

This transcript is provided for the convenience of investors only, for a full recording please see the Goldman Sachs Communacopia + Technology conference [webcast](#).

Thomas Kurian, CEO of Google Cloud at the Goldman Sachs Communacopia + Technology Conference on September 7, 2023

Eric Sheridan (Goldman Sachs): Okay. I know we're moving from session to session, so appreciate it if everyone can find their seats. And thanks so much for being here.

For those that don't know me, my name is Eric Sheridan. I'm Goldman Sachs' U.S. Internet Analyst. and it's my pleasure to have Thomas Kurian, the Google Cloud CEO, for our next session.

Thomas Kurian joined Google Cloud as CEO in November 2018. He has deep experience in enterprise software. Prior to Google Cloud, he spent 22 years at Oracle, where most recently he was the President of Product Development.

And before Thomas comes up on stage, I will read a Safe Harbor Statement. Some of the statements that Mr. Kurian may make today can be considered forward-looking. These statements involve a number of risks and uncertainties that can cause actual results to differ materially. Any forward-looking statements that Mr. Kurian makes are based on assumptions as of today, and Alphabet undertakes no obligation to update them. Please refer to Alphabet's Form 10-K for discussion of the risk factors that may affect its results.

Please join me in thanking Thomas Kurian for a presentation and fireside chat today.

Thomas Kurian, SVP Google Cloud: Thank you, Eric.

Google Cloud is the Enterprise Software Division of Google. We take the technology that Google's building and make it available to enterprises and governments to help them with their digital transformation.

At the foundational level, we build on a global infrastructure. We have 39 regions around the world live, many more under construction, connected over a broad network.

On top of it, we build a foundational set of AI models. There are many generations of foundational AI models that we have, and we expose it to clients through five product lines: Infrastructure to help them modernize their core systems; developer platforms to help them build new cloud-native applications; data analytics to help them understand data, analyze it and manage it more efficiently; cybersecurity tools to keep their systems, users and data safe; and collaboration to help them communicate and collaborate better.

I want to talk through a little bit why customers choose our cloud, and how, with the creation of Generative AI, we're seeing increased demand for various pieces of our portfolio.

So why do we lead in infrastructure, and why do customers choose our cloud for infrastructure? When you are modernizing IT systems, you're looking for two primary things: Performance at

scale, which give you better cost efficiency because you can run your applications more cost efficiently if you have a better performing and scalable system. And then from a risk point of view, you're worried about security and reliability, meaning uptime and security risks.

We're highly differentiated in a variety of areas. Just two examples: if you look at raw disk, we deliver ten times better storage performance when measured on an IOPS-per-disk volume than competitors, which means you get much faster and cheaper access to disk.

From an uptime point of view, we're materially differentiated. We have two to three times fewer outage hours, which means for mission-critical systems, a lot better experience for our customers.

As Generative AI came along, we've evolved our infrastructure portfolio. Remember, Google has been building large-scale infrastructure for AI for over ten years, so we have over a ten-year head start on any competitor.

As AI matures, we see a range of needs on infrastructure. First, we offer very high-performance systems for training models. Model training happens on large-scale infrastructure that we have. Over 50% of all funded AI companies are our customers on Google Cloud. And over 70% of AI unicorns use Google Cloud. In other words, those that know how to do AI understand our differentiators.

When you shift from model training to serving, the needs are quite different. From inferencing, you need a range of different types of accelerators to accelerate performance. You may have dense models, sparse models. You may be using a technique called vectorization. You may be doing embeddings to expose data from your databases in.

We offer 13 different types of accelerators for clients. The value is that they get the best choice of system-for-price performance. We also have the most diversified supply, because these accelerators give us a choice of what to offer clients.

In addition to having a range of accelerators, we have made many innovations as Google over ten years around the system design. So just saying "I have an accelerator" doesn't mean that it runs well. There's many elements of system design - from optical switching, how we handle dense memory configurations, what we do with network design and even water cooling, which delivers over 30% improvements in system throughput.

And, finally, as Google, we've invented many pieces of software with DeepMind and Google to parallelize training and inferencing across a large cluster of machines, which gives people way more efficiency on a per-unit basis when they're running.

So you see these results in our ability to offer two times better efficiency from training and serving models compared to the best systems of any competitor.

Developers love to build applications with Google Cloud. There's two reasons for it. First, developers love using open-source tools, and we have the deepest and broadest collection of open source tools.

Second, we also give them the ability to build applications using a set of technology that can run on any environment that they have. When we say "on any environment" - at their premise, on our cloud or on any other cloud. So, in other words, they can learn once, write once, deploy anywhere; and we make money no matter where they deploy.

