping we are going to be talking to you today. The world is not standing still, and neither is our response.

Our transformation is not technology-led for the sake of technology. It is strategy-led with increasingly sophisticated interplay between people, processes, and technology to deliver very differentiated outcomes for our clients, colleagues, and for our shareholders. Noelle will get in a minute to talk about the tech architecture, but I really wanted to highlight the fact that this is not just a tech story for us. It's a process, people, technology story. That's one of the key reasons why Noelle and I have chosen to do the session jointly today.

I want to be clear from the start - and Bill said this as well - that this is not a defensive cost action for us. Our whole transformation agenda is about creating operating leverage to deliver exponential growth. That's the objective of our transformation work. There are cost targets that I'm going to get into, which is an outcome of the work, but that's not the real objective. A lot of the work that we are going to be talking about today is the work that's already been done, which is enabling growth, and how we feel, by enabling AI on top of it, that growth is going to be delivered further.

The other thing that Bill said, which I want to double-click on, is that we are a global bank. We don't have one or two home markets. Our network is our home, and that is one of our biggest structural strengths. So, to build an operating model that leverages the value of the network to deliver on our super-connector aspiration requires a response, which is a very distinctive response. So, like I said, we are going to be telling you today the investments that have already been made, the outcomes they have achieved, and what happens next, especially with the advent of AI.

At its heart, our ambition is to deliver a bank that is simple, connected, and fast. Simple means global core platforms, fewer variants, and standardisation where it creates scale. Connected means shared cross-border capabilities that serve multiple markets. Fast means executing at pace but - Jason will love this - with very, very clear guardrails. So, that is simple, connected, fast. We're going to talk about simple, connected, fast a lot today.

We are not measuring our transformation just by program volume, KPIs, end dates, et cetera. We've got very clear financial metrics that we are linking this to, very clear commercial outcomes. A 20% increase in income per employee by 2028, greater than 15% reduction in our corporate function headcount - this is technology, operations, all of our support areas - resulting in, ultimately, a structurally lower cost base, so a cost-income ratio of 57% in 2028.

Our transformation is not new to us. This has been a journey that's been on. The first slide that Bill flashed today in his presentation outlines the financial outcomes. We started our transformation by tackling our operating model, because we firmly believe that technology follows the operating model. If your operating model is complex, technology will be complex. I think, candidly speaking, if you go back a couple of years, we did operate in a matrix with a very strong local orientation.

What that did for us, is it resulted in duplication, inefficient capital allocation, and unscalable investment decisions. So, one of the big pieces of work that's been happening for a long time, predating Fit for Growth (FFG) but accelerated by FFG, was the work that we have done in simplifying our operating model. We stripped out regional layers, radically simplified our executive and leadership bench, and clarified decision rights between our global businesses and market leaders.

To just give you a sense of numbers, we moved in my time in the bank from eight regions to four and now three. We've got our market CEOs double-hatting with one of the two global businesses. Today, almost 80% of our market CEOs double-hat with one of the two global businesses. One of the big changes we did was aligned all of our markets to one of the two global businesses. So, this is not just a reporting line change. This is the way we do capital allocation. This is the way we report on performance. This is the way we take investment decisions. It's been a pretty fundamental change.

What this has done - and again, not just a cost outcome. What this has done for us is accelerated decision-making. It's reduced the path to decision-making, but what it's also done is brought accountability much closer to our clients. That's been the two big outcomes of the operating model changes that we have done. One number on this slide I want to double-click on is the size of our workforce in our global capability centres. We call them GBSs, Global Business Services. 43% of our headcount now sits in one of our capability centres.

Again, this is not just a defensive cost action play. What it does for us is, by co-locating a critical mass of our processes and people, it puts us in a position to be able to standardise, automate, and deploy AI at scale. Far more difficult to do when you're trying to do it in a much more geographically fragmented model. So, these organisational changes have been the foundation of everything that we have done until now. With that, I'm going to pass on to Noelle, who's going to talk about modernisation of our tech architecture.

Noelle Eder: Thank you. Hello, everyone. Process of elimination, I'm Noelle. Here's the good news for you today. Bill and Manus were so excited about the talk that I'm about to give that they gave a good portion of it already. So, what you all know, is we have BCP in our management team. In case anything happens to me up here, we know they can step in. So, I've been in technology and operations for a little more than 30 years, and so I've been leading change across technology cycles and across multiple organisations in different sectors.

