

Edited Transcript

SpaceX Third Quarter 2026 Webcast October 20, 2026

The information contained in this event transcript is a textual representation of SpaceX's webcast. While efforts are made to provide an accurate transcription, there may be material errors, omissions, or inaccuracies in the reporting of the substance of the webcast. In no way does SpaceX assume any responsibility for any investment or other decisions made based upon the information provided in any event transcript. Users are advised to review SpaceX's webcast itself and SpaceX's SEC filings before making any investment or other decisions.

Corporate Participants: Elon Musk, Gwynne Shotwell, Bret Johnsen, Andrea Williams

Management Discussion Section

Andrea Williams: Good afternoon, everyone. Thank you for joining us today for SpaceX's third quarter 2026 earnings presentation. I'm Andrea Williams, Head of Investor Relations. Joining me today are Elon, Gwynne, and Bret. We are coming to you live from our Starlink factory and future Starlink production site here in Bastrop, Texas, and we are livestreaming on X.

Our third quarter financial results were announced just after 3:00 p.m. Texas time. And all relevant materials have been published to our Investor Relations website at ir.spacex.com.

For today's presentation, our prepared remarks will be followed by a question-and-answer session. We will take live Q&A from equity analysts and then move to questions directly from our investors that were polled on our dedicated earnings Q&A platform, which the SpaceX team built with Grok.

During this webcast, we will discuss our business outlook and make forward-looking statements. These statements are based on our predictions and expectations as of today. Actual events or results could differ materially due to a number of risks and uncertainties, including those mentioned in our earnings release and the most recent filings with the SEC, including our final prospectus.

Additionally, during this call we may discuss certain non-GAAP financial measures. In today's earnings release and our filings with the SEC, you can find additional disclosures regarding these non-GAAP measures, including all of the reconciliations with comparable GAAP measures.

For the third quarter of 2026, we announced the following financial results as compared to the prior-year period:

- Revenues of \$14.6 billion, up 161% from \$5.6 billion
- Net income of \$4.4 billion, an improvement of \$5.5 billion from a net loss of \$1.1 billion
- And Adjusted EBITDA of \$8.8 billion, up 306% from \$2.2 billion

Now, I will hand the call over to Elon.

Elon Musk: Thank you. We are in an inflection moment for SpaceX across launch, connectivity, and Super Intelligence.

Starting with Starship, Flight 14 was a critical milestone. Starship made it to orbit for the first time and deployed operational Starlink V3 satellites. We have a recap video we're going to play.

(Video)

For Flight 15, which we expect to happen next month, we're going to attempt to catch both the booster and the ship. If we succeed, it will be the first time that an orbital launch vehicle has been fully recovered. This is a massive technology milestone. The only way that we can extend consciousness beyond Earth on a sustained basis is with a fully reusable orbital launch vehicle. This is, I think, a profound technology milestone in the history of civilization and consciousness.

Let's look at what a simulation of Flight 15 would look like if we succeed. If we don't succeed on Flight 15, we'll get there. Success is not guaranteed on these first flights, but I'd say it's more likely than not that this succeeds next flight. And if not, we have a steady stream of Starship flight attempts through the end of this year and then accelerating next year.

With that, the Flight 15 simulation.

(Video)

We've also announced the development of the Louisiana spaceport. So that's going to be an incredible asset to the country and to the world that will, I think at some point, probably have on the order of 20 or more launch towers and enable the transport of millions of tons per year to orbit and beyond. That's really going to be a big deal.

We expect Starship to fly weekly at some point in the first half of next year, and then in the second half of next year, twice weekly, or maybe even more frequently than that by the end of next year. The rate at which we launch Starship will increase dramatically as we solve rapid and complete reusability.

We'll have complete reusability. Then at first, there will be a little bit of refurbishment that's needed between each flight, and we'll iron out the need for any refurbishment such that you can immediately land the rocket and launch it again, just like an airplane. As I said, this is a profound technology milestone that most people in the world don't realize how important it is, but it will be obvious in hindsight.

Starlink is also just getting started. Flight 14 deployed 26 terabits per second of Starlink bandwidth. That's equivalent to 10 Falcon 9 launches. A single Starship launch, which wasn't even fully loaded — it was about maybe half loaded — and it did the equivalent of 10 Falcon 9 launches.

What I expect over time — in the not too distant future, probably within three years — is that we will have about 100x the bandwidth of the current Starlink constellation. Because we will have launched satellites with an order of magnitude more bandwidth, and we will place into orbit an order of magnitude more Starlink satellites. So probably something on the order of 100,000 satellites, each of which have more than 10x the bandwidth of the current Starlink system.

