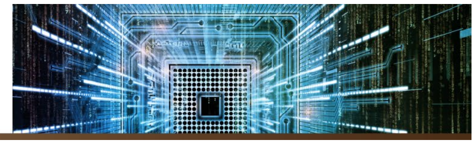

IonQ Inc. presentation delivered at the 2026 Global Technology, Media and Communications Conference on Monday, May 18, 2026 at 9:45 AM

Harlan Sur: Good morning, and welcome to J.P. Morgan's 54th Annual Technology, Media, and Communications Conference. My name is Harlan Sur. I'm the U.S. Semiconductor and Semiconductor Capital Equipment Analyst for the firm.

Also with me today is Mayur Ramdhani. He helps us cover our small to mid-cap semiconductor franchise. Very pleased to have Inder Singh, Chief Financial Officer, Chief Operating Officer of IonQ here with us today. Inder will kick us off with a brief overview of IonQ. It's been a pretty earnings season.

I've asked him to also give us just a brief overview of the March quarter, June quarter outlook, and then we can go ahead and kick off the Q and A. Inder, thank you for joining us today. Let me turn it over to you.

Inder Singh: Thank you, Harlan. Pleasure to be with you today, and Mayur, thanks as well. I've been with IonQ as a board member, originally when I was still CFO of Arm, and had sold Arm to Nvidia, and we were in the waiting period for approvals, which of course never came.

During that period, I joined the company's board as the first independent director, and then basically watched it grow from almost zero revenue to last year's $130 million and then joined late last year around September, once we had appointed a new CEO to help him execute on a platform strategy.

It's been a pleasure to watch the company go from being a lab experiment maybe in 2021 to being commercially deployed more and more. As Harlan noted, we had a strong year last year, $130 million, strong growth year on year. Guided this year for a top limit of $270 million in revenue. The company has a track record of beating and/or raising over the last five years.

We hope to keep that model going over time. We provided the guidance to indicate that we were going to again double year on year at the midpoint of the guidance this year, with strong organic growth, meaning our computing business, which was the primary driver for the first few years, continues to power the top line.

In the coming years, we're going to roll out the rest of our platform business, which includes sensing, it includes atomic clocks, it includes the ability to secure networks against quantum, etc. A portfolio story very much. We also indicated that we have strong RPOs, which are a measure of future revenue, $470 million.

Feeling pretty good about where we are, and looking to invest in the ecosystem around us to make sure that quantum can continue growing.

With my CFO hat on, it's obvious things. With a COO hat on, it's making sure we have the manufacturing, the supply chain, the IT, the procurement, the supply chain security as well, all of the things that are needed to scale the company.

Those are the opening comments. I know you have a number of questions you may want to try to get into here. Happy to try to address as many of those.

Harlan: I appreciate you participating today. The team has advocated that the trapped-ion systems, your architecture of choice, offer superior fidelity, connectivity versus other alternatives. There are three or four other alternatives out there, as most of you probably know.

As the industry moves towards what we call more fault-tolerant quantum computing platforms, how confident is the team that trapped ion will remain competitive from a scaling, manufacturability, and overall simplicity relative to the complexity, simplicity of the platform itself?

Inder: That's a terrific question. As we looked at the company about two years ago, the company was still using lasers to control these ion traps.

Ion traps have a natural advantage over other modalities. They begin with less errors. They begin with higher fidelity. They begin with higher coherence. These are all some of the drivers of compute power over time, and also, they enable computing at scale if you can get the right number of logical qubits.

About two years ago, we learned that using lasers is great up to a point. You can scale up to a point. You can scale up to maybe 100, 200 qubits before the machine becomes too big, too expensive, too bulky, requires too much maintenance, too much downtime.

That learning curve that we went through -- and we now have our fifth-generation machine rolling out -- is what was behind our acquisition of Oxford Ionics, which puts us onto a semiconductor roadmap.

A much more proven modality that exists today, 30 years of history, 30 years of scaling, 30 years of knowing how to scale something from a few bits to a billion bits and more. AMD, NVIDIA, obviously, all of them leverage that. Now, we are also.

Going forward, our intent is to use a semiconductor roadmap, beginning with our 256-qubit machine, which we are developing already, and our 10,000-qubit machine, which we've begun to turn our attention to.

