Operator: Welcome, everyone. Thank you for standing by for Alphabet's AI Responsibility & Data Center Energy Strategy call.

All participants are in a listen-only mode.

I would like to hand the conference over to your host today, Jim Friedland, Head of Investor Relations. Please go ahead.

Jim Friedland, Head of Investor Relations: Thanks. Hi, everyone. Thanks for joining our AI Responsibility & Data Center Energy Strategy call today. I'm Jim Friedland, Head of Investor Relations for Alphabet.

We're pleased to be joined today by our subject matter experts, David Bledin, VP, Trust & Safety and Responsible AI, and Amanda Peterson Corio, Global Head of Energy & Power.

Also joining us is Nathalia Valdes from our Investor Relations team, who many of you know, and she will be moderating today's discussion.

And with that, I'll turn it over to Nathalia.

Nathalia Valdes, Investor Relations: Thanks, Jim, and thank you all for joining us today and for submitting your questions ahead of time.

We've grouped today's discussion around the most common themes you shared across both AI responsibility and our data center energy strategy.

We'll begin on the AI responsibility side with David.

David, you've been with Google for over 15 years. Before we dive into questions, could you share a bit about your current role and what your team focuses on here at Google?

David Bledin, VP, Trust & Safety and Responsible AI: Sure. But first off, I really appreciate the opportunity to connect with this group on Google's Responsible AI approach.

I've been at Google for some time, as Nathalia mentioned, the last nine of which have been in an organization called Trust & Safety, focused on a variety of different roles. I led Product Policy, focused on egregious harms like child safety and violent extremism. I led our Incident Command and Elections Integrity teams, which have been quite busy over the last couple of weeks with the U.S., Brazil, and Israel elections, as you can imagine. And now I lead our Responsible AI team.

My current team covers several areas of Responsible AI: End-to-end governance across the dozens of models and the hundreds of applications launching across Google every year; testing, including teams focused on running scaled evaluations, red-teaming our models for unknown and emerging risks, and a novel testing team developing new approaches for areas like agentic and personalization; and several other teams like

focus on areas like fairness, neutrality, and health.

I'll note that while my team is called Responsible AI, this work is really infused across our entire Trust & Safety team and beyond. Nowadays, everybody at Google is working on Responsible AI.

In that vein, we partner extremely closely with many teams. For example, with Google DeepMind, we collaborate on testing their models and co-developing mitigations to ensure those models align with desired behaviors set by our policies and frameworks; and with the many Product Areas across Google partnering closely with product and engineering teams to ensure the launches are bold and responsible, and equally important, post-launch, running monitoring and funneling key insights back to those teams for continuous improvement.

Nathalia Valdes, Investor Relations: Thanks, David. The pace of change is so incredible here. How does Alphabet ensure its governance, policies and processes remain fit for purpose, and how much of our approach to Responsible AI has changed over the last two to three years?

David Bledin, VP, Trust & Safety and Responsible AI: I'll address the question in reverse, maybe starting with how our approach has changed over the last three years, and then I'll focus on how we remain fit for purpose.

So, over the last two to three years, Alphabet has evolved from high-level aspirational guidelines -- our 2018 AI Principles, we were one of the first companies to set principles around AI, our north star for all of our products and models -- to an institutionalized, full-stack operational framework to ensure those principles are reviewed in hundreds of Gen AI launches across the company every year.

We now have governance for every stage of the lifecycle, from pre-training dataset curation, to real-time application-layer filters and post-launch telemetry.

And we institutionalized this accountability across all business units through a program called the Responsible AI at Scale, or RAISE, initiative. And RAISE is a systematic framework for identifying, mitigating and documenting responsibility risks across Google surfaces. It includes pre-launch risk assessments baked into launch infrastructure. This is the same infrastructure that Product teams are using, they're already using. We're meeting them where they are. And with these risk assessments, we set and ensure proportionate requirements based on the risk dimensions of a given launch.

And then as we get closer to an actual launch, attestation and gating "bit flip" to ensure all of those requirements are met and documented for audit readiness.

So just as an example, if a launch is focused on a sensitive vertical-like Health, via RAISE, we require health-specific evaluations supported by our in-house clinical team to ensure the launch is meeting pre-established launch thresholds for Health safety and

responsibility.