I think revenue for Starlink will be at least somewhat proportionate to the total bandwidth delivered. And we will also have high-bandwidth direct-to-cell capability, and we will be building out our terrestrial network. There's broad-based connectivity anywhere that systems like Starlink are allowed to operate. Anywhere in the world that we're allowed to operate. And I think the demand for bandwidth is going to be dramatically higher than it is today.

I think in three to five years, there's probably a demand for 1,000x more bandwidth than there is currently, because AI or SI agents will, on a per-agent basis, use far more bandwidth than a person, and there will be

far more agents than there are people. This ends up requiring a several-order-of-magnitude increase in communication bandwidth in order to satisfy that communication demand.

Talking about SpaceX Super Intelligence. SpaceX SI. Sexy... whatever.

As I think people are starting to realize, the constraint on deployment of super intelligence is more in power and cooling and infrastructure than it is in chips. What we've done at SpaceX is we've taken a team that has been able to achieve a market share in mass to orbit that is almost 90% of all mass to orbit of Earth, where China is about, I think, roughly 6%, and the rest of the world combined, including the rest of the U.S., is about 4%.

A team that has been able to achieve results that are an order of magnitude greater than the rest of all competitors in space combined, and applied that same team to the deployment of terrestrial compute. This is a bit like the New York Yankees playing in Little League. It's not fair. But the net result is being able to deploy power, cooling, and infrastructure at a rate that is vastly faster than anyone else, and may at some point be faster than everyone else combined. So, just let that sink in for a minute.

I've been spending personally a lot of my time in Memphis, often four days a week, night and day, just getting the Memphis supercompute cluster, Colossus II, online. We're tracking to be able to bring roughly 200,000 GB300s online per month for every month for the remainder of this year, and continuing that pace next year, with next year being primarily Rubin, Vera Rubin compute.

At this point, I would say I'm fairly confident of being able to bring an incremental 10 GW of compute online in 2027. That would be in addition to the roughly 3 GW maybe 2.5 to 3 GW, that we'll have online by the end of this year. So, let me put it this way: I would be surprised if we brought on less than 13 GW of compute by the end of next year.

And as I mentioned earlier, power, cooling, and electrical equipment infrastructure is somewhat of a limiting factor. So, to make sure that is not a constraint for our build-out, we're actually going to aim for an incremental 20 GW of power, cooling, and infrastructure next year so that the amount of power and cooling that we have is far in excess of what is needed to turn the GPUs on.

There's a non-zero possibility — maybe it's a small possibility at this point, but it's not zero — of SpaceX having over 20 gigawatts of Super Intelligence compute online next year. A small possibility, but it used to be zero. It's not zero anymore.

We're also planning to have the most powerful training cluster online by the end of Q2, which will be 440,000 Vera Rubin computers deployed for optimal training. This will, I think, be extremely helpful for getting Grok to be potentially the best model, but certainly to be extremely competitive.

On the Grok Bot side, we're seeing the fastest growth of a product that I've ever seen in my life for Grok Bot. I've never seen anything grow this fast. We're also seeing that the usage of Grok Bot grows over time per user. It does not decline. As we make Grok bot more useful, we're seeing users use it more per week, not less. And we'll be deploying Grok Bot at scale for enterprise customers, which is a very big deal, so they can be used effectively in teams by companies as though it were a virtual human. We already are in early deployment with some trusted customers, and we expect to be in wide deployment within a few weeks for enterprise customers for Grok Bot.

I think Grok bot may end up being what uses most of our Super Intelligence compute.

In a nutshell, what this implies for SpaceX annual recurring revenue by the end of '27 is somewhere between \$500 billion and \$1 trillion. Bret? Are you going next?

Gwynne Shotwell: That's okay, Bret. You can go instead.

Elon Musk: Sorry. I thought — on that crazy note, I thought —

Gwynne Shotwell: On that crazy note.

Elon Musk: Yes. I don't want to give Bret a heart attack.

Gwynne Shotwell: Thanks, Elon.

We had an exceptional execution and quarter, again. And today I'll focus both on the operational and the commercial highlights, as well as talk about the priority initiatives for next quarter.

Starting with our Space business. We remain the leading launch provider for the world, flying the vast majority of payload mass to orbit. During the third quarter, our 35 launches included 21 flights of Starlink on Falcon 9, two flights of Starship, a Falcon Heavy, which I'll talk about in a second, and 12 high-profile customer missions, including the Falcon Heavy launch, which flew NASA's Roman Space Telescope into deep space. We also flew six national security space launch missions.