As we look at the ability to deliver what we call time to solution, which is really what matters, how quickly can you get to useful answers.

We've proven through a paper that we've published on our website that you can all see, that for many practical things, ion traps offer you the best path to that and if you have enough cubits, logical cubits, you can do some pretty amazing things.

We sit at a crossroads here where we are now moving from laser-based systems and tempo, which we are in the market with right now, is our last laser-based system, it's 100 cubits.

The 256, we've started development on already. On the last quarter, we announced we had successfully gone through developing a prototype and now we're building the system around it. We have a clearer path to being able to scale to 10K, 20K, and even a million K by 2030.

Harlan: I do think it is a big differentiator for the IonQ team, the optical/laser-based control approach to your innovative electronics/RF control-based opportunity. We'll get into that a little bit later.

Going back to my earlier comment, which is that, there's multi modalities that exist now for enabling quantum computing architectures. Do you expect a future where there will be multiple modalities that exist, superconducting, trapped ion, photonic-based quantum computing for different workloads?

Or do you think that the industry consolidates to one or two modalities given the compute related applications that are required?

Inder: Terrific questions. I think there are a number of modalities in the market already. There are a number of companies that are either using ion trap as you mentioned, superconducting is another very popular one. There's photonics, there's neutral atom and so on. Multiple modalities.

Usually, what you find is over time in industry, and I've lived through a few industries in technology, you end up consolidating around a few, so I expect there will be a few. I don't think there is one.

Ion trap definitely starts with an advantage is there. I think superconducting with names like IBM and others behind it probably is also there. As for the others, they have more development work to do. These two are furthest along.

Ion trap for sure, because we've been investing in it for five years now, and then superconducting as I say as well. We are preparing for a multimodal world in the future. Five years from now, absolutely, there should be a heterogeneous environment.

Our networking capability, which we are unique in being able to connect quantum with quantum, connects an ion trap to another modality. Our ability to secure against quantum is also agnostic of the platform itself. We're preparing for that. We think ion trap will play a key role. The other modalities may play a different role.

Harlan: Going back to the differentiation that you guys have brought into the portfolio with the acquisition of Oxford Ionics and this whole notion of how do we control the qubits? Trapped ion historically, you've used light/laser-based techniques to control the qubits.

You guys with Oxford Ionics have brought this very elegant, very simplified electronic means to control the qubits where, like you said, you're leveraging classical semiconductor based, existing, very, very mature semiconductor technologies.

As you move from your current platform, which is tempo, which is still optical based, laser based to your next generation 256 physical qubit platform where you will be integrating your new electronic control architecture, what are the key milestones that we should look for between now and expected 2020 timeframe for your next generation solution?

What are some of the milestones? What are some of the key KPIs that we should be? What are you looking for in terms of bringing this solution to the market? Gate fidelity, reliability, uptime, calibration, packaging yield, manufacturability, what are some of these metrics?

Inder: Fantastic question. We've begun to reveal some of those milestones already. The last two earnings calls, we've talked about the fact that even while we're putting the Tempo system into the market and that's going to drive the majority of our revenue in computing this year, we've already developed the 256K chip prototype, already developed.

It's gone through tapeout A, B and C. D is now complete as well. It's got feature-rich ability at this point. We are now surrounding it with the rest of the system. The chip is about the size of your thumbnail. The machine itself is much bigger, of course. All of it has to come together and work together.

There's a compiler that has to be part of it. There's other electronics that have to be part of it. All of that is now being put together into multiple prototypes for the rest of this year. Each quarter, we'll be telling you how we're doing.

Last two quarters, we've been ahead of schedule. What we thought we would take nine months to do has been done in just a few. We're working with a fab here in the United States that allows us to accelerate our roadmap, unlike some of the fabs we were looking at overseas previously.

Not only is it secure from a US government standpoint in terms of supply chain -- the government feels comfortable with it, potentially being a customer in the future -- but also, it gives us the ability to do parallel prototyping, not just one at a time.

To your point, the ability to get from 100-qubit laser based system to a 256-qubit electronic control system is already there. It's already on a chip. To go from 256 to 10K is the next milestone that we'll be talking more and more about.