Our governance also ensures that we're aligning with international regulatory standards. Our internal governance has evolved to directly operationalize emerging statutory compliance regimes, such as the EU AI Act and California SB-53, ensuring our internal benchmarks and transparency reports are fully aligned with external global standards.

Now, Nathalia, getting to the second part of your question, how we remain fit for purpose, we conduct continuous threat modeling and evolving our policies accordingly.

First is ensuring we have a holistic understanding of the harm journeys we are concerned about in the AI space. We have developed an RAI risk taxonomy of all risks across safety, security, privacy, and other dimensions, and treat this risk taxonomy as a living document.

We continually pressure-test our models using specialized internal testing teams alongside external domain experts, especially in areas where we need that external domain expertise, areas like Health, Kids & Family, frontier risks, which I'll talk about in a moment, to discover emerging vulnerabilities and feed those findings directly back into our safety policies and baseline requirements.

For the class of frontier risks that are especially severe, risks like Cyber; Chemical, Bio, Nuclear, Radiological (we call it CBRNE); Harmful Manipulation; Misalignment; Loss of Control, we've established dynamic capabilities, triggers. Instead of relying on static, rigid checklists, our Frontier Safety Framework, which we've externalized, uses empirical capability-based triggers, which means that as models develop new, measurable capabilities, they automatically trigger proportionate levels of safety and security reviews and mitigations, ensuring our guardrails -- especially in these frontier risks, these very sensitive areas dynamically scale up and step with raw model power.

Nathalia Valdes, Investor Relations: David, how do you balance the pace of AI innovation with safety, trust, and responsible deployment, and how do senior management and the Board oversee that balance?

David Bledin, VP, Trust & Safety and Responsible AI: This is where we navigate every day, and it's actually at the core of our AI Principles, and our approach is to be bold and responsible, and we do this work together, together with, I mentioned experts, researchers, the AI safety community to get that balance right.

And really, the fundamental premise of this is a belief that innovation and safety are not a zero-sum tradeoff. We are convinced that being bold about AI's opportunities require being deeply responsible in our execution throughout the AI development lifecycle. Getting this balance right is what builds the user trust required for long-term commercial success.

And we achieve this balance via a defense in-depth approach. We protect innovation velocity via a defense in depth by baking safety into the model during training, so things

like alignment tuning, data curation, while layering adaptive outside-of-the-model guardrails, such as real-time application safety filters and other inference time techniques to steer the model towards desired outcomes. And we address emerging risks without retraining.

We also invest in post-launch monitoring and enforcement, such as disabling access for adversarial users repeatedly attempting to jailbreak our models.

Now, how senior management and the Board oversees this balance. We have a committee called the Trust & Compliance Committee, or TCC, which acts as our primary senior executive mechanism, ensuring higher severity safety or compliance concerns cannot be unilaterally bypassed by Product teams. We have this RAISE process that I mentioned earlier in the call.

And while the full Alphabet Board oversees our overall AI strategy and execution, two committees provide specialized independent oversight of AI-related risks.

Our Risk and Compliance Committee focuses on evolving regulatory and operational landscape, while the Audit Committee oversees the reliability, capacity, and business continuity of our underlying infrastructure.

Nathalia Valdes, Investor Relations: David, when internal teams disagree or raise safety objections during launch readiness, who holds the ultimate authority to delay a rollout, and how frequently is that lever exercised?

David Bledin, VP, Trust & Safety and Responsible AI: This is such an important question.

And at Alphabet, Product Areas and product owners are directly accountable for their launches. This is really important. For Google to be as innovative and to move with high velocity in this market, product owners, they own both the commercial success and the risk profile of their products, ensuring safety is integrated directly into the product development lifecycle, as I mentioned, rather than treated as an external hurdle.

Specialized review teams, including my team, Trust & Safety, Legal, our Privacy team, engage with these Product teams early in the developmental lifecycle, and getting that balance of bold and responsible right is not a single gateway, but it's an iterative, collaborative process of continuous testing, adversarial red-teaming by my team and others, and safety mitigations in advance of any launch.