We've also secured important new business, winning three additional human spaceflight missions on Dragon, which will take us through 2028, along with the Starlink communications contract on NASA's Artemis III mission. In addition, we signed 18 new national security space launch missions, a contract to deliver payload to the lunar surface with Starship, and three new Starfall missions on Starship. We just did a Starfall mission today, by the way, on Falcon 9. Starfall is our space capsule, which provides routine, relatively low-cost, access to microgravity environments for scientific research, including improving our understanding of diseases like cancer and diabetes and developing better protocols for patients here on Earth.

During the final week of the quarter, we flew Starship Flight 14 to orbit, as Elon talked about, followed just after, on October 1, by a triple header. We flew three Falcon missions across our three Falcon launch sites in just 13 hours: Crew-13, which we brought safely to the International Space Station; a Falcon Heavy NRO mission; and a rideshare mission flying 130 payloads. It's an incredible display of what we can deliver in a single day across crew, commercial, and national security missions.

We're also incredibly excited for what's ahead for Starship: flying the largest payloads to help build the space economy in orbit and on other planets and the Moon. As Elon said, Starship's Flight 14 was a major technical milestone, and we remain on track with our goals to bring back both the first and second stages to the tower this year, refly the second stage by early next year, and ramp our launch cadence as we make progress towards rapid reusability. We'll attempt an on-orbit refueling early next year, and we'll begin flying Mobile V2 and Starlink satellites in 2027.

We have had significant constructive engagement with the FAA, and we expect the license for our Flight 15 to include return to launch site and tower catch for both stages. I'm also extremely excited about expanding Starship operations to Louisiana with the second Louisiana purchase. The geography of Louisiana, the community, and their work ethic are so aligned for an amazing Starship future.

Moving on to Starlink. We had another standout quarter, driven by continued growth across both our consumer, enterprise, and government businesses. We ended the quarter active in 174 markets, including new market launches for Vietnam and Uganda.

Looking at subscriber growth, we added 1.6 million consumers — not total users, but consumers — in the quarter, reaching just under 14 million global broadband consumer subscribers. As I mentioned, this

number does not include additional hundreds of thousands of enterprise subscribers and aviation customers.

We remain incredibly inspired by the impact of Starlink on communities around the world, including providing service to hundreds of thousands of students and teachers in schools globally, as well as for emergency response during natural disasters. Over the third quarter, Starlink's rapid response capabilities enabled support across 25 crises, including the Venezuela earthquake, the earthquake in western Colombia, and Hurricane Lowell in Hawaii. In Venezuela alone, Starlink donated over 1,400 kits to response agencies, providing critical Starlink broadband connectivity to over 250,000 people.

As we mentioned last quarter, we continue to see significant growth opportunities on the enterprise and government side of the business, with aviation a continuing major focus and fixed site gaining momentum. Starlink is now contracted for roughly 12,000 planes, almost one quarter of the available aircraft, and we're expanding quickly with both new and existing customers. We signed agreements with three of the four major U.S. carriers, along with a growing list of leading international airlines. During the quarter, Lufthansa, Swiss, Copa Airlines, Korean Air, Austrian Airlines, and Asiana Airlines all flew their first Starlink-equipped aircraft.

To keep pace with this demand, we've quadrupled aviation production rate since the beginning of the year, and we expect to double before the end of the year. Our aviation partners report that Starlink has become a key reason customers choose to fly with them, both business travelers who increasingly rely on full connectivity in the skies and individuals who want to stay connected all the time, everywhere. Our data shows, and frankly our customers tell us, that Starlink-equipped flights carry higher load factors, and full connectivity is quickly becoming table stakes for passengers.

On the government side, Starshield continues to deliver unique mission-critical capabilities, and we see further room for growth. We remain on pace for enterprise and government revenues to exceed consumer connectivity revenues in the future. Each. The successful deployment of the operational V3 broadband satellites on Starship's Flight 14 marks an important step and an inflection point for Starlink. The increased capability from these satellites, as Elon talked about, and the Starship launcher itself, will provide additional capacity across markets and enable new sales in high-demand areas.

We expect the market for communications to grow far beyond just human communication that we traditionally think about. Connecting cars, robots, and the adoption of Super Intelligence will increase communication demands, not in percentages, but in many orders of magnitude. In the years ahead, we expect Starlink alone to represent a significant portion of global Internet traffic.

In Mobile, we continue to expand international partnerships, turning on service with new carriers across multiple countries in the third quarter, including Uganda, the Democratic Republic of Congo, Kazakhstan, and Costa Rica. We surpassed 22 million monthly users on Starlink Mobile as of the end of the quarter, up 37% quarter-over-quarter. Our near-term priority remains flying the Mobile V2 constellation through next year and integrating our spectrum purchases to turn on service in early 2028. Just a few weeks ago, we received approval from the FCC to launch 15,000 Mobile V2 satellites and announced an agreement to acquire a portfolio of low-band spectrum, which will pave the way for Starlink Mobile to become a major mobile carrier in the U.S.