So, I know when you hear the word transformation, it sounds a bit lofty. Maybe there are a couple of sceptical people in the room. What Tanuj and I are going to share with you today, I think, based on my experience is a very

credible story about people and technology coming together and changing what this bank is capable of doing. So, I'm going to start by talking about significant work that's been underway for the last few years - Bill talked a little bit about it, Manus talked a little bit about it - where we have been modernising our technology foundations for resilience and growth, and that work is now starting to pay off.

So, Bill talked about our journey in cybersecurity and our journey in financial crime or anti-crime. We then focused on three additional pillars. All of these things together form the very core of this bank. So, those three critical pillars are what I'm going to talk to you about for the next few minutes. First, in the fourth quarter of last year, we completed our global private cloud. We now have 10 times the processing horsepower that we had previously, and we have geographic resilience against domestic disturbances, datacentre outages, subsea cable disruptions, climate issues, with unprecedented levels of automation.

We have redesigned and rebuilt our connections to third parties. We use a concept called abstraction, which means that we can move them to any location. So, put simply, we can be in any market quickly, in any market with internet access. So, that's the first thing, and I want to pause there for just a moment because that wasn't just a feat of engineering. It was actually a feat of engineering, but to have a global private cloud, to have geo-resilience and to have third-party flexibility the way that we have it, is really quite differentiated and very, very important for a network bank like ours.

Second, in March of this year, we migrated our largest market - which we're all sitting in, Hong Kong - from the mainframe to our global core banking platform. We now have more than 90% of our markets on a single platform. What that means is we can build, we can test, and we can deploy code much more quickly. Because we own the source code, we are no longer beholden to third-party roadmaps or geopolitical tensions. We can and have started integrating digital assets, blockchain, and AI natively, on our own terms at commodity prices and sustainable scale. Okay, that's the second thing.

Third thing, our modern payments platform. So, we're at the infrastructure, core banking, now we're up at payments. Payments is an out-of-the-box service offering the same levels of choice and service excellence consistently across markets. What this means is that our businesses can open new digital services, new payment corridors, and facilitate trade quickly and easily. So, those are the three things that we've been focused on for the last several years. Significant investment has gone in, and you might be able to tell that we're quite proud of them.

So, these systems are future-ready and resilience is built in, and I'm going to give you four reasons why that's true. First, they're commodity based. So, what that means is they're scalable, they're standardised, they're performant at the best possible price. Second, they're cloud-native, so we are significantly moving away from mainframes. Third, they're API-first, so no more bespoke integrations. Last, the source code is owned by us, as I already mentioned.

Now, we did not set out to rebuild everything everywhere all at once, even though there might be somebody in the room who published that just a little while ago. We focused on creating operating leverage for the bank by targeting the systems, modernising the systems that all of our products rely on. We made them industrial-grade and future-ready. Now we can extend up above that core, where differentiation matters, at the client interface level. So, we're standardising it at the core and we're differentiated at the client edge, and the balance between the two is critical. Standardisation gives us speed and efficiency, and specialisation and configuration gives us new products, new services, and new experiences for clients.

Our transformation program doesn't have an end date. It is increasingly the way that we operate. We've gone from large, infrequent releases to small, frequent changes every day. Feedback loops, whether they're customer, operational, or risk are built into these platforms, and data is not an output. It is a live signal. What that means is, it streams and it's analysed in real time. It drives decisions and it drives improvements every single day.

Now, let's talk a little bit about the results we're seeing from all of the work that I just described. So, as Tanuj mentioned earlier, we've been on a continuous improvement journey. So, over the last couple of years, we've seen a 30% improvement in operations in throughput per FTE. Our digital services are now always on in any market, 24/7. We've seen a ninefold increase in transactions processed per second, supported by, as I mentioned earlier, a 10x increase in processing power, and our downtime has been reduced by 80% while run costs have remained flat.

So, we think our results are getting better. We expect more. As you probably heard from Bill, we're constructively dissatisfied fairly frequently. So, the key point is this. Scale, speed, and resilience are no longer constrained by linear cost increases. For a network bank, that's pretty transformational. So, that's our modern foundation. Now we're going to talk about the people and processes that sit on top of it and bring it to life, because tech for tech's sake is not what we're about. We are an applied technology company, and people and process bring this to life for us.