If a product launch does present high-severity and residual risks after we have a good understanding and we've synthesized all of the plans and executed mitigations, the launch cannot proceed unilaterally. And so this is the exception where PAs are accountable for their launches, but if there's a high severity residual risk, that decision is reviewed by the Trust & Compliance Committee that I mentioned earlier, which encompasses many of our Google Leads.

Because our safety reviews are deeply collaborative and iterative, it's this constant cycle of leading up to a launch and beyond, outright rollout delays due to unresolved disputes are pretty rare. Instead, our development timelines natively build in these iterations for safety mitigations, treating risk calibration as a standard, active feature of launch readiness.

Nathalia Valdes, Investor Relations: Thanks, David.

Talk to us a little bit about what operational processes we have in place to give investors confidence that safety and risk concerns are addressed proactively and before the commercial rollouts happen.

David Bledin, VP, Trust & Safety and Responsible AI: I mentioned some of these already, but repetition doesn't spoil the prayer.

I mentioned mandatory RAISE gates. Every AI-enabled commercial launch at Google must adhere to this Responsible AI at Scale framework, completing intake, risk assessments, required testing and pre-launch attestation and risk reviews.

We also have a two-pronged pre-deployment testing system where systems undergo structured, quantitative evaluations to benchmark performance against known risk areas, areas like safety, security, child safety, neutrality, alongside extensive adversarial red-teaming by internal specialized units -- again, my team is one of the bigger teams at Google, but there are others -- and external domain experts to identify unknown vulnerabilities.

We also leverage controlled phased rollouts. So, higher risk commercial releases follow staged delivery schedules -- you saw that with some of our recent models, like Argon and Carbon -- progressing through trusted-tester feedback programs, before launching to general availability.

And really important, our work does not stop with launch. We also invest in continuous post-launch telemetry and incident management. Launch is not the finish line. For operational safety, it's maintained through real-time automated monitoring, actioning adversarial users that violate our Prohibited Use Policy, and formalized incident remediation procedures to address emerging issues immediately.

Nathalia Valdes, Investor Relations: Now, as Gemini increasingly integrates personal context, connected applications, and the ability to take action on users' behalf, how are you evolving privacy, consent, and permission controls to ensure that users maintain meaningful oversight over this data and their AI interactions?

David Bledin, VP, Trust & Safety and Responsible AI: Especially with these emerging capabilities like personalization and agentic, trust is a core strategic differentiator for us. As Gemini evolves from answering questions to executing personalized, multi-step agentic tasks through Workspace, Search, Chrome, Android, this user trust is

paramount.

And our architecture is grounded in a core principle. Users maintain absolute control over what context Gemini accesses, what it remembers, and what actions it executes.

Just to delve a little bit deeper into that, we require explicit opt-in for context access. So cross-app integrations, when Gemini or other surfaces are using things like Gmail data or Calendar data for personalization, this is all disabled by default. Users must provide explicit, granular opt-in to connect those individual apps, managed directly through existing Google Account privacy controls.

We also maintain rigorous data boundaries. We maintain a strict boundary between user context and foundational model development. Alphabet does not train generative models on raw personalized data from Gmail, Drive, Calendar, or Photos.

We provide verifiable grounding and granular memory controls. So, to ensure transparency, Gemini cites specific connected sources so users can independently verify answers.

Users also retain full oversight of AI memory with controls to view, edit, purge, remember context at any time, alongside ephemeral Temporary Chats for off-the-record queries.

We also require mandatory human confirmation for high-impact actions. So, if you're using an agent like Spark or an agent via Workspace or other enterprise capabilities, our agentic workflows operate within isolated execution environments that provide real-time "stop" and "take control" handoffs.

Explicit human confirmation is strictly enforced before completing high-stakes irreversible actions, such as sending communications, editing shared documents, deleting data, or completing financial transactions.

Nathalia Valdes, Investor Relations: Thanks, David.

Shifting to the regulatory side, what policies and frameworks is Alphabet actively supporting or lobbying for, and how do you view the outlook for federal AI regulations?

David Bledin, VP, Trust & Safety and Responsible AI: As I mentioned, Google has long called for an approach to AI that is both bold and responsible, so those AI Principles, they govern our own north star behavior, but also how we hope that the regulatory environment around AI shapes up. And we've said this a lot, that AI is important -- AI is too important not to regulate and too important not to regulate well.