An example of that is a recent product we introduced called AlloyDB. It's the fastest-performing relational database in the market. We run it in all four environments: Our cloud, on-premise and on other clouds. And it's the only relational database that can run in any of those configurations.

You see that in our adoption, both at the top end of the market where a system - for example, like Spanner, which is our large-scale database - scales 20 times better than the largest scalable system of any competitor. So for high-end, we work extremely well. And, also, because we made it so easy to use, startups and small businesses are growing very quickly in their adoption of our platform.

When we introduced our AI systems, we introduced a platform called Vertex AI. Vertex is used today by 50,000 companies, and it's grown 15 times quarter over quarter in its adoption because of the interest in Generative AI.

Vertex offers four primary capabilities. It offers the broadest collection of models. We've got over 100 models from Google, open source, third parties like Cohere, Anthropic and others all available through the platform. So people have a choice of the model they want.

Vertex also gives you all the services that you really need to use and build applications with AI. Things like grounding to ensure that you're [getting] fresh but also factual results from a model, watermarking, automation of reinforcement learning human feedback, synthetic data generation, ML pipelines. Google's been the longest in the market in using AI in products, and the tools that we use internally were made available to developers through Vertex to use for their own needs.

We also recognize that to use AI, you don't just want models; you need other capabilities like search, conversational AI, and all those are made available through Vertex.

And, finally, we add various kinds of controls in the system so people don't have to worry about the privacy of their data, intellectual property protection, areas around, for example, responsibility controls so that you don't have to worry about the tone of the models.

Analysis. Many people are using AI for the purpose of analyzing data. We start from a position of real strength, because we've always said you want to keep as few places for your data to be analyzed. And for many years, we have consolidated all the different things that people want for analysis in one set of systems.

So we offer a single system to analyze structured and unstructured data. A single system that is a data warehouse and a data lake. A single system that supports the ability to analyze using SQL, Python, Spark, and Generative AI tools. And you can analyze data from any SaaS application, on-premise environments and any other cloud without needing to move the data to

Google Cloud.

That's led us to manage over 40 times as much data as any other data cloud provider, and we have five times more customers than other data clouds.

When Generative AI came along, we've always said that we wanted to provide a digital assistant -- think of it as a digital expert -- to help people with their tasks for analysis. So we created something called Duet AI, which is a realtime AI system to help you do all your tasks you need for analysis.

What does Duet AI help you do? You can ask in natural language for the system to help you clean and prepare your data for analysis. You can ask it a question: "Tell me how revenue's trending." It not only tells you how it's trending, it also helps you analyze why it's trending.

It can generate visualization for you to see it. It can even write a set of slides for you, put the charts in it and write the narration for you.

Now, two core things with that. We already run large-scale AI models in our data cloud. We're running over 300 million prediction operations a year in BigQuery, so we know how to run it super efficiently.

Additionally, because our AI platform is right in line with our data cloud, you don't need to copy data out of our environment to another cloud or to another AI system, which saves you cost in network egress and other things, and can be five times more cost efficient.

We want to help customers from a security point of view, and we do it in two ways. The first is if you run your systems on Google Cloud, we want to keep you secure.

From that point of view, the data speaks for itself. We have far fewer security issues than competitors. And so we start from a strong foundation.

We also provide solutions to help an organization secure their data and systems across their enterprise, on-premise, in other clouds and on Google.

If you think about it, there are four pieces that you need to solve for that purpose. To get the best threat intelligence on what is going on and what are the new threats emerging, the reason we acquired Mandiant was for that purpose. Mandiant has the best threat intelligence in the market. They bring real expertise from the front line handling threats.

We take that and put it into a platform for analysis. And what that platform allows you to do is What security threats are emerging? Which ones are going to affect me? How is an intruder likely to come in and attack me? What do I need to take as an action to protect myself? How can I validate that I have closed the door? And the combination of Mandiant and our security platform gives us the ability to offer customers that.

When AI came along, we added, similar to Duet AI for Analytics, a Duet AI for Cyber. What it allows you to do is to get an AI-powered security pro sitting right next to you.

It allows you to get a summary of all the threats that are emerging. It prioritizes the threats; often organizations are overwhelmed by the number of threats coming in, so this can help you prioritize the incidents.

It gives you what we call attack-path management, which allows you to see what's the best path that an intruder is likely to come in with, and then it allows you to automate and document your controls.

But in macro, if you think about it, cyber is a relatively simple problem: understanding threats and threat intelligence, searching for patterns that are occurring and the ability to lock down those patterns. We have expertise as Google in Search, and we have the world-leading threat intelligence from Mandiant. So that combination allows us to offer material differentiation.

