Nelson Gonzalez: Hello, my name is Nelson Gonzalez. I'm the head of Global Impact Computing at Amazon Web Services. Delighted to join you for the Six Five summit, and I'm particularly delighted to see that there is an ESG and DEI track here at the summit. That is the job of my team to mainstream social and environmental impact into our advanced compute motion, so very glad to see that I am among close colleagues who are doing the same.

At AWS, we believe that the next decade of business change will be defined essentially by the sustainability transition that we're all in the midst of. Companies that are leveraging this transition, and we call these companies, "Twin Transformers," are, we found, over two and a half times more likely to be among tomorrow's strongest performing businesses. And with these twin engines of growth, sustainability, and innovation, this outperformance of their peers shows us that sustainability is good for people, for the planet, and also for business. So often, we think about sustainability as being something that has to do with philanthropy or CSR. Right now, we know that it's bottom line, that it's a driver for innovation and a driver for growth.

The enterprise sustainability conversation has generally centered on the very important task of achieving net-zero targets, by reducing carbon footprints across our workloads. This is critical, however, sustainability challenges resulting from climate change and resource use are posing opportunities for us to actually become much more intentional about both mitigation and adaptation. So we are committed to both developing an increasingly sustainable compute infrastructure, but also using that infrastructure to create computing solutions that address sustainability challenges, that our customers and partners for whom social and environmental sustainability is their core business, is their bottom line, we want to empower them as well.

So let's begin by thinking about how we are committing to making cloud compute itself more sustainable. The cloud itself, as you may know, is by itself a great step in moving toward a more sustainable net-zero strategy. When compared to data centers across the world, AWS can lower a customer's carbon footprint by nearly 80% today, and once we achieve our targets that I'll talk about in a second, that number will jump to around 96%, and we're on the path to being 100% renewable by 2025, which will allow all of us to achieve those targets. So the cloud itself is a move towards sustainability.

At AWS, we are committing to do as much as we can on our part to ensure that we can pass on our carbon efficiencies to you, our customers and partners. We recently announced that our commitment to return more water to the communities in which we work than we use, in direct operations by 2030, and we're going to meet that goal by further improving water efficiency, increasing our use of sustainable water sources, like recycled water and reusing cooling water, and supporting projects that increase water availability in communities around the world.

I'm really proud to also mention that our teams are innovating at the silicon level itself. So we're innovating at the infrastructure level in our data centers, innovating how we use resources like energy and water, but we're also building entirely new chips that are much more carbon efficient than their predecessors. AWS' global infrastructure's built on our own hardware, which includes purpose-built servers, routers, and processors that are optimized for our various workloads. But Graviton Three and Inferention are EC2-based instances that are up to 60% less energy consumptive than comparable other EC2 instances. So not only by moving to the cloud,

but by using some of our own silicon, you too can achieve these net-zero efficiencies. So we're very proud of the ways in which both our infrastructure and our compute technologies are actually moving us toward greater efficiency.

So this is great news, and in some ways, it's also table stakes, because we're committed to reducing our carbon footprint, and, by extension, yours, and we're also committed to driving innovation in our compute services, so that they can power the solutions to social and environmental sustainability challenges that together we're facing. And so, we think about sustainable computing, but also computing for sustainability.

The AWS Impact Computing team, which I'm very proud to lead, was created to enable compute solutions to these social and environmental challenges that we're facing, as we think about both adaptation and mitigation. Our team leverages AI/ML, including generative AI and very large scale machine learning, high performance computing, and even hyperscale computing, in some cases, as I'll talk about in a second. Quantum computing, as well as autonomous computing, and that includes digital twins in areas that I mention here, climate risk and resilience, food security, biodiversity and conservation, circular economy, health equity, and ESG analytics, moving beyond greenwashing to very rigorous and meaningful, both measurement and monitoring, of our obligations around ESG and beyond.

We begin with data, and therefore my team is also very proud to host the Amazon Sustainability Data Initiative, which focuses on reducing the costs, the time, and the technical barriers, to extracting knowledge and insights from very large data sets, the kinds of massive data sets that are used to predict weather or forecast climate risk, and increasingly get us to an early warning for all system, as the UN has recently called for.