Tanuj Kapilashrami: Thanks. Thank you, Noelle. Continuous improvement is a word that we have spoken about quite a bit. Our transformation is not a program with an end date. We were transforming before Fit for Growth, but Fit for Growth was a very important accelerant, and we can pick up your question when we get to FAQ later today. We are on track to deliver $1.3 billion, which is going to be a 1:1 ratio in terms of spend and cost saves. Over 300 initiatives across all parts of the organisation, driving improved customer experience, increased straight-through processing rate much faster turnaround time on our applications.

Perhaps what's been less visible on Fit for Growth is the amount of investment we have done in securing our plumbing. I wasn't going to use the word, but then Bill used it in his opening, so I am going to use it. We have mapped 100% of our processes in the organisation, and we have mapped a consistent set of skills that underpin all of those processes. That roadmap, heat map, blueprint that we have developed becomes incredibly important when we talk to you about AI deployment going later.

So, 100% of processes mapped, consistent set of underlying skills mapped, which gives us a very good sense of how work gets done. Some really clear outcomes on the slide. ~53% straight-through-processing rates increase in wealth solution. What's not here is a 10% reduction in our technology estate that has happened because of all of the work that we have done. So, what Fit for Growth has helped us do is build that muscle of continuous improvement. So, it's not just a one-off efficiency gain, but a muscle that sustains beyond the program, which is going to finish by end of this year.

So, moving on, we call this the bridging slide because we will talk about, having secured the foundation, what happens to us next. I will go back to why simple, connected, fast is not just a transformation tagline for us. It's an economic logic that helps a super-connector bank deliver to its full potential and beyond. So, again, I've said this a few times, we are not anchored to a single market, and for us that's a structural strength. That means we are not constrained by one growth cycle, one domestic balance sheet, one regulatory environment. Quite the opposite. We are very well positioned to be able to capitalise on opportunities that arise anywhere in our network.

To do that, we do need an operating model which is simple, connected, fast, which basically means we want to scale without multiplying cost. For that, the bank has to be simple. To harness the power of our network, the bank has to be connected. To do this safely and repeatedly at the speed required across the bank, across the globe, the bank has to be fast. What that does collectively, for all of you, is high income growth, greater operating leverage, and a workforce that is much more upskilled to be able to compete with the future direction of the bank and compete with the future of banking.

So, like I said, not just a transformation tagline for us, but an economic logic that helps a super-connector bank deliver on its aspirations.

Noelle Eder: So, let's go just a little bit deeper into how we're transforming to be simple, connected, and fast. We've modernised our technology through a simplification approach that retains optionality, standardised at the core, differentiated at the client edge. Our principles are straightforward. Build once and deploy across markets and businesses, differentiate through configuration rather than custom code, and reduce fragmentation. This lowers the unit cost to serve, it improves resilience, and it makes growth more scalable, easier to achieve.

Let me just give you a couple of examples to try to bring this to life. The first one is the one I've mentioned already. Standardising on commodity infrastructure, one pattern across the estate. Second, we have a single identity and access management layer for Standard Chartered. Three, we have one payments backbone for the Company, and fourth, we have a single enterprise AI platform with reusable services.

In operations - this is a good example - we've unified over 100 applications into 30 standardised workflows accessible for our operations people through a single user interface. Over time, that interface becomes a single pane of glass to manage client onboarding, servicing, risk and governance workflows, and more around the world.

Tanuj Kapilashrami: What it really means for our client is that a capability that we build in one market, let's say a real-time payment solution, can be deployed in multiple markets at a fraction of cost and in a far more speedy deployment way. So, that's the real value of the model that we are doing. We are constantly simplifying, standardising our processes and data, and a lot of that, leveraging our Fit for Growth investments, has happened in our capability centres.

So, we have demonstrated reduced timelines, more efficient processes in our KYC, onboarding, customer due diligence processes.

Noelle Eder: So, we've talked about simple. Now we're going to move on to connected. So, the power of a super-connector is in the connections. A super-connector creates network value, so a single capability can be reused across clients, products, and markets. Our platform architecture is designed so that one client engaging with us in a single product in a single market has access to our full capability globally.

So, trade finance will link to payments, and cash management will link to foreign exchange. The client doesn't see seams. They see a network. The real unlock over time is deeply integrated client and transactional data, which gives us advanced analytics, more precise personalisation, more seamless payments across geographies, more straight-through-processing for our teams. That reduces friction, it improves the client experience, and it helps our businesses deepen relationships with their customers.