You can see some of the analysis that's come from actual customers. Our security operations platform allows them to search four times faster to find issues than competitors. And our Duet AI tool allows them to handle security incidents by improving the productivity of security professionals by seven times.

Finally, our collaboration platform, Workspace. Workspace has 10 million paying companies as customers. People choose to use Workspace, because it is easy to deploy and manage. All you need is a browser or smartphone to access it. It's secure by design. It's easy for first-line workers. First-line workers are people in a retail store, people like nurses in a hospital, pilots in an airplane -- in an airline company, concierges in a hotel. They don't want to carry around a laptop in their backpack in order to do their work. They have native access, secure access from these environments, and it's super easy to operate.

If you compare from a security point of view since 2019, independent studies have shown we are far, far, far, far more secure.

We've also integrated AI inside Workspace since 2015. So this is our eighth year, and we run it at scale. Just to give you a sense, in email alone, we handle 45 billion operations a quarter. So we know how to make AI run inline extremely cost effectively.

When Generative AI came along, we thought about providing people through Duet AI for Workspace a helper: an author that can help you write; a digital graphics artist that can help you design images and slides and music to go with it; a project manager that can help you do analysis using spreadsheets; and a meeting assistant that record a meeting, transcribe it, summarize it, assign action items and catch you up if you're late for meetings.

Recognizing that customers want AI through solutions that they buy, we're also working with a broad ecosystem of partners. I want to touch on three types of partners we're working with.

Data providers. Last week at our conference, we announced partnerships with Bloomberg, Dun & Bradstreet, MSCI and others to provide cleansed, high-quality data for both customers who want to train their models with it, or they wanted to use it for grounding. Grounding is validating an answer that a model gives you.

We also work with SaaS companies - Workday, SAP, Box, UKG and many others - to take our AI models and embed it in their platforms, so that they can empower the business process or business-line application that they want.

And, third, we're working with a broad network of system integrators. Just the five largest system integrators in the world have committed to train 150,000 people on our Generative AI models, and they have trained over 100,000 since March alone.

So how do we make -- how do we monetize Generative AI? Generative AI gives us five key opportunities. The first is to win new customers. And we are rapidly gaining customers who may be using another cloud, maybe still on-premise. And because they're choosing AI as a platform, they're switching to us given the strengths we have in AI.

It also allows us to upsell products, and there are three flavors of them. Infrastructure products - for example, our machine learning systems - are sold similar to general-purpose compute on a per-virtual CPU per hour.

Our platform Vertex is sold with a platform fee and then there's fees for the different services: the models, grounding, et cetera.

We also have the opportunity to sell Duet to our installed base. Duet for Workspace is sold in a subscription. Duet for GCP, we are in preview. We will talk about pricing when we make it generally available. We have a very broad portfolio of products that we have an opportunity to upsell into our installed base.

Over and above that, we see two additional opportunities. One, many of the solutions being deployed are outside of IT: in marketing, for example, to do print advertising; in human resources, to automate the employee helpdesk; in sales and customer service response; in product design; in field service. So we have many more opportunities outside of IT to sell our technology, and we also see the rapid growth of more projects within the IT department.

All of our Generative AI products - whether it's infrastructure, Vertex, our models and Duet for Workspace - are generally available. Duet for Google Cloud Platform is in preview. And we said at our conference last week, we will make it available very shortly.

As a result of our product differentiation, our go-to-market reach, we are seeing rapid growth both in existing and new customers. On average, customers use over 14 products from Google Cloud, which tells you the depth of the relationships we have with them. Customers - we have over 100,000 partners in our partner program, giving you a sense of the scale of our ecosystem.

We have expanded our go-to-market organization by more than four times in the last four years in a very disciplined way. And we are seeing growth from new customers as well as in existing customers.

Finally, we are being very focused on cost discipline as we grow. And the four primary drivers of

cost discipline: How efficient is your engineering team in building products? How efficient are you in deploying capital, particularly machines and data centers? How efficient is your go-to-market organization and salesforce? And how differentiated are your products? You see all these results in our performance.

In 2019, we were a very small organization. Four years later, we're one of the five largest enterprise software companies in the world as a stand-alone business. We've also grown both top-line and operating income over that period.

So just in closing, the market for cloud is still in its very early stages. As we have proven, every single customer who chose Google could have chosen another cloud, because the other clouds have existed for longer than we have. So they chose us because we have differentiated products; a strong go-to-market organization and partners; we have strong customer momentum; and most importantly, we have a proven track record for growing both top-line and operating income. Thank you.