Moving to Super Intelligence. The industry continues to demand ever more massive amounts of compute. We're rapidly expanding capacity to meet both our own needs, as Grok and related products scale, and those of our cloud service customers. Enterprise feedback on recent model releases has been positive, and Grok Bot, in particular, is seeing strong adoption out of the gate. Grok Bot, our agentic assistant,

launched in August, and we already had paid weekly users of 800,000 earlier this month, more than doubling over the last month.

B2B is our fastest-growing category, with companies increasingly using Grok Bot as a force multiplier across sales, marketing, IT, and engineering, treating it as an always-on teammate for complex end-to-end workflows. We're also accelerating our Grok model release cadence, with significant performance gains expected in the near term — which have been, and more are expected, in the near term.

At the same time, the integration of Cursor continues to progress well, strengthening both our product capabilities and our go-to-market reach with a scaled enterprise sales force. Early evidence of this momentum is already visible. Grok Enterprise customers with more than \$1 million in run-rate revenue, including Cursor, were up 50% quarter-over-quarter on a pro forma basis. Last week, we further integrated the portfolio by announcing a new scheme that brings together Grok, Grok Bot, Cursor, and X Premium on a single plan.

In review, our business is experiencing extremely strong momentum across all three segments. Our engineering teams are scaling Starship and next-generation satellites across broadband, mobile, and Starshield, while our commercial teams expand Starlink across aviation, enterprise, government, and global markets. At the same time, we're scaling our Super Intelligence business in the key areas of critical compute infrastructure, leveraging our hardware expertise.

Now, Bret, you get to go.

Bret Johnsen: Thank you, Gwynne.

We delivered another strong quarter in Q3. Revenue grew an impressive 161% year-over-year to \$14.6 billion, fueled by the continued momentum in the connectivity business and our clear inflection point in the Super Intelligence, or SI, business.

Adjusted EBITDA once again grew faster than revenue, to \$8.8 billion, up 306% year-over-year. The combination of our launch leadership, recurring Starlink revenue, and the significant expansion of our SI infrastructure business have combined to drive meaningful operating leverage in the business.

Our balance sheet remains very strong, and our focus on capital allocation continues to keep us in a position to accelerate the pace of innovation that underpins the significant revenue opportunity across all of our businesses.

I'll now walk through the financials by segment.

Space revenues of \$1.8 billion were up 84% sequentially and up 6% year-over-year, anchored by a favorable customer mix. We're on track to deliver strong growth for the year based on our robust customer manifest for Falcon and Dragon vehicles.

As far as capital allocation, we continued to invest in our Starship program, accelerating Raptor engine and vehicle production, advancing Gigabay construction, and progressing multiple launch pads at both Starbase and the Cape. And importantly, we secured the location for the next phase of United States space launch leadership in Louisiana. These investments remain essential to unlocking the step change in cost and capacity that Starship will deliver.

Connectivity again delivered robust revenue growth and profitability, with revenues of \$5.2 billion, up 69% year-over-year and 22% sequentially. As Gwynne mentioned, we reached 13.6 million Starlink consumer subscribers, adding 1.6 million subscribers just in Q3 alone, driven by growth in more developed markets. During the quarter, we began ramping production of our V5 kits, which supported our progress on

alleviating inventory constraints. We strategically managed our channel distribution and prioritized delivery of kits to retail channels ahead of the holiday season, which we expect will lead to increased activations in Q4.

Starlink consumer average revenue per user, or ARPU increased 5% sequentially to \$69 per month, benefiting from a favorable customer-mix dynamic in the quarter. Enterprise and government revenues grew 91% year-over-year and 18% sequentially, highlighted by growth in our Starshield program with the United States government and the continued momentum in the aviation industry. Aviation remains a major growth opportunity, as we're still early in the installation phase with many of the largest carriers worldwide and have line of sight into an accelerated activation path into '27 for our key partners.

Connectivity gross profit doubled year-over-year, with gross margins of 55%, including a one-time benefit of roughly 3%, primarily driven by tariff refunds in the quarter. Excluding the one-time benefit adjusted EBITDA margins expanded roughly 2% year-over-year.

Looking ahead, the satellite roadmap is a key driver of both the capacity and capability of our Connectivity segment. Through the end of the year and throughout 2027, we expect to deploy V3 broadband satellites in meaningful volume, with each Starship launch ramping towards a payload that delivers over 20x the downlink capacity of the current generation of Falcon 9 launches today. In parallel, we will ramp Starlink Mobile V2 satellites on Starship starting in 2027 and leverage our newly enhanced wireless spectrum portfolio to expand our direct-to-cell service and become a major mobile carrier, as Gwynne mentioned, in the United States.