Using AWS and its scalable infrastructure and associated computational services to address this problem, we're using ASDI to advance innovation in climate science and climate modeling around the world. Through the program, we also provide cloud grants to democratize access to data, and to compute for innovators and researchers and all geographies to use this data for their own climate mitigation strategies and actions. And I'll have a link at the end of my presentation so that you can apply for these grants, or get in touch with me if you'd like to collaborate on this data as well.

ASDI the currently works with scientific organizations, like NOAA, NASA, and the UK Met Office, as well as the Government of Queensland and Australia, to identify, host and deploy key data sets in the AWS cloud, including weather observations, weather forecasts, climate projection data, satellite imagery, hydrological data, air quality data, and ocean forecast data, and these data sets are all publicly available to anyone. And we've also begun to work with the Amazon Data Exchange to ensure that you can leverage this data on an open platform, but also make some of that data closed, or even monetize it, to facilitate business model innovation across the industry as well.

So let me talk a little bit about how we are leveraging advanced compute for some of these social and environmental sustainability use cases. And I'll begin with high performance

computing, because in a way, HPC and supercomputing of different kinds really powers our ability to do the magical things that we can do with AI/ML, quantum, et cetera. So with high performance computing, we're leveraging, as these data, to build tooling that democratizes access to advanced compute, by increasingly automating and facilitating easier ways to leverage what is essentially a supercomputer, and democratizing access to that kind of compute power.

We're pushing the bounds of supercomputing in the cloud, specifically currently for food insecurity early warning systems. We're really proud to be working with Descartes Labs, we have an R&D initiative at Harvard, we're working with organizations like Cropin and the World Food Program HungerMap, to develop early warning systems that predict food insecurity and then avert famines. We're incredibly proud to be taking away the undifferentiated heavy lifting of this massive scale compute, so that we can focus on the intersection of climate and crop yields in their political and economic context, so that the most vulnerable populations in the world, particularly in Africa, can have a chance of having the kind of food security that we should all enjoy as a basic human right.

With artificial intelligence and machine learning, and that includes very extreme edge applications of AI/ML and generative AI, we are working with the Natural History Museum, and some very interesting modeling and simulation partners, like Lucid Mines, as well as the Nature Conservancy, to use the most advanced forms of AI/ML for biodiversity monitoring and reporting, as well as increasingly to develop natural resource economies that allow populations around the globe, and including indigenous communities, to essentially securitize their natural assets for economic benefits, extending the logic of carbon credits, in a way.

With analytics, which are essentially what we can do with the data on ASI and some of our other platforms, we are creating new kinds of data lakes and analytic methods to respond to the complexities of, for example, the intersection of public health and climate change, or as I mentioned earlier, food security and climate change, including applying living ecosystem relevant data science to multi-variate analysis and very complex causal inferencing around circular economy intelligence and analytics, to assess and address the health impacts of climate change. We're delighted to be working with Silver Lining on an initiative, together with the Brookings Institution and the Rockefeller Foundation, which will present to the UN General Assembly this September, as well as working with the Met Office to innovate the data proximate compute that actually brings compute to the data, rather than forcing us to bring the data to the compute, therefore creating efficiencies and carbon reductions that are very important to the previous net-zero conversation that we had.

One of the most exciting areas, for me, recently, is our work with our simulation technology group. We are working on developing what we call, "Natural Digital Twins," and living system modeling for regenerative design and circular economy. We, in cloud computing, tend to be incredibly good at addressing the needs of linear, industrial, mechanistic systems, even very complex ones like jet engines or wind turbines or the development of new vehicles like the America's Cup ships or boats or new Formula One cars. Incredibly complex and sophisticated work, but not as sophisticated and complex as modeling, for example, the behavior of populations vis-a-vis climate change, how population health and epidemiology needs to change

as we think about the intersection of, for example, redlining and asthma, in the context of pollution.