Eric Sheridan (Goldman Sachs): Great. Thank you, Thomas. As Thomas makes his way over, we're going to have a fireside chat for the next 12 minutes or so.

Thomas, thanks for being part of the conference. Thanks for all the information in the slides. And congratulations on a really interesting Next event last week.

Thomas Kurian, SVP Google Cloud: Thank you.

Eric Sheridan (Goldman Sachs): Maybe to kick us off, the public cloud space and Google Cloud, in particular, has been through a lot of evolution over the last couple years. What is your view of where we are in terms of cloud adoption, cloud usage and where Google Cloud sits competitively inside that world view?

Thomas Kurian, SVP Google Cloud: Cloud adoption is still in its early stage, if you look at all the industries, all the countries. And the need -- and many industries are being reshaped as software-powered. If you just look at the percentage of spend that's in public cloud today, it's relatively a small percentage, so there's a long way to go.

The second thing, we've always said, if you look at what we see with Generative AI, it's just an evolution of cloud computing. So cloud computing was always about simplifying technology to make it accessible to everybody. The first version was infrastructure. So that instead of having to buy data centers and stand up machines, you simply have an API or UI to go to, to get compute on demand.

The second step was managed services. Now, with Generative AI, you can go into a chat system and say, "I'm building a mobile app. I need four nines of availability. I need less than half a second of response time, or one millisecond response time. Please create a cluster for me that's Kubernetes-based, manage it for me." Think of how simple that is if you can ask it like that, a question in English or Spanish or whatever language you want and have it created. And the more we make things simpler, the wider the opportunity.

Eric Sheridan (Goldman Sachs): Understood. Okay.

In your slide deck, you had a chart of the revenue growth you've seen over the last four-plus years since you took over Google Cloud. Can you look backwards, first, and help us understand what some of the key growth drivers were looking at that ramp in revenue trajectory? And what are you most excited about as primary growth drivers looking forward that will allow Google Cloud to win new clients?

Thomas Kurian, SVP Google Cloud: I mean, the growth drivers have been relatively simple. You have to have great and differentiated products, because customers can always choose somebody else. Two, you have to have a great go-to-market organization. And, third, you have to have a big partner ecosystem.

When we started, we were so small, we didn't have a partner ecosystem. We have over 100,000 partners now. When we were started, we had a very small go-to-market organization. To scale a go-to-market organization to the size we are, with the discipline to grow top-line and operating income during that period, obviously took a lot of work. But we have a very strong foundation.

Looking forward, I think - you know, we've been working at Google on AI since 2004, so that's our 20th year. There's a lot we have learned in not just building models but building products that have been enabled by AI. And you're seeing that. Most people ask us how come you guys launched so many products at Next last week. It's because we've been working on it for a while, but we also have extraordinary expertise in doing it over the last ten years.

Eric Sheridan (Goldman Sachs): Understood.

And referencing Next, you did talk a lot about Generative AI last week. Can you lay out your worldview of how Generative AI tools will be adopted and utilized by customers? And how you, as an organization are sort of helping the deployment of those tools into your customer base?

Thomas Kurian, SVP Google Cloud: Broad brush, we see people using it for a couple of purposes. One is within their organizations, to become more efficient. For example, to give their software engineering team the models to help them code. Helping their marketing organizations to create advertising copy using our image model and using text, for example, generation to do print ads. Automating the employee helpdesk. Employee help desk is where employees go to ask questions about benefits and all these things.

Looking at the procurement system. For example, "Find me all my contracts that have no indemnification in warranties, so that I can make sure that I have the right contracts in place." There are hundreds of these. We have over 500 use cases that customers have already solved with our Generative AI platform. And we are making those available. And those are all around. So that's the first thing that people are doing, is how can I use it to become more efficient?

The second is how do I transform the way that I'm working with my customer base? And the speed of the projects are super quick. Orange, a telecommunications company, the national carrier in France, built a customer service agent in two weeks. It handles over 10,000 questions

a day. And so it gives you the sense of just transforming the customer interface, transforming the way that products are created, et cetera.

Eric Sheridan (Goldman Sachs): When you take a step back, how should investors think about the sizing of the opportunity for Google Cloud between a few companies needing to train foundation models versus the application layer and other AI tools and services that you referenced in your presentation?

Thomas Kurian, SVP Google Cloud: So the way we think about it at the infrastructure level, we offer a range of different kinds of accelerators. Think of it as another flavor of compute. It's much more differentiated because we have material differences in the way we design systems. And systems design is far more important than just the chip itself to get the right performance.

Because we have a range of accelerators, we can offer the broadest choice. We're not supply constrained because, for example, manufacturing process for one type of chip doesn't depend on the chip -- on wafer and substrate, which is bottlenecked. But it's priced on a compute-hour basis.