Tanuj Kapilashrami: So, basically what happens is, clients don't see product market boundaries. They see one Standard Chartered. At the beginning of the slide, when I spoke about the work we have done on our global capability centres, that's a really good proof point that helps deliver on connected. By centralising processes and people in big, shared service centres, we are able to drive that connectivity across the network, which was not possible in a geography-by-geography model. So, a really good example of leveraging our shared service centres to drive that level of connectivity.

Noelle Eder: Agreed. So, we've talked about simple. We've talked about connected. The two together have a multiplicative effect on making us fast. So, we're not just digitising. We're building the capability to operate in a fundamentally different ecosystem. Bill talked about it this morning. Agentic commerce, the digitisation of money, the speed at which financial services will move, I think in the not-too-distant future, is something that we are preparing ourselves for.

The winners will be those who can sense, decide, and act in real time, dynamically changing products, dynamically managing risk, and strengthening with scale. The most important part about this ecosystem from my perspective - and perhaps this is somebody from technology talking - but serving customers has to be quick, easy, and fast, whether they're human or machine. We've referred to Jason in this conversation. He happens to be our Chief Risk Officer. Always with the guardrails that this industry requires.

Architecturally, we've separated our foundations from our product delivery, and we're running those foundations as utilities. So, our run costs are predictable, time to market is faster, and scaling across borders is easier. One example for you. In technology delivery, we've moved from 18 manual approvals to get code across our markets to clients. We now have an automated pipeline-based process, so that's the from-to in the kind of ecosystem we're building, and we're preparing all of these systems for AI by rolling out standards-based APIs across them all.

What that means is that these systems can be orchestrated by humans and by machines. The result? Our product owners can make changes much more quickly, and be able to respond to the dynamic markets that we operate in.

Tanuj Kapilashrami: So, in summary, simple reduces friction and risk, connected removes latency, and together they multiply so we can execute much faster, at scale, and with the right level of control. Just leaving you with some numbers before we move to AI. Our clients are already seeing results. 97% of our tech releases are now fully automated. We are deploying products 30% faster than we have done previously, and we have reported 28% - Judy will share some of the customer satisfaction data later - but 28% reduction in manual client payment queries in a relatively short period of time.

So, that's been the impact of this work that has been seen by our customers, all of this resulting in much better client experience, which is what this work is in service of ultimately. The foundations are now in place for us, and the focus goes from building to scaling. That's where the beauty of AI comes in.

Noelle Eder: Yes. So, now we're going to talk about AI, and I wrote down some of the questions that were asked so, if I don't get to them, we're going to have a little session at the end where you can ask us questions. So, the scale that Tanuj talked about really only matters if complexity is reduced and it stays reduced. AI exposes complexity immediately through data quality issues, through hidden dependencies, through technical debt, Its capability

curve is exponential, not linear. So, the pace is not just fast. This thing is structurally different from any technology advancement that comes before it.

So, we are increasingly simple. We're more standardised and more connected across platforms. At our core, we are increasingly one unified financial platform across 54 markets. That creates a simple ambition that we work on every single day. Build a bank that gets better with every transaction and every client interaction.

Our next frontier is to take that AI platform that Bill described and embed it deeply into the foundation as a structural capability, not as a bolt-on. So, the bank, in every transaction, learns from each one, and we're doing that through three reinforcing elements. First, one enterprise AI platform, as I mentioned earlier, with intelligence built in, so AI can scale consistently across markets and businesses. Second, an operating model that identifies and automates routine, repeatable transactions while keeping human judgment where trust matters most. Third, a data and AI architecture that learns, improving accuracy, improving cycle times, and improving marginal cost as we scale.

So, what this enables is a business that can sense change, adapt faster, and act more quickly. So, let's double-click on the platform for just a moment, because it has three key characteristics to it that are really important to us as we continue to scale it. First, risk management. As I mentioned we are, after all, a bank, and so regulatory change, regulatory standards, the control environment are critical to us, and they absorb significant capacity across the organisation.

An AI platform in a network bank in particular must be able to codify controls, monitor them through automated guardrails, and be able to absorb regulatory change that happens very frequently with ease. Second, productivity. Our AI platform can handle routine initiation, validation, approvals. Again, that frees up people to focus on trust, clients, advice, relationships. Third, precision. In my opinion, this is the game-changer for an AI platform, because it's all about data and the quality and accessibility of data.