These living systems, and more importantly, the conversation among these living systems, create for a very complex data science analytics, and indeed, modeling and simulation challenge, and we see that as an opportunity. And with our simtech group, and partners like [inaudible] and modeling and simulation to take account of these dynamic living systems. We're very excited to be working toward developing what I'm calling not just, "Human-in-the-loop," but, "Life-in-the-loop" approaches to agent-based modeling and simulations of large complex adaptive systems, and enabling agent-based modeling at scale for applications like climate risk analytics.

And finally, with quantum, the absolute edge of what we're developing at AWS, in fact, we have our own quantum computer that we're developing in partnership with the University of California, and are doing some amazing work with Harvard on quantum networking, when quantum machines begin to talk one another. Absolutely changing the game, even at the edges of what high performance computing can do. And we're beginning to apply quantum for social and environmental impact, initially around virtual screenings through hyperscale compute, toward quantum chemistry, to, on one hand, remediate PFAS. These are forever chemicals that you may know are highly toxic and are in just about everything that we wear, touch, walk on, and eat, from and with. And so, we are beginning to work with big chemistry to successfully, actually, begin to change the molecular dynamics of some of these plastics, so that we can actually return them safely to the environment, and for the first time ever, create circularity around what have been indestructible and incredibly toxic and dangerous chemicals.

What's interesting is we're using a very similar set of processes, virtual screening, with the Dana-Farber Cancer Center and Harvard University, to think about how we can use virtual screening to begin to democratize access to very large scale drug discovery methodologies. In the case of the Dana-Farber work, we actually beat, as far as we know, every cloud record in terms of cloud computing. We hosted a 5.8 million VCPU run last year, that enabled Dana-Farber to analyze over 80 billion molecules, which is exponentially more than they've ever been able to do. Not only is this good for curing the hardest to treat cancers, but it's great for enabling Global South scientists to address rare diseases in ways that give them the power that, to date, has been in the hands of very few. So we're very excited about what we're doing with quantum around virtual screening and beyond.

So with that, let me close by saying that the opportunity that we have together is to actually create an entirely new field, and perhaps industry. Many of us know that about 20 years ago, the financial services were revolutionized by an idea called, "Impact Investing," when investors on Wall Street and the City of London and beyond began to realize that they were serious market rate investors, but they were using the market intentionally to drive social environmental impact. And now, we know that impact investing is absolutely mainstream in finance, and impact is becoming a requirement as particularly new generations force us into a more conscious approach to investment.

Well, we're on the brink of creating something akin to impact investing in impact computing, and that is to mainstream social, environmental and ethical considerations into our code, into the very stack through which we develop our solutions, and then applying those solutions to some of these, until recently, intractable problems, in a social and environmental perspective, that I mentioned, with our partners and customers, we're beginning to actually make much progress on.

And so, for us, we're beginning to think about what an impact computing stack might look like, and that is to begin to think about the ways in which our databases, our data lakes, our analytics, and our insight generation, our modern data architecture, could begin to be deployed in new ways to help the constant impact measurement and continuous improvement that our customers need as they seek to improve lives, whether those lives are human, other species or entire ecosystems.

So we have much that we can use, but there's also much that we need to invent, because as we think about large, complex, living data systems, as we think about the dynamics of human cognition, population behavior, economies, natural ecosystems, oceans, air streams, et cetera, vector-borne diseases, we actually have to think about data, data lakes, analytics, and indeed, the insight generation compute that we use differently. We need to humanize this technology, and we need to make it increasingly accessible, particularly in the Global South, so that we're building capacity that begins to undo some of the data divides, technology divides, compute divides, that are utterly necessary if we, together, are going to achieve our goals toward more sustainability.

So it's a pleasure to be able to enter into this conversation with you. Here are some links that I hope you will follow so that you can come and help build with us. Feel free to reach out to me, my email address is here. And there's some links here to the Amazon Sustainability Data Initiative, so that you can learn more about the data sets there, as well as how you might apply for some of the grants that are deployed through our cloud program.

In my team, leverages credits and specialists to do fully funded POCs, so if you would like to use advanced compute on AWS Cloud to drive innovation in your own services, either for more sustainable compute or compute for sustainability, I would be more than happy to hear from you, and to work together to ensure that we are improving the life of our fellow humans, of other species, and of our natural ecosystems, so that we can together achieve a more sustainable future. Thank you.