And so we strongly support a pragmatic, dynamic, and evidence-based regulatory framework that recognizes the unique challenges and opportunities posed by frontier AI on the one hand and widely-deployed AI on the other.

In the U.S., we support a federal framework for frontier AI risks, which will identify priority requirements and critical safety practices, rather than unhelpful patchwork of state laws.

We also foresee that new organizations, such as a public-private standards organization -- Demis had a blog post about this earlier in the year -- could be very useful to advance AI standards development.

Internationally, beyond the U.S. for frontier AI risks, it would be important to collaborate on international standards and benchmarks. Establishing common definitions and shared baseline standards will strengthen the technical rigor of evaluating frontier AI risks.

From widespread AI applications enabled by models at lower levels of capabilities, such as chatbots, we believe the government can draw on and in some cases amend existing laws and rules to address real world outputs and specific harms, harms related to the Workspace, children's safety, privacy and information integrity among others.

And finally, we believe safety is best advanced by clear, consistent standards rather than a patchwork of disparate and fragmented requirements.

Nathalia Valdes, Investor Relations: Thanks, David.

And last question for you. How are we evaluating alignment with recognized AI governance frameworks and standards, and what are our plans for expanding future disclosures in this area?

David Bledin, VP, Trust & Safety and Responsible AI: To answer this question, it helps to break this down into different categories of frameworks and standards, both current and emerging.

So, first, there are binding laws, like California's SB-53 or the EU AI Act. Those come with clear compliance guidelines, and we, of course, comply with laws in the jurisdictions where we operate our business.

Many of those laws include disclosure requirements, largely around frontier AI model capabilities, so that California and New York, for example, one of the new frontier model transparency requirements, the publication of our Frontier Safety Framework. I mentioned that earlier. That's the framework that governs the dynamic capability triggers that guide mitigation and other safeguard implementations when you release a frontier model.

And we've been publishing our FSFs for several years, and we also released a safety and security report with every new major model release. You might have seen that with Argon. It was a pretty in-depth report, over 200 pages. We also released one for

Carbon. And those reports, they provide a lot of information, both in terms of the testing that we do, but also in these emerging areas like chain-of-thought reasoning with agentic, things around eval alignment, et cetera.

Second are voluntary commitments made to the public and governments around the world. These range from pledges around election safety to the White House cybersecurity executive order. We take these commitments really seriously and work closely with governments to ensure our Responsible AI practices are aligned with these frameworks.

And then third, these proposals and public policy debates that have not become formal instruments for either voluntary guidelines or binding regulation. Given the increase in frontier model capabilities over the past few months, there are a number of these new proposals being debated, and we're studying each of these proposals carefully, including their impacts on their business and safety priorities.

Finally, I want to point out that we've been long proponents of AI transparency. We've pioneered the content of a model card and routinely published our AI safety research with the aim of helping the entire AI ecosystem build increasingly sophisticated safety and security approaches to AI.

And we publicly share our annual Responsible AI Progress Report, which represents our multi-layered approach to AI governance, both some of the things that we talked about today, this RAISE governance approach of the framework, as well as how we think about safety across the AI development lifecycle, but also in terms of this range of mitigations that we have to align our models with desired behavior.

Nathalia Valdes, Investor Relations: Well, thank you so much, David. We really appreciate your time today.

David Bledin, VP, Trust & Safety and Responsible AI: Thank you.

Nathalia Valdes, Investor Relations: With that, let's pivot to our data center energy strategy. I'm pleased to now welcome Amanda Peterson Corio to the call.

Amanda serves as Google's Global Head of Energy & Power, and she has been with us for over ten years and oversees the energy infrastructure and strategy required to power Google's data center operations.

Amanda Peterson Corio, Head of Energy & Power: Thank you so much for the warm introduction. It's great to be back with this group again. I'm looking forward to answering your questions.

Nathalia Valdes, Investor Relations: Well, Amanda, you joined us last year for a conversation on data center energy strategy, and a lot has changed since then. Before we dive into the specific investor themes, could you give us a quick overview of what

has shifted in your space in the last year?

Amanda Peterson Corio, Head of Energy & Power: Certainly. There's definitely been a lot of change in the last 12 months. Last year, as the scale of electricity demand was really moving to the forefront of market conversations, I shared our "energy first" framework, which is the idea that smart power development is a foundational driver of how and where we build.