So, we're moving from 150 fragmented data lakes across the state to a single global data supply chain that's cloud-native, standards-based, with compliance built in. Better data improves model accuracy. Model accuracy gives us more straight-through-processing. Straight-through-processing frees up capacity for clients and for growth. So, what you see on the slides behind me are really just examples. They're not an exhaustive list, but they are the kinds of outcomes that we're experiencing now as we scale into this capability.

We can see this very clearly across the bank as AI is applied end to end. In payment operations, our AI platform can route routine transactions while surfacing exceptions in context so that people can make decisions more quickly. That improves our client experience, and it lowers our unit cost per transaction over time. AI-enabled software engineering. So, in our SDLC - software development lifecycle - AI is helping us develop faster, increasing our time to market with fewer defects across the estate.

We've also deployed Copilot across the enterprise - you can see the numbers behind me - and we are targeting work that has low value where time can be reallocated to higher-order, judgment-led, client-focused work. So, for us, the logic is very clear. We lower run costs. We lower production incidents. We increase time to market and capacity for differentiation.

So, stepping back, a platform changes the economics of change itself. Each new capability builds on what's come before it, so AI deployments become faster and cheaper over time. It also changes how work scales. So, growth no

longer requires proportional cost increases because machines handle speed and volume, while humans focus on clients, trust, relationships, and growth. That creates flexible, AI-enabled capacity, and it's a powerful lever on cost to income over time.

Because the system learns, the economics improve in a reinforcing loop. Better data, better models, more straight-through-processing, more capacity free. So, the investment thesis is straightforward. We're building a bank where the marginal cost of growth declines with scale, and we believe that's what an AI-native enterprise delivers. It's already underway here.

Tanuj Kapilashrami: Thank you, Noelle. This brings us to a really critical inflection point, and I'm conscious we are running on time, so I'll make this very quick. We have secured our foundations. The work now is for us to accelerate transformation by embedding AI into our business processes. That's the work that we have already kickstarted, and that's really going to be the focus over the next couple of years.

A key enabler of this work has been the work we have done to identify our processes and pivot the organisation to becoming a much more skill-based organisation. This is not a tagline, but by deconstructing work into a set of consistent skills and identifying the activities that sit under the work, we can be very precise on what gets automated, what gets augmented, and what needs to remain human-led. That's really the embedding of AI into our processes that sits on top of the foundation.

So, just to bring this to life by giving you some examples, in Singapore and India, where we have more than 50% of our hiring demand, we have launched agent-enabled employee onboarding. So, AI agents coordinate end-to-end onboarding, manage tasks, data, exceptions across systems, while humans retain the final decision on risk management and the end decision of the process.

This work has reduced hiring manager effort by 35% and has led to a value creation framework which we are deploying into sizing the size of our HR operations outfit. So, it is a really tangible example of deconstruct work, decide where we deploy agents, what gets left behind on humans. That is the process that we have been deploying across multiple of our processes.

Treasury is another area where there's been some fabulous AI deployments that we have done. So, we have driven AI-driven decision-making in enhancing liquidity management with faster, clearer insights, demonstrating how we can scale AI, deploy AI in complex, regulated environments. Across both of these examples, returns compound as each AI deployment accelerates the next, and that's when our economics start improving. That's when we structurally start decoupling volume growth in our businesses from headcount growth. So, that is the way we see - we already are, but we see - deploying AI across our business processes.

So, look, bringing you back to where this all started, our transformation. It's a journey. It's a continuum. What we are sharing with you today is a point in time as we see it today. What are the outcomes? A very clear set of productivity metrics. Bill's alluded to the numbers. I won't go through the numbers, but you can see it. A 20% increase in revenue per FTE, at least a 15% reduction in back-office headcount over the longer term, and a structurally efficient organisation with a much-improved cost-income ratio.

There are three messages before we open for Q&A that Noelle and I want to leave with you today. First, we have delivered significant structural change in our organisation, simplified the organisation, modernised our core, and addressed complexity and risk. It's creating a more resilient and scalable foundation for the bank.

Second, our transformation is about building an agile operating model, an agile operating system. This is what allows us to operate as a single global network so we can scale without rebuilding cost every time. The third, we are not standing still. We are continuously improving, innovating, and harnessing AI to accelerate better outcomes for our clients. Thank you very much. Noelle and I will take any questions that you might have for us now.