Then Vertex, which is our developer platform, you pay a platform fee and then you pay a fee for every operation you do. And an operation could be calling a model. An operation could be calling the grounding service to validate your result, et cetera.

And then when you go up one more layer, you are using Duet. Duet for Workspace is priced on a per-user, per-month basis, similar to the subscription.

Now the thing you should know is that because we have great efficiency in the way we run models, we have the ability to choose the most efficient model to answer a question.

So, for instance, if you are -- if you use Gmail, we have on the mobile phone a model that's very efficient because it runs on the phone. When you ask "help me write," which a number of you can try, it actually writes on the phone. It doesn't go to the cloud.

If you can't use -- if you can't satisfy the result locally, it will transparently go back. But that's an example of how we're making models run in the most cost-efficient way in the right place. And so every step on efficiency also helps overall for us in terms of thinking of the opportunity.

Eric Sheridan (Goldman Sachs): Okay. If we take the theme of Generative AI, what synergies do you see between Generative AI and other aspects of your portfolio: Data analytics, security, data? How do those work together from a synergistic standpoint?

Thomas Kurian, SVP Google Cloud: Two elements. We've always felt that AI in the end was going to be -- our goal was to build all the skills that humans have in an AI-powered system. And so when you look at Duet, it's a manifestation of that vision. We offer an analyst, a professional developer, a cybersecurity expert, a graphics designer, et cetera.

So the first thing, because we make it so much more efficient for people to do analysis, people are willing to pay for that productivity benefit.

Second, because we're making it much easier to use, many more people will use the system. For example, very few people outside of professional analysts feel comfortable going to a reporting - or dashboarding system. CEOs certainly don't go to that to get their numbers.

But if a CEO could type in "tell me how revenue's trending and give me the answer, not only on what's happening but why it's happening," the seat count in terms of the number of people using it will also grow, because you are making it available to many more people.

So in addition to capturing the productivity benefit, we also see it widening the aperture of users within these accounts.

Eric Sheridan (Goldman Sachs): Okay. When you were here a year ago, you were in the midst of the Mandiant acquisition. We spoke about that a year ago.

Can you talk about your focus on security and how you think that differentiates Google Cloud? And bring us up to speed on how Mandiant fits into that broader security strategy.

Thomas Kurian, SVP Google Cloud: As I said, two basic things: if you are a security provider, you got to keep your house clean first. Why would anybody trust you as a security provider if you don't know how to secure your own cloud? And whether that's in Workspace or Google Cloud Platform, the numbers speak for themselves.

Now, when we acquired Mandiant, we said the rate at which threats are growing and the sophistication of the threat actors - even prior to AI, but now being armed with AI - requires you to have three basic elements. One, the best threat intelligence which is why we acquired Mandiant; second, a super scalable search and analytics platform; and then the addition of AI to help you prioritize and analyze and remediate.

When we talk to customers, for example, they have invested in lots of cybersecurity tools. When you turn on alerting, one customer showed us they had 15,000 alerts. We put Duet on it. It was able to find the two that are the most likely threat vector. And so it helps them prioritize and act and, as a result, keep themselves more secure.

So that was the rationale behind Mandiant. We always knew that we were going to integrate it with our security operations platform and our AI capability, to offer a differentiated cybersecurity offering.

Eric Sheridan (Goldman Sachs): I know we're coming to the end of our time. But if I could just squeeze in one more question. I think a key topic for investors is custom silicon.

Can you talk about your strategy to develop custom silicon? How should we think about the journey you're on from where it is today to where it might go long term? And talk a little bit about the benefits of offering both TPUs and GPUs to customers.

Thomas Kurian, SVP Google Cloud: So we've been working on AI systems in production environments for over ten years and huge environments: Search, Gmail, et cetera. So we've

always known that to accelerate performance and lower cost, you will need a range of different kinds of accelerators. And it's not about TPU or GPU, but a variety of different kinds of accelerators. So we started that.

And there are many, many elements of system design that we have within our systems. For example, if you are running an application to do inferencing but access a vector, we have optical switching that dynamically translates that and can find the right amount of memory. The way we do floating point is materially different. And we also allow you to train on a TPU and serve on a GPU, for example.

So that gives us much more efficient systems overall, which means that we can be more efficient in serving workload. It gives us flexibility and supply. And it allows us to offer not just customers, but our own internal needs, with the most cost-efficient capital utilization.

Eric Sheridan (Goldman Sachs): Understood.

Well, Thomas, thanks so much for making yourself available to have another conversation this year.

Please join me in thanking Thomas Kurian for being part of the conference this year.