Across all three of our connectivity applications — Starlink Broadband, Starlink Mobile, and Starshield — the launch cadence on Starship is expected to drive a step change in total capacity, coverage, and service quality.

SI, or Super Intelligence, was again the standout growth driver for the quarter. Segment revenue increased substantially to \$7.6 billion, up \$5 billion quarter-over-quarter and \$6.7 billion year-over-year, primarily powered by the ramp of our cloud services agreements. We meaningfully improved the segment's profitability profile as we monetized newly deployed capacity and delivered powerful operating leverage for the business. We ended the third quarter with 1.7 GW of nameplate compute capacity, with additional data halls already online in early October.

For SI Solutions and Infrastructure, revenues were \$7 billion, of which \$608 million came from the Cursor acquisition, closed roughly mid-quarter. And even excluding Cursor, SI Solutions and Infrastructure revenues grew 192% sequentially, driven primarily by the activation of new cloud services agreements. There are many customers asking for compute, and we're actively building the supply to meet both internal and external demand.

We expect to have additional new customers in future deals, and we'll continuously evaluate the highest-margin usage for our compute assets across both new and existing cloud service agreements and serving inference for our own SI products.

What we're seeing in the market is extraordinary. At current pricing, new compute deployments are generating a payback of less than a year. Pricing is currently tracking slightly above the high end of the \$30 to \$50 per watt range we were discussing just last quarter, which further supports the attractiveness of these terrestrial SI compute investments.

We'll continue to be demand-led and prudent in the pace of deployment, matching our disciplined and proactive capital planning for procurement of semiconductor supply to the near-term customer demand. Our current customer contracts for SI compute have the benefit of allowing us to flexibly manage our

capacity for our own model-training purposes, as well as be able to monetize our SI cloud service agreements at favorable pricing as market availability continues to tighten.

Advertising revenues grew 19% year-over-year as we build more momentum on the new advertising tech platform that we rolled out just last quarter, in Q2.

SI segment adjusted EBITDA was \$5 billion, up \$3.8 billion quarter-over-quarter and \$5.4 billion year-over-year, driven primarily by higher-margin cloud service agreements. SI segment operating income was \$431 million, but excluding the \$1.2 billion of stock-based compensation expense related to the Cursor acquisition, the segment delivered \$1.6 billion of operating profit just in Q3.

Double-clicking on the share-based compensation expenses: of the \$60 billion equity value for Cursor when we purchased it, \$23 billion was attributed toward retaining key employees and replacing outstanding equity awards, including revesting agreements for key executives. These equity awards will be recognized as stock-based compensation over roughly the next three-year vesting period.

Total company capital expenditures in the third quarter were \$19.6 billion, primarily in support of our SI compute build-out, Starship infrastructure, and satellite production. We ended the quarter with substantial backlog totaling \$63 billion as of September 30, of which \$17.5 billion was recognized as deferred revenue. Approximately \$42 billion of this backlog is expected to be recognized within the next year.

Looking ahead, demand remains robust across all three segments, particularly in our cloud services business. We remain on track to deliver our previously stated goal of \$100 billion in ARR by the end of this year, defined as annualized revenue based on the month of December.

SI infrastructure continuing to come online at Memphis and Southaven gives us visibility into an accelerated build path through the remainder of this year. As a result, we're on track to reach nearly 2 GW of nameplate compute at the end of this month alone, in October, and have a clear line of sight, as Elon highlighted, to nameplate compute in the range of 2.5 to 3 gigawatts just by the end of this quarter.

Because of the additional capacity coming online this quarter, we expect fourth quarter capital expenditures to now be in the range of \$35 billion to \$40 billion for this quarter. We also expect total company gross margins in Q4 to be more consistent with the second quarter, given some one-time benefits and customer mix that we highlighted in Q3.

Given the increased visibility into our compute contract pipeline over the next six months, we have the opportunity to provide additional look-forward guidance on the acceleration we see in our business in Q1 2027. We now see total company revenues in the range of \$33 billion to \$37 billion for Q1 of 2027.

In closing, we are in a transformative phase for the company's financial profile. We're investing aggressively across SI compute, Starship development and production, and next-generation Starlink broadband and mobile constellations — the three areas we believe will generate the highest long-term returns. These are demand-led investments in markets with visible, tangible returns in the near term.

We're building the infrastructure of the future with discipline, focused on capital efficiency, strong returns on invested capital, and maintaining a strong investment-grade balance sheet, while still moving with the velocity to capture the opportunity in front of us. The combination of our diversified earnings base, focus on extreme vertical integration, and the robust demand in the market for Super Intelligence gives us conviction in the path ahead.