The ability to have electronic control means fewer lasers. Fewer lasers means lower cost. Fewer laser means less complexity, less downtime, less bill of materials cost. As the machine becomes more powerful, it becomes simpler and potentially cheaper.

Over time, our strategy with our five-year roadmap that we've also laid out is to make our machines modular so that after the 10,000-qubit machine, which we've already started work on, you start to do modular upgrades. You don't have to replace the machine anymore. You do swap-outs of a few modules within the machine. The system stays intact.

You get customer stickiness that way. Our customers benefit from lower total cost of ownership.

Our machines don't require being operated at zero degrees Kelvin or close to that. We don't have to have dilution refrigerators, helium access, etc. There are some advantages.

The cost of buying the machine and then operating the machine, which is really what a customer looks at, TCO, much, much lower. Then modular upgrade strategy makes it stickier with the customer because we forward-deploy engineers and app developers to make our machine become part of the customer's revenue stream, not just their cost equation.

Harlan: Before we go into some of the forward roadmaps, you have already gone through one transition. You're focusing on your fifth-generation Tempo platform, moving to...We've been talking about your next-generation 256-qubit platform.

Going from fifth-generation to sixth-generation, just help us understand the scale of, from your customers' perspective, the scale of applications and complexities that you've been able to unlock for your customers in making that move from fifth-generation to sixth-generation?

Inder: Terrific question. In terms of classical computing, when we think about increasing the number of bits in a processor, for example, the amount of Level 1 cache, things like that, there are a few things that actually drive up computing power, but it's still Moore's Law. It's basically doubling over time and lowering cost over time.

With quantum, it's exponential growth in computing power. It's 2^N, not 2*N. When you go from 100 to 256, you basically have exponential increases in computing power. You just have to make sure that you are able to do that in a manufacturable way, sustainable way, which we're doing, but also you have to make sure there are algorithms and applications ready to take advantage of it.

We've already done a lot of proof concepts with our older-generation machines, things around life sciences, like protein folding; things around drug discovery, accelerating drug discovery, for example, in partnership with NVIDIA, a number of things that you can do with fewer qubits.

With 256, you unlock so much more. You can just think, with 10,000 qubits, it's a leap up in the ability to get to fault tolerance.

Harlan: Beyond your 256 sixth-generation system, on the last earnings call, you guys said you guys are already pre-selling some of these platforms right now. Your next shift would be...The team is already starting to focus its sights executing on your next-next-generation platform, is your 10,000-qubit solution. That's a pretty significant jump.

What underpins the team's convictions you can deliver that step change? Which are the leading indicators, KPIs, that would signal to us that the team continues to be on track to execute that?

Inder: Absolutely. As we have turned our eye to the 10K, we've not taken our eye off 256, of course. As we've turned our eye to that, it's around multiplexing, something that the semiconductor industry knows how to do very well. It's leveraging a CMOS environment to actually scale from 256 up to 10K. Very proven path over decades. We're following that path.

To your point earlier, which was really important, we're using mature nodes. We don't have to be three-nanometer or two-nanometer, ever. We're talking about things that are 128, maybe going to half that as we go to shrinking die size over time, never having to need those advanced nodes, which means fully depreciated plant and therefore lower cost for us over time as well.

The milestones will be very similar. It'll be the ability to demonstrate that we can have a prototype of a 10K -- that'll be step one -- to have multiple iterations of that to work out the yield over time, of course, and then to build a system around that, just as we're doing with five generations of experience of building systems with 256.

Harlan: The way that you described it is, 256 to 10,000, that's classical leveraging semiconductor expertise. As we all know covering semiconductor companies, the circuit complexity will become more. It's probably still monolithic chip-focused.

If we think about your roadmap now to two million physical qubits in that 2030 timeframe, you guys have articulated a number of different potential strategies. It could be multi-chip. It could be still very much monolithic chip-focused.

Help us understand, for the two million physical qubits 2030 timeframe, how much of what remains is still fundamental science/innovation versus engineering and scaling work that's largely derisked that can take advantage of either your semiconductor expertise or your optical networking expertise? What are the additional technologies or breakthroughs that are still required to achieve that and so on?