As historic compute demand collided with grid and supply chain constraints this past year, the strategy matured into one defined by three imperatives: Speed, scale, and delivery certainty.

At the same time, we've seen a noticeable shift in the public sentiment. Communities are rightfully asking harder questions about power, water, and who pays for grid growth.

All of this has required us to reflect on how we show up locally, how we gain more direct control of our pipelines, and how we play a more active role in solving local energy resource and affordability challenges.

For decades, Google has been committed to being a responsible energy partner. This year, strategic moves, like our acquisition of Intersect, alongside our pioneering work on advanced energy solutions, and tools like being able to have demand response, or contract with virtual power plants, gives us the operational muscle to execute with speed, certainty, and care in today's environment.

Nathalia Valdes, Investor Relations: That brings us to one of the key questions we've heard from investors following last week's U.S. election. In state after state, data center growth has become a focal point. How is Google managing the risk of local moratoria, delays, and potential challenges to the pace of our domestic build-out?

Amanda Peterson Corio, Head of Energy & Power: Maintaining domestic technology leadership is a priority that isn't going away, and investing in the U.S. compute infrastructure is foundational to achieving that vision.

At the same time, voters, politicians, and local communities have made it clear that it's paramount that data center providers take a responsible approach to infrastructure development. We agree and have long been committed to this approach.

When done right, large loads like data centers should not be a burden to communities. They should help reduce electricity costs and create a more reliable and resilient electricity grid.

As we at Google look to the future, we'll continue to partner with governors, state legislators, and local leaders to ensure that data center developers are held to high standards of practice.

For us specifically, responsible data center development really means first engaging transparently and proactively so we can build community trust right from the beginning.

Second, protecting electric bills by paying for 100% of the electricity our data centers use, as well as the direct infrastructure cost required to connect them to the grid.

Third, we are minimizing our local impact on resources by engaging in responsible water stewardship and mitigating any noise or light pollution from our data centers.

We're also investing deeply in the community. For all-new campuses, we're committed to paying full real property tax on our data center buildings and land. We're also investing in community services and continue to support local businesses.

And finally, we're hiring locally. We are prioritizing hiring employees and suppliers that live within the community and continuing to fund long-term worker training and skilling programs in our data center communities.

It all goes to our goal that has always been to set the industry standard and serve as long-term partners everywhere where we operate, who can be counted on to do the right things in these places that we call our home.

Nathalia Valdes, Investor Relations: Amanda, with public sentiment increasingly linking AI data center growth to rising local power bills, how is Google protecting everyday ratepayers while safeguarding our long-term license to operate?

Amanda Peterson Corio, Head of Energy & Power: We believe, first and foremost, that energy growth and ratepayer protection must go hand in hand. As I mentioned, that starts with us paying for 100% of the electricity that our data centers use, as well as the energy infrastructure cost required to connect us to the grid.

We back that up with electricity structures that are fit for purpose for these much larger loads, like our Capacity Commitment Framework and our Clean Transition Tariff, which is a great cross-design for how we interconnect to the grid.

These new pathways are specifically designed to insulate everyday customers from the cost of our growth.

But beyond simply covering our own footprint cost, we invest in new energy generation. This expands capacity for the entire system and helps to lower the cost of electricity for others connected to the grid.

You see this reflected in several recent utility filings where our data centers are located. For example, in Entergy Arkansas, their cost modeling shows that our agreement will return more than $2.1 billion in savings back to the utility's 738,000 customers.

In Minnesota, Xcel Energy's analysis estimates that our agreement will result in $1.1 billion in net benefits for other electricity customers from 2028 to 2043, while making the

world's largest long duration battery storage project announced to date cost-effective to the entire system.

And in Georgia, our agreement with Georgia Power to support their uprates of their nuclear power plants Vogtle and Hatch are going to unlock an additional 96 megawatts of around-the-clock carbon-free energy and is expected to enable approximately $900 million of benefits to Georgia Power customers over the life of the unit.

In addition to this, we are also working to improve energy affordability directly in our data center communities through our Energy Impact Program. We began this program in 2023 to collaborate with governments, utilities, and non-profits to fund important energy home weatherization, efficiency upgrades, and local workforce training, and we're already seeing meaningful results.