With that, I'll turn it back to Andrea.

Question and Answer Section

Andrea Williams: Thanks, Bret. We're now going to begin the Q&A section. We will take the first question from Adam Jonas at Morgan Stanley.

Adam Jonas, Morgan Stanley: First, Elon, thank you for making flying cars real.

Elon Musk: You're welcome.

Adam Jonas, Morgan Stanley: That was freaking insane. It's my first comment. It's super inspiring for people all over the world, especially young people that want to build.

I got a two-part question on Terafab. You want to make chips. You want to make crazy amounts of chips onshore, or in a safe and secure place. So first, can you confirm if SpaceX has placed orders for EUV machines? If so, when do you expect delivery? And just as a second part to the question, when you take a first-principles approach to the semi fab supply chain — I know you've toured TSMC, ASML, AMD, Lam; these are very impressive companies, as I'm sure you'd agree — where do you see the biggest areas to improve their processes and supply chain when you do it your way?

Elon Musk: Well, I mean, we're just figuring out chip fabrication. But I'll tell you what our internal expectations are, which obviously I put big error bars on, because we don't know if we will actually achieve this or not. But what we're aiming for internally is to make our first sort of test chip functional in approximately May or June of next year. That's just a kind of SRAM chip, something that basically tests the line. That's at the research fab in Austin. And our aspiration is to reach volume production around Q3 of '28 at Terafab at Gibbons Creek, the main Terafab location.

Now these are very aggressive dates, and we're very appreciative of ASML and all of our other suppliers in accelerating equipment deliveries in order to achieve this outcome. Let's just say that outcome of achieving volume production is at least in a set of possible outcomes for Q3 or Q4 of '28. These are obviously insanely fast timeframes for chip production. We do have, I think, some significant potential breakthroughs in chip fabrication, but I don't think I should discuss that on an earnings call. Let's not give away the shop here. But I am optimistic that we will dramatically accelerate chip production and that Terafab will ultimately do what the name suggests, which is to make a terawatt of super intelligence chips per year.

Andrea Williams: We'll take our next question from Doug Anmuth at JPMorgan.

Douglas Anmuth, JPMorgan: Great. Elon, can you talk more about just how you're working through those constraints in power, cooling, and infra, and just what gets you comfortable with the 13 GW ending '27? And then, Bret, if you could also help us understand how you're thinking about the funding requirements just as you look out to that kind of capacity?

Elon Musk: Yes. Well, I mean, 3 GW of that, or very close to 3 GW, is what we expect to have done just in the next few months. You can just look at the rate at which we are actually deploying and turning on compute.

You don't hear a lot of detail from other hyperscalers about how many GPUs they're actually turning on every month. And with us, you are. I'm being as precise as possible about this. It's absolutely insane to be deploying — providing the power and all the electrical equipment, the cooling, the batteries, the networking, and everything else that's needed — to turn on 200,000 GB300s a month, but that is actually what we're doing right now. And that pace will obviously accelerate next year. We're not going to slow down.

If you say, like, how we're going to get all this power, let's just say we actually do have a plan that gets to 20 GW of power and cooling by the end of next year. But that is not something that I should spell out. That is highly proprietary, because if I tell you about it, others are going to try to do the same thing. And I'm not saying we'll have 20 GW of power. I'm saying we have a plan where, if everything goes right, we'll have 20 GW of incremental power, I should say, next year. So, 23 GW roughly in total. Now I don't expect everything to go right. That's why I said I'm confident of 10 GW, highly confident of 10 GW, next year. And 20 GW is difficult but not impossible.

And I mean, some of the things we're doing, I have mentioned publicly. So, we're casting the blades and vanes for the turbines. SpaceX already has a very advanced castings capability that was developed for rocket engines. We have, I think, actually the most advanced casting capability on Earth, because the demands for castings for rocket engines exceed the demands for everything and anything else you could possibly imagine. Taking that incredible team that has, I think objectively, the most advanced casting technology in the world and saying, hey, can you guys also make some blades and vanes for terrestrial turbines, is much easier than what they have done in the past.

And when we talk to the turbine suppliers, they tell us that the blades and vanes are 12 to 18 months behind the rest of the turbine, meaning they can deliver everything, except the blades and vanes, 12 to 18 months before they can deliver the blades and vanes. Casting is this really difficult bottleneck. And if anything, it seems to be getting worse. That means that we can bring power online 12 to 18 months faster than anyone else. That's the kind of advantage that's difficult to replicate. Unless someone else is building an advanced casting shop, they will be unable to compete with SpaceX in this regard.