Inder: A lot of the science breakthroughs that were needed are behind us at this point. Now it's about engineering and manufacturability.

The last remaining milestone in the science breakthroughs was getting to four 9s, and we did that. 99.99 percent fidelity, which means the lowest possible error rate, which means basically it's on par with classical computing at that point. We have to maintain something close to that as we scale.

The engineering part of it is around a triplet strategy, moving beyond multiplexing, scaling up, and at some point, maybe even going beyond two million physical qubits, we would use interconnects and things like that if we needed.

The team feels very confident about even getting to the two million. To be candid, at two million with 80,000 logical qubits, you can do some really, really impossible in very, very short periods of time.

We feel confident that we are executing the roadmap. What I like is that we're ahead of schedule. We're ahead of schedule on the 256 development, and now we're ahead of schedule and actually starting to think about the engineering design of the 10,000-gigabit system.

I'm not predicting anything yet. There's always things to do. Working with SkyWater, which is our foundry here in the US, we've had very good success in being able to demonstrate that you can make a chip-based ion trap system and scale it.

256 is far more than anyone else has been able to do so far. For sure, using electronic controls were unique. We think that over time, that will be a natural advantage in terms of cost coming either further down. Either we pass that to our customers or keep some of that ourselves, that's a decision still to be made.

Harlan: This sector is somewhat very highly technical. We hear terms often being used, highest fidelity, physical qubit, physical gates. We hear about things like error correction and so on.

The bottom line is the endgame is to build a compute system that is fully fault-tolerant. The team did put out a blueprint for that. You call that your Walking Cat Architecture. Spend a few minutes talking about fault tolerant quantum computing and what it means for IonQ and then it also appears that as a part of this fault tolerant roadmap.

You are potentially moving towards what we call a QCCD like architecture, shuttling ions into dedicated zones for computation. Maybe you can also talk a little bit about that as well.

Inder: Yeah. The ability to do any connections is also quite unique to ion trap. It's harder to do with the other modalities. If you have 10,000 or 20,000 or 200,000, being able to entangle ions that are not physically next to each other, and being able to do that in this Cat state that you talked about, which is part of our Walking Cat Architecture, is unique.

Because that operates in a way where you can not disturb the quantum entanglement, still be able to look for errors and be able to correct those errors. That's a fault tolerant machine.

We think when we are at the 10,000 and beyond, we can start thinking about fault tolerance, which is why we published this paper called the Walking Cat Architecture. It's a cute name. It's named after Schrödinger's cat. It's about 100 pages, so it's not a light read.

You can have an AI agent summarize it for you and make it easier to understand. Essentially, it involves modularity, it involves making sure that we can have manufacturability as we do this. It involves making sure that we have a compiler system and a micro architecture, all that come together. It's all published.

Not only are we talking about the ion trap, not only are we talking about the number of cubits going to 10,000, we're talking about a fault-tolerant machine, which essentially, if you think what that is, it's a self-healing machine. If it finds an error, it corrects the error itself without intervention, and that's what you need for industrial scale.

Harlan: Just pivoting to business strategy, you're pursuing computing, networking and sensing simultaneously, often with different technologies.

How are you integrating these into a unified platform from a hardware, software and go-to market perspective? Can you share perhaps a few concrete examples of applications where these capabilities work together?

Inder: Great question. Again, the company is unique in terms of the platform it's put together, which includes the ability to network machines together. Going back to my days at Cisco Systems and learning that you need to be platform agnostic, connect everything to everything, our approach to networking is exactly that.

We've demonstrated from a technology standpoint and now from a deployment standpoint in a number of countries we've announced that we can deploy the network irrespective of the compute platform, and even if you don't have a compute platform.

We've also demonstrated the ability to secure against what quantum computers will be able to do one day, which everyone calls Q-Day. The ability to break encryption, which is you probably have been reading and I've been obviously tracking, is getting closer and closer and closer.

Even a year ago, people were saying it's 20 years away. Now Google and others are saying it's a few years away, so it's a question of whether it's a few years or less. We're preparing for our customers to be quantum secure and have quantum computing at the same time.

We're also preparing for the ability to provide networks and sensing P&T networks that are jam proof in an environment where you have GPS being spoofed and jammed every day, as we've seen.