In the Carolinas, for example, our Energy Impact Program delivered annual cost savings between $545 and $729 per year for the homes that participated.

And in Nebraska, where we recently announced a $6 million contribution to the Omaha Public Power District and the United Way of the Midlands to enhance energy efficiency and weatherization, our partners estimate the investments will lower energy bills by up to 30%, saving households between $840 and $1,260 every single year.

Put simply, we want to ensure our growth strengthens local grids without everyday families paying the price.

Nathalia Valdes, Investor Relations: Amanda, now give us a sense of how Google's approach to power development for data centers is most differentiated, and how are these approaches translating into measurable advantages for securing power and accelerating deployment?

Amanda Peterson Corio, Head of Energy & Power: So, at Google, our approach has always started with thinking about what makes sense to the business, what makes sense for our users, and the electricity system in the long run.

We know that we're building long-term assets, and we think about what that means for power supply over the entire life of the asset. And that intentional design and thoughtfulness combined with our 20-year-plus operating history of powering data centers defines how we are leading in this moment.

Given the uncertainty in the near-term, this approach has never been more important. If we focus on building smart assets, remain committed to our principles of responsible power design and continue to diversify the multiple pathways to power, the speed and scale to capacity will follow.

In our regulated energy markets, our contract models, specifically the Capacity Commitment Framework and the Clean Transition Tariff, have now become industry blueprints. Both help balance rapid demand growth, utility planning and consumer

protection, which is why utilities and public utility commissions are adopting them across the country.

In our deregulated markets, we ensure that we are Bringing Our Own New Generation, whether that's by co-locating power at our data centers, or through system delivery of power through our wholesale deregulated electricity markets.

In today's environment where power assets are scarce and timing is critical, we've been able to cut the time required to negotiate and execute power purchase agreements for generation capacity by roughly 80%, and this is through our work with LevelTen and the co-development of the LEAP power power purchasing model, which standardizes and streamlines our contract execution timing.

This combined with our Bring Our Own New Generation pathways have enabled us to protect the system affordability while also moving at a speed that our growth requires.

And finally, we aren't just waiting on long-lead transmission lines to become available. We're unlocking immediate speed to power by getting more out of the existing grid. We've already integrated a gigawatt of demand response capabilities into our long-term energy contracts with multiple utilities across the U.S. This unique flexibility enables us to connect to power faster.

We've also established a first-of-its-kind Virtual Power Plant collaboration that helps to support the grid during periods of peak system stress, while also helping to lower consumer energy bills.

Last month, we announced an expanded agreement with Voltus to build a gigawatt of virtual power plants within the PJM grid region. This region serves 67 million people across 13 states and Washington, D.C. We've scaled this from an initial hundred megawatt pilot to a gigawatt of Virtual Power Plant supply, representing one of the largest distributed capacity portfolios developed by a hyperscaler to date.

Ultimately, by combining transaction velocity, commercial and grid innovation and thoughtfully building a diversified portfolio power supply, we are creating the capacity headroom that allows us to scale responsibly ahead of the market.

Nathalia Valdes, Investor Relations: Amanda, you mentioned Bringing Your Own New Generation, and this is a topic we received several questions on. Could you walk us through Google's approach here and where behind-the-meter solutions fit in?

Amanda Peterson Corio, Head of Energy & Power: So, our starting principle is simple. There is no one-size-fits-all model. Electricity market dynamics and, again, regulatory structures vary greatly across many markets where we operate.

But where the regulations and market structures permit, behind-the-meter solutions can serve as a practical and efficient bridge for unlocking speed to power and having that more within our control.

At the same time, we're very clear-eyed about the role of the centralized grid and the role it plays in the long-term as we continue to build out these large amounts of new generation and loads. Our grid system was designed for a reason. It provides essential scale, diversity of supply, lower cost, and long-term optimization. This is why our overarching approach focuses on strengthening the grid while also bringing new resources to meet speed and scale of capacity.

So we work in close coordination, both with our utility partners and with generation owners, as well as developers and our in-house development team with Intersect, and regional grid operators to plan and develop solutions that are both additive and responsible.