Q&A

Company Representative: Thank you, Noelle and Tanuj. So, we're going to do questions, but I think we already had - Perlie had asked a couple of questions on this, so why don't we just take those first? There's one around Fit for Growth, so how this was different to Fit for Growth? So, maybe if Tanuj could take that. Then there was a second one which was around AI, and how much AI adoption is embedded into the revenue FTE targets that we've announced. So, Noelle for that one. So, Tanuj first.

Tanuj Kapilashrami: I exceeded the time limit on my presentation to answer your question in more detail. We are going to complete FFG by end of the year, and I think the point I was hoping that I made, it's not just on the $1.3 billion that we are delivering, but it's the investments that have been made in our foundation work that we've done, including a lot of the core investments that we have made on AI. So, things like AI factory, the foundational work on AI, did come out of FFG, so it was a very important accelerator. The program will finish by the end of the year, and we'll give back $1.3 billion in terms of cost savings. Thank you.

Noelle Eder: Let me ask if we answered your second question already, or if there's a nuance to it you'd like us to address.

Perlie Mong, Bank of America: No, I think you've very much answered it already. I think it's just about - because AI is moving very fast, right? It's one of those things that we are all trying to adapt to. In terms of the way you think about cost, at what point when you see the next version of Claude, or whichever technology comes through, how do you think about investing in that versus maybe just working with what you have?

Noelle Eder: Yes. Okay, so you want me to take that one? Okay. So, let me talk about how we think about our AI stack, and maybe that will help a little bit. So, at the foundational layer, I've mentioned commodity. I mean, we're very, very interested in ensuring that we have best possible pricing for GPUs, et cetera, and so forth. Above that, at the data layer we have a partnership with Databricks to help us really get to that high-quality accessible data in the global data supply chain that I talked about.

Above that is the model layer. We are agnostic at the model layer, and the reason we're agnostic is because we agree with you. The capability curve here is like nothing we have seen before, and so the idea that we're going to out-innovate the market, or any particular third-party relationship we have is going to out-innovate, the market is unlikely. So, we deploy a concept called portability as much as possible. We are not perfect, but we deploy portability as much as we can. What that means is, there are architectural patterns where you can move your models from one place to another.

So, we try to help our businesses be agnostic and be able to ascertain for themselves, based on the value created for their clients and for their businesses, which model is best over time. We don't think for a moment that the world will remain static.

Nick Lord, Morgan Stanley: Thank you very much for the presentation. The first is just about Mythos and how you are thinking about that. I mean, I presume you've not seen it yet, but I'd be interested. From what I do hear there's quite a lot of work that is required once you've seen it, so I'd be interested to know how you're thinking about that.

Second, you might've partly answered this already with Perlie's question, but which models are you using at the moment? Are you using open-source models as well as your Anthropics and your GPTs, or are you using some of the Chinese models as well as some of the US models? I'd just be interested in more detail on that.

Noelle Eder: Sure. So, Mythos. We've talked a lot about this as a management team. We've talked a lot about it with our Board, and obviously with our CISO. I think - we think- Mythos can best be interpreted as a signal in a much broader trend around the curve on vulnerabilities. I think everybody knows what Mythos is, but just in case you don't, it's a frontier model that reports to be able to detect and move to exploitation on zero-day vulnerabilities.

So, the bank has a two-part response. The first part is very operational in orientation. We have vulnerability management practices. We have software development practices. We have a defence in depth set of controls implemented in the bank, and so we continue to tune those and advance those. Both our software development organisation and our architecture and cybersecurity organisations have been on top of this trend since long ago.

So, Mythos is a point in time on a curve, but the trend has existed for quite some time. So, we have really been putting our energy into shifting left across the capability curve, so software and how AI is used inside the software development lifecycle, as I mentioned earlier, to help our software developers sort of identify, predict, determine where they might need to resolve something before it goes out the door.

From a more strategic standpoint, vulnerabilities take advantage of complexity. They take advantage of technical debt, legacy, architecture, et cetera. So, you heard us talk today about our technology and really our bank gameplan, which is to resolve the questions about technology life cycle funding and obsolescence and really take that estate to the next level of modernisation. That is as large a defence against vulnerabilities as anything else.

So, our cybersecurity perimeter defences, our threat intelligence, is really, really quite strong. We have a defence in depth strategy. We're shifting left, and then from a technology standpoint overall, the strategy helps by eliminating complexity and legacy technical debt.