And we're looking at many different sources of power, not just gas turbines. We obviously want to use solar, wind, and every other source of power that we can obtain. We think there's also a lot of opportunity for inference compute if you manufacture modular inference compute at the factory level, such that the power and cooling are built in at the factory, and then you distribute modular inference compute clusters all over the country, if not all over the world. That's about as much as I should say about scaling power and cooling. But like I said, at this point, I would say I feel highly confident of 10 GW incremental next year, so 13 GW total. And some number over 20 GW — the probability of achieving that is rising rapidly.

Bret Johnsen: Doug, the second part of your question related to funding. I think we certainly have the advantage related to others in that we are investment grade and have a very strong balance sheet. And there are a number of available pools of debt capital that we are talking to at the moment related to funding this incredible growth. The other really important thing to emphasize is the payback period is very short. And so, this capability starts to fund itself in a pretty near-term situation. The fact that we're finishing this quarter with more capability starts the funding mechanism even into the first half of next year.

You have your operational cash flows, first and foremost. You have multiple pools of capital available in the debt market, especially because of our investment-grade status and strong balance sheet. And then you also have the opportunity, from what we've seen, to partner with the whole value chain on the vendor and customer side. There's nothing easy about SpaceX, for sure... we might as well have the finance piece be part of the challenge at the same time. But we have a lot of conviction that we're going to be able to do this in a very measured fashion and continue to keep ourselves in that IG status along the way.

Elon Musk: And certainly, if you look at our track record with respect to capital efficiency with Starship and all of our other projects, you'd be hard-pressed to find any company in the world that is as good at capital efficiency as SpaceX.

Bret Johnsen: Exactly.

Andrea Williams: Next question from Edison Yu at Deutsche Bank.

Edison Yu, Deutsche Bank: I'll echo Adam's enthusiasm about the Roadster. There's been a lot of reporting, a lot of anticipation of a merger with Tesla. I realize maybe you may be limited on what you can say. But if we take a step back, how do you think about the pros and cons of a potential combination? And as part of that framework, are there any roadblocks that you would need to address before embarking on it?

Elon Musk: I mean, this is definitely kind of a question where I feel like I'm tap dancing through a minefield. So, I have to be thoughtful about what I say here.

Both SpaceX and Tesla have formed special committees to discuss this idea at the request of shareholders. These special committees consist of independent directors and because I'm a conflicted party, I can't actually be in the discussions of the special committees. It's sort of an odd kabuki dance where two groups of people are talking, and I can't be in the conversation.

I think some people out there assume that a SpaceX and Tesla merger is a certainty, and I think that would be overstating it. Do I think it would be a good idea? Yes. But at the end of the day, the shareholders have to be supportive of this. And particularly on the Tesla side, we will need a 50% vote of all shares, not just those who vote, but of all shareholders. And that's quite a high hurdle to try to achieve. I can just say at this point that the special committees have not reached agreement on the terms of any transaction. And obviously there's no guarantee that they will.

But I do think there's a lot of merit to combining the companies, as we are pretty much in the super intelligence singularity and there's a lot to be gained from cooperating on developing super intelligent software, chip design, the Terafab, obviously robotics, mass to orbit with the Starlink constellation.

Gwynne Shotwell: Connectivity.

Elon Musk: Connectivity. I think it will be a good idea, but it's really up to the special committees of both companies. They've got to come up with something, and then that's got to be put to the shareholders as a vote.

Andrea Williams: Next question from Kutgun Maral at Evercore.

Kutgun Maral, Evercore: I wanted to ask about Starlink Mobile, which has come a long way in a short time, including the recent low-band deal. Even so, there's still some uncertainty about how your spectrum portfolio, even though it's much improved and possibly will continue to grow, can compete with incumbents in the U.S. And part of the unknown for us on the outside is how much coverage and capacity the satellite layer can really provide to bridge the terrestrial gap to peers. What roles do you expect the satellite and terrestrial parts of the network will play? And how should we think about the traffic mix between them in delivering a 5G-like experience in the U.S.? And maybe more broadly, how much of the connectivity focus is on traditional retail wireless versus what could be a massive physical AI opportunity?

Gwynne Shotwell: Elon can fill in some details, but let me point out that the vast majority of mobile connectivity is actually Wi-Fi, right? I think it's something between 85% and 93% is done through Wi-Fi. And keep in mind, Starlink Broadband does a lot of Wi-Fi work as well. As far as the mix goes, I think we expect that the S-band spectrum will provide the majority. And then obviously you need the low-band spectrum, the 800 megahertz spectrum, to basically penetrate buildings. So, if you're not on Wi-Fi, you can still receive the signal. I think — and by the way, I'll correct myself if I'm wrong afterwards — but I think the S-band will carry the majority, and then the low-band spectrum will carry less, but provide a critical piece.