That platform or that suite of products that we bring, some people start with one thing and go to another. Some people start with two things. Some people start with more than one thing.

We have the ability under one roof now to have a customer start their journey by buying the network first and then the computer or vice versa, or in the case of a customer, QuantumBasel, buy a computer and the next generation and the next generation and the next generation.

That's a huge lock in for us over time, and it gives us visibility through the RPOs that we've talked about, to be able to serve those customers over time. Having $3 billion of cash available also helps as well in terms of our ability to invest for the long term.

Harlan: I want to make sure that we address any questions in the audience. If you do have a question, feel free to raise your hand. Please wait for the mic to come to you if you have any questions.

We've got a question right here in the middle here.

Audience Member: Thank you. Putting your COO hat on, in the topic of supply chain vulnerability and quantum technologies as a sovereign technology, how are you building your roadmap for a shifting regulatory landscape?

Inder: Terrific question. Thank you for that.

The fact that we are developing a roadmap that I think is without parallel, with all due humility, the ability to have 10k and have 20k in the time frame we're talking about, means that you will create machines that can do amazingly great things, things that classical just can't do, and amazingly bad things potentially at the same time as well.

It's important for us to have, therefore, to your point, not only a secure supply chain in terms of availability, right, but also in terms of provenance of the components, the manufacturing being secure itself. We were asked by certain national security customers to have that in place before they start to even think about deploying some of the things that we have for those types of applications.

We were looking, just like every other company in quantum is, which foundry do you use, how do you scale, can they move fast enough, is it secure? I'm not going to name any particular ones. You can probably figure those out yourself.

Most of the foundries that are out there are for semiconductors. They don't have experience with quantum. What I was finding with my CEO hat on as I negotiated with some of these foundries was, they were amazed by how much volume we were predicting we would need and how quickly we would need it.

They were struggling with their own parent company to be able to justify funding that part of for quantum alone, so they are asking for things like revenue share and stuff, and I said, "Over my dead body."

We started looking at a US foundry at that time as an alternative. It turned out to be SkyWater. SkyWater brings with it the highest level of military security for many of the applications that they already do for the government.

We felt comfortable having them manufacture for us because we could look at provenance, we could look at making sure that the people that would work on our machines in terms of developing the chip itself, we would have clear line of sight, there's no embedded malware. As you know in semiconductors are things called secure enclaves. They're not always secure.

Those are the things that we can now focus on. Surety of supply and security of supply for our compute platform in particular is something that we took very seriously because we figured might as well do that now rather than having to do that later. It will be very hard to change foundries two years from now versus today.

We're starting our chip roadmap entirely in the Skywater foundry to your point. The need for sovereign ownership of machines is something also we're seeing. Every country that I've spoken with, that Niccolo, our CEO has spoken with or our sales team is looking for a machine to be owned by them, which is why we're selling more and more systems.

They're happy to get cloud access to learn, to understand how quantum works, to train people. But for hybrid workloads, which we're seeing more and more of, every country is saying next to my AI factory, next to my GPU cluster, I want a QPU and I want to be able to do hybrid computing. For that, I need access to the machine itself, not cloud access.

There are certain things you can do just fine on the cloud. There are many more things you can do if you own the machine. We've moved very much into providing those machines. To your question though, which is a good one, I come from semiconductors most recently in other areas, we recognize that when we have a 10k, 20 thousand and beyond, we may be not allowed to sell those machines to certain countries.

We're operating already as if we have export controls, even without them being in place today. I wouldn't want to promise something to a customer and say buy our 256 and not be able to sell them a 10k next or 20k next. Thank you for the question.

Harlan: Any other questions? Oh, I've got one up here.

Audience Member: Thank you. Could you give us a sense of how important software and algorithms are compared to multiyear?

Inder: Hugely important. Obviously, the question around algorithms and software, and that's not lost on us for sure. We have one of the largest application development teams, if not the largest in the world, that we've built and are building. We've identified about half a dozen areas, end markets like life sciences, material science, financial services that we will develop algorithms for ourselves.

Then others that we will do it through partnerships. Protein folding, things like that, will do with someone else perhaps. Drug discovery, we'll do for someone else perhaps. Whereas in material science and battery chemistry and things like that, we might develop that ourselves. You can't do everything, but we are investing in the ecosystem at the same time as we're investing in our products.