Tanuj Kapilashrami: Nick, can I just add? One of the points I made when we were talking about FFG is the 10% reduction in our tech estate as one of the outcomes of the work, just in the last two, two and a half years. We are doing the same with very clear targets on reducing our third-party suppliers. So, this idea is, how do we contain that estate in a way that we can mitigate against the risks better

Noelle Eder: Yes, and it's a great point, because the third-party ecosystem is obviously not under immediate and direct control by Standard Chartered. So, advancing our third-party security assessments in partnership with Tanuj's organisation, who has procurement and relationship management for us with third parties, has been an integral part of the last year. We're well served by it now, in my view.

So, our businesses predominantly make choices around models because they really create the business cases. They create the value propositions. They are closest to our clients and can really ascertain. So, my suggestion would be

that Judy is going to be speaking this afternoon. She's very well versed on this and can discuss it at length in terms of WRB and what we're doing there with regard to models.

I would just say to you, yes, in answer to your question about which models are being used. We're quite agnostic. We're a pretty large bank. We have a large number of businesses, and they have different needs and different interest levels. So, over time, from a technology strategy standpoint, what we endeavour to do is make sure that we are precise with regard to model utilisation so that the value matches the cost of the model, the improvement of the model matches the value and the expectations of the client, and so that's what we try to do to enable the businesses.

Jason Napier, UBS: Tanuj, with Fit for Growth, restructuring charges and growth saves, we've slightly lost track of what organic cost inflation in the bank is. If you could give us a sense as to what the underlying rate of inflation costs is. Then Noelle, within that context, I think we all agree that the AI companies of today are not making any money and are spending a lot of it. What proportion of Group costs is IT, broadly defined, and is it a problem that we don't really know how they'll be charging in a year or two from now?

Company Representative: Okay. So, I think on the first one I'll probably turn actually to Manus to help come in on that, and then on the second to Tanuj.

Manus Costello: Thanks, Jason. I mean, obviously what we're doing is, Tanuj and I and the rest of the team, we're working very closely together on managing the costs of the bank going forwards. We've given you an indication in the past of what the inflation and growth rates of the bank are, and you've seen those walks previously. I think you should assume that that's a natural run rate that we'll be at, but you should also assume that we have flexibility ourselves, both in terms of what we see in the environment and in terms of how we want to invest and the pace of investment that we go at.

So, you've seen- we've called it out in historic cost walks in the past, and I'm not calling out anything differently now, but you should place that in the context of what we think we can achieve in terms of the top line and what you think we can achieve in terms of improving on this efficiency as well.

Tanuj Kapilashrami: So, managing the cost of tech, and I'd like Noelle to come in on this as well. I mean, one of the big things that we are doing is being very, very thoughtful on the cost of AI. So, the big one was Copilot, which we very recently rolled out to a large percentage of our workforce. I'm going to be very honest to say we took our time doing it quite intentionally. We took our time to roll out in a way which was a much more structured way. So, deep analysis across job families. What is the value creation framework for a job family that we roll it out in, et cetera.

So, yes, we are incurring the cost but developing a very clear value creation framework to ensure that we have the right return on investment that we are doing, with our colleagues, from an augmentation perspective. In terms of buying tech with the overlay of AI, I know that's something you've been thinking a lot about, Noelle.

Noelle Eder: Yes. I think Bill referenced the points earlier today that I would give you. First, I think it may appear to people that we might be a little slower than others with regard to AI. We don't publish numbers on massive amounts of use cases and all these kinds of things that we see in the press. What we spent our time doing was investing in a platform, because we fundamentally believe that the platform gives us scale, it gives us reusable capabilities, and it helps us sharpen our value creation framework, whether that is with regard to revenue and growth and client experience, or whether that is operational efficiency.

Our AI book is about 50-50 across both, but the marginal cost of reusing these capabilities is much lower than a significant new build every time we have an idea. So, then one more thing I would add, which is, I mentioned in my remarks that we have moved from large, infrequent releases to small, frequent changes daily. So, in my experience, and in my opinion, that's how AI works best, because oftentimes in generative AI in particular it doesn't come out of the box behaving itself like one might hope.

So, what you're doing is constantly tuning your hypothesis around the problem statement, and you're tuning the model to get it to work. So, if you go big in that frame, you're going to spend a lot of money trying to get to an outcome that you could have proven at much lower cost. So, we try to have that discipline, and the fact that we took our time on this platform, I think, gave us the opportunity to really establish those frameworks.

We have two co-heads of AI. We have two folks leading it on the tech side. They work together on this framework to help our businesses and our functional groups really use this effectively.