We are not seeking to divert existing grid-connected resources away from the electricity grid into our data centers. Our focus is on bringing new capacity online in many different forms that may show up, and that's tailored to the regional needs and realities of each market where we operate.

Nathalia Valdes, Investor Relations: So now, how is Google approaching the public calls for better transparency when it comes to our data center projects?

Amanda Peterson Corio, Head of Energy & Power: Across the industry, we've seen that pushback can happen when development feels secretive, unilateral, or abrupt. So, engaging transparently and proactively is core to our approach in every community when we enter.

To ensure that we're being open and transparent in our work, Google has eliminated the use of NDAs with state and local governments. Beyond that, we only seek confidentiality in limited cases, such as for trade secrets or information impacting the security of our operations.

But we've adapted our approach so that we're disclosing our involvement in projects earlier, either with our first application or first public meeting, whichever comes first in the process.

We're publishing additional project-level data about our use of resources in each community, such as our water use, and how we're meeting sound and light limits and any traffic impacts.

And throughout our data center development process, we host town halls and establish local advisory councils with the community, which provide direct forums for our teams to listen, to adapt, and to solve for local concerns.

Nathalia Valdes, Investor Relations: As AI related infrastructure continues to expand, can you talk about Google's efforts to maximize energy and water efficiency in our own data centers?

Amanda Peterson Corio, Head of Energy & Power: This is -- first and foremost, at Google, we are committed to maximizing efficiency at every layer of our operations. We take a full-stack approach by designing every element for maximum performance per watt.

So, for example, at the facility level last year, Google's data centers were delivering three times more computing performance per unit of energy than five years previously, while also using 83% less overhead energy than the industry average for a data center. That efficiency is accelerated by our TPUs.

So, our Ironwood TPU, which launched in early 2025, is nearly 30 times more power efficient than our first Cloud TPU from 2018. And our latest generation of TPUs announced this year, TPU 8t and TPU 8i, deliver up to two times better performance per watt over Ironwood.

So when it comes to water, we also look at efficiency as well. Google has decades of experience in water stewardship, including an ambition to replenish more water than we consume at our sites by the end of this decade.

At each new data center campus, we apply a data-driven water risk framework to assess local watersheds, and only consider water cooling if local resources are healthy and resilient. If there is any water stress, we cool our data centers with air cooling, which uses very, very minimal water.

Nathalia Valdes, Investor Relations: This last question, Amanda, ties back to our climate ambitions. Given the infrastructure expansion is increasing faster than many local grids can decarbonize, how is the team assessing the net-zero and carbon-free energy goals?

Amanda Peterson Corio, Head of Energy & Power: At Google, hopefully the market has seen, we have never stopped expanding access to clean, affordable, reliable energy. And today, we see our growth really as an opportunity and has been an opportunity over the past few years to drive innovation that accelerates the development of electricity decarbonization solutions that haven't yet existed at scale, but that the world needs.

When we first set our environmental ambitions, including running on 24/7 carbon-free energy and achieving net-zero emissions, we knew they were moonshots. We intentionally chose those ambitious goals knowing that the broader energy system and advanced technologies that are required to fully achieve this goal didn't even exist yet.

Since then, we've seen our capital and commercial commitments have helped scale and lower the cost of the technology needed for a reliable around-the-clock clean power, including things like our next generation geothermal, our work in advanced nuclear, and our investments in long-duration energy storage.

And while the path to these ambitions continues to be complex and nonlinear, our long-term commitment to reducing our environmental impact and achieving 24/7 carbon-free energy continues to remain the north star.

Nathalia Valdes, Investor Relations: Well, thank you so much, Amanda. We really appreciate you joining us today.

Amanda Peterson Corio, Head of Energy & Power: Thank you for having me.

Nathalia Valdes, Investor Relations: With that, I'll turn the call back over to Jim.

Jim Friedland, Head of Investor Relations: Well, thanks, everybody. Thanks, Nathalia, and thanks, everybody, for joining us today. Appreciate everyone from the investment community joining us, and especially appreciate our subject matter experts, both Amanda and David. It was a great, very helpful call.

Thanks again, everybody, have a good day, and I'll turn it over to the operator to close us out. Thanks, everyone.

Operator: Thank you, everyone. This concludes today's call. Thank you for participating. You may now disconnect.