One of the stats I read just this morning actually was a research study that came out today and I was really surprised. In the world, it said there are only 5,000 quantum engineers. It sounds like a lot, but not really. There are lot of quantum physicists, not many quantum engineers. What do we need at this stage?

Both. Physicists, of course, but much more quantum engineers. We started investing in universities, certain universities, not everyone, where quantum engineering we think can be something really big, where they can have our machine, train engineers on our machine, you can see the benefit, and then graduate as quantum engineers.

And so 5,000 hopefully becomes 10,000 and much more. If you think of all the companies in this space and you think about even the Googles and the Microsofts all having quantum folks, they all need engineers at some point. Fortunately, they're not building machines, that's not their business model. IBM is of course.

We want to do something that helps the entire industry and us. One way to do that is to have people able to develop quantum algorithms. And so if you think about the iPhone and the App Store is the example that I try to use. If we're building more and more powerful iPhones, we're building the App Store that goes with it at the same time.

Some that we will have ourselves, some that we will curate that to work on our machine. Thanks for the question.

Harlan: Any other questions? We've got one right there.

Audience Member: Great insights. Thank you. Two questions. First of all, in terms of obviously, you've heard Quantinuum is going public. They're also in the trapped ion modality. Would love to understand from your perspectives where you think IonQ is going to be differentiating. And then you were talking about various modalities.

Google, a couple months ago, announced that they were working on, I wouldn't say necessarily abandoning their superconducting program, but moving towards neutral atoms as an alternative. Obviously, scalability was one part of that. Would love to get your comments and perspectives on those two points.

Inder: As I mentioned earlier, I think I would love for all these modalities to really have a market in the future, let's say five years out. I'd love for all of them to coexist. I think the reality is probably some of them have more science breakthroughs to do than others that will just take a little bit longer to get there. That's not a knock on any modality.

I'm an engineer, so every engineer thinks they're doing the best thing at the right time. All of them are as well. To your point, Quantinuum is also an ion trap company. In the US, I know of ourselves and them, and I wish them well. I do think that we need to have a number of really successful companies three to five years out for this to become an industry.

I'd love for them to do the same investments we're doing in the ecosystem. We're just ahead in terms of the fifth generation, the sixth generation, the seventh generation and selling at scale and manufacturing it scale. We'd love to see all of these actually take off. As I said earlier, I think two are already on the trajectory.

Superconducting, yes, Ion trap, yes. Neutral atom has certain advantages and I'll let those companies speak for themselves. Some of them are here, here today. Of course, photonics over long distances offers lots of promise, still has some science breakthroughs. Having light travel and being tangled over very long distances is a non trivial matter.

The networking that we have, the security that we have, the sensing that we have is meeting customer needs today, which is flowing through our revenue stream right now. As I said, we are investing in making sure there's actually an ecosystem. And that's how companies always need that ecosystem for success.

This is a nascent industry. We draw AI engineers from the trillion-dollar tech companies. They choose to come to work with us, because they think they're going to build the most cutting edge, leading edge applications that can't be simulated in a classical environment. I'm not saying QPUs will replace GPUs.

People would love for me to say that, I'm not saying that. I think it will be a hybrid world. I can't come from a CPU company arm. There are still more CPUs than GPUs, believe it or not, but they all coexist together. I think that every modality begins with some benefits and some disadvantages.

Ion traps begin with probably more advantages and our founder 30 years ago, whether lucky or smart, chose Ion trap, allowing us to be able to be where we are today. We are a merchant supplier. In fact, we sell components to the other quantum computing companies. They don't talk about it. We don't talk about it.

Some of the things that they require for their machines, not all of them, some of them, their machines wouldn't work with our components. We want all of them to see it, actually. The competition to me is not any of them, candidly. The competition is probably a sovereign nation on the other side of the planet, maybe a few of them, trying to get to the same Q-day that this country is racing to as well.

Harlan: Great. Well, we are just about out of time. Inder, thank you for your participation today. I look forward to monitoring the progress of the team as the year unfolds. Thank you very much.

Inder: Thanks for having us.

Harlan: Pleasure.