Katherine Lei, J.P. Morgan: I have two questions. One is that STAN operates across many jurisdictions, and I think different governments may have different guardrails when it comes to AI implementations. So, how do we reconcile that? So, say, for example, one example in China, I think they have very specific requirement on data storage and how to use data. So, something that you build outside of China may be very challenging to implement it onshore.

I'm not so sure about the other jurisdictions. I think they may have their own thing. So, how do you reconcile that and make sure that, within this connected bank, this will still work for STAN?

I think for the second question, it's not easy. Let me think about how to put it. It's about layoffs. It's about staff management, right? Now we have AI, I think just now you mentioned a very good point, is that AI handles the process where the human handles the trust, the relationship, the growth, right? I think that's a very ideal picture, but in reality, I think some of our existing staff may not be the best fit in this type of model.

So, how do we handle that relationship? Some of the tech companies in the US are making very, I'll say, chunky cuts, like 20% of the workforce and all those. As a bank, how do we think about human resources and staff management on that front? Thank you.

Noelle Eder: So, it's a good question, and we are in 54 different markets so there are a lot of rules. The majority of rules are around data sovereignty and data storage. That's the majority of the rules. The AI frameworks, from a regulatory standpoint, are coming forward, but I don't think they're substantiated enough for me to comment on them really at the moment.

So, let me tell you about our architectural approach to AI, because we saw this coming. We've had a lot of experience over a lot of years in all of these markets. So, we use two concepts that are really quite important to this thesis. The first is orchestration. What that means is essentially where data is rule-based at rest, meaning it needs to be kept inside of a particular country and come to rest there and stay there and be encrypted, we honour that. What orchestration does is, it will pull that data across the network if a client who owns that data wants access to that data. So, we manage the complexity in the network itself, and so the network really becomes a big part of the strength.

The second concept I mentioned earlier is portability. Portability is a concept that says, essentially, if something happens, I can move from one market to another market. I'll give you an example. It's not in AI, but it is in data, and Bill referenced it this morning. When Amazon had three availability zones taken out in the Middle East - and we do a lot of business in the Middle East, we have data stored in the Middle East, we have regulatory considerations about that data in the Middle East - what we did is partnered with our regulators and our risk teams and our businesses to move that data to the geo-resilient data centres that I spoke about earlier in my remarks, and we did it in a matter of hours.

So, that kind of concept is the portability concept, and so we are employing that concept. Again, we're not perfect, but we are employing that concept as much as we can to be able to facilitate on behalf of our clients across multiple jurisdictions, which is really one of the huge value propositions of the bank for them.

Tanuj Kapilashrami: Let's go to people. I mean, before that I'd say we had a responsible AI council in our bank before generative AI became a thing. So, even before generative AI came up, we were talking about exactly the kind of questions you're asking today. How do you compete with AI-native companies when you have multi-geography, multi-regulatory frameworks? So, this is something we obsess about a lot, as you would expect us to.

On people, I spoke in my presentation about becoming a much more skills-based organisation, and we've been on this journey now six, seven years. It's really not a tagline, but it is this realisation that with AI coming in, the construct of jobs that we understand today is going to become irrelevant. So, AI is going to impact every job. All of our jobs. We believe there'll be very few jobs that will fully go away, but we believe there'll be loads of activities in the jobs that we all do that will go away.

So, I think that mindset of what does it mean to become a company where people's work is defined by the jobs they occupy, to a company where you underpin all of the work by a consistent set of skills, has been a journey we've been on for the last five years. We've made a huge amount of progress on it, and what that has led is a big focus in our company on re-skilling and re-deployment. I looked at the numbers just last week. If I look at last year, over 50% of the new jobs that have been created in the bank have been filled by people internally by re-skilling and re-deployment. That number was less than 30% even 18 months ago.

So, this idea that these are skills needed for the future, we are going to up-skill our colleagues and deploy them in those roles is going to go away. We have been very clear in our presentation that structurally our back office or our corporate functions will be 15% lesser, and that means people's jobs are going to be impacted. We are going to give our colleagues all the support for them to be able to re-skill themselves for opportunities within our bank or opportunities outside. So, that's going to be the philosophy within which we will operate.

Company Representative: Thank you. I know there's a few more questions, but I think we are going to have to stop there. Noelle and Tanuj will be around at the next break if you want to ask them any other questions. So, thank you, Noelle and Tanuj.