And also, we talked last quarter, at the last earnings call, about smaller, kind of not quite tower-capable, but smaller radios that we could put on already installed homes for Starlink, businesses for Starlink, towers, as well. It's not just the spectrum that we've talked about with the S-band. We will basically leverage the real estate from Starlink installs and other towers.

Elon Musk: I think we probably should do a longer presentation on the game plan for Starlink connectivity and why we think it will ultimately be a business that generates several hundred billion dollars in revenue a year, because you are missing some of the pieces that we know about internally. I can say we do see a clear path to several hundred billion dollars a year in connectivity revenue. But I think we should probably describe that in more detail in a forum that's not limited to a few minutes at an earnings call.

Andrea Williams: I'll take the next question from Eric Sheridan at Goldman Sachs.

Eric Sheridan, Goldman Sachs: Maybe if I could turn to the scaling of compute you laid out at the beginning of the call. When you think about a framework for applying that compute to the external opportunity that presents itself from cloud services, relative to the internal needs of that compute for scaling Cursor and Grok and Grok Bot, how has your view evolved about striking the right balance in that mix? And how to think through the output, from a monetization standpoint, of what that mix might produce for the business in the next 12 to 18 months?

Elon Musk: Well, the relative allocation of compute to third parties versus internal is something that's going to change from one month to the next. I guess as a rough approximation, our goal is to kind of split it 50-50, more or less, where half of the compute is used internally and half of it is rented to others. But that number could shift 60-40 one way or the other. It's going to be dynamic. And obviously our approach is to rent at effectively the spot-market rate, as opposed to signing these long-term deals.

Our belief is that the value of compute per watt will increase over time, not decrease. And the reason for that, I think, should be obvious to people, really, is that as you improve the AI software, as the models get better and better, you get more useful intelligence per watt. And the value of that compute is proportionate to the useful intelligence per watt. And that's why you see things like even old computers, like H100s, renting for more than they rented when they were new. And I think that's likely to continue. I think we're going to face a shortage of super intelligence compute. And that shortage is going to remain for a long time. I think the smart move is to do short-term contracts and essentially to place the bet on useful intelligence per watt increasing rapidly, which it obviously has done.

So that's what's happened so far. And also, if you look at, say, how are our short-term contracts renewing? Are they renewing at higher dollars per watt or lower? And we see them renewing at higher dollars per watt, not lower dollars per watt. We are also, frankly, hounded relentlessly by every hyperscaler and by dozens of smaller companies to rent more of our compute. We could rent out our compute several times over if we said yes to everyone. That's just the situation here.

That's why I posted that sort of video of me, of the Monty Python cheese shop sketch, where the guy goes into a cheese shop to get some cheese and they actually have no cheese. And so he goes like, what about this kind of cheese? What about that kind of cheese? And they really just don't have any cheese. So we've actually tried to offload some of our inference compute to other cloud providers so that we can use more compute for training. And we get the cheese shop sketch. We'll ask them about, what about some GB300s? Don't have any of those. What about some GB200s? Don't have any of those. What about some TPUs? No. What about some Trains? No. What about some — okay, how about some H100s? And like, well, we might have some overseas somewhere, and then they'll check and come back and say, no, we don't have that either. It's ridiculous.

I don't think we're going to have a demand issue here. And it's not impossible. I'm not saying this will happen. It's not impossible that we will have a six-month payback period for compute next year.

Andrea Williams: We're now going to turn to some questions from individual investors that were posted directly onto our public Q&A platform.

We've talked a lot about compute deployment terrestrially. Let's move to a question on V3 sats: when will capacity gains from V3 satellites actually show up as meaningfully higher speeds and lower latency for users? How many V3 sats are needed before the network feels noticeably better in high-demand areas?

Gwynne Shotwell: I think we'll need a few hundred V3 satellites for people to really feel the impact, because right now, with very few satellites, you'll get an impact and then 10 minutes later it won't be there. So, a few hundred satellites, and I think we can achieve that by mid next year.

Andrea Williams: Can you provide an update on Starbase Louisiana? What's the timeline? And what's your current confidence level in achieving a first Starship launch from the site as early as 2029?

Elon Musk: Well, our internal aspiration for first launch from Louisiana is actually in two years. So '28 would be our goal. But I think '29 is highly likely. I think we'll be launching from Louisiana at scale in '29. Highly confident.

Andrea Williams: Thanks, everyone. This wraps up the Q3 2026 earnings report. Thank you all for tuning in. And we will see you next quarter.