

Edited Transcript

SpaceX Second Quarter 2026 Webcast August 4, 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 second quarter 2026 Earnings Conference Call. I'm Andrea Williams, Head of Investor Relations. Joining me today are Elon, Gwynne, and Bret. We are speaking to you from our facility in Bastrop, Texas.

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

For this webcast, 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. Thank you, to everyone who submitted your questions in advance.

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 most recent filings with the SEC including our final prospectus and today's published Form 10-Q.

Additionally, during this call we may discuss certain non-GAAP measures. In our press release and our filings with the SEC, you can find additional disclosures including all reconciliations with comparable GAAP measures.

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

- Revenues of \$7.8 billion, up 92% from \$4.1 billion
- A quarterly net loss of \$541 million, an improvement of \$467 million
- And Adjusted EBITDA of \$3.5 billion, up 191% from \$1.2 billion.

Now, I will hand the call over to Elon for his opening remarks.

Elon Musk: Thank you. So, it's been another milestone year for SpaceX so far. We're making great progress developing a fully and rapidly reusable, reliable rocket with Starship.

We're connecting millions more people globally across 170 markets with fast, reliable Starlink Internet, we're building AI compute capacity at scale faster than anyone else, we believe.

And we're significantly improving our AI models, most notably with the release of Grok 4.5 last month. And now, I'll dive into each of those topics separately.

So with Starship, we completed two successful flights of Starship V3 in the past 90 days. Flight 13 demonstrated core capabilities necessary to achieve an orbital mission and return to Starbase for catch. On Flight 14 will be our first flight to fly our Version 3 Starlink satellites, our communication satellites, to operational orbit.

We'll try to actually catch the first and the second stage of Starship V3 this year, and we could possibly catch the ship as soon as the next flight. And it's difficult to actually explain to most people, just the incredible significance of Starship.

I think perhaps describing it in terms of tonnage to orbit is a good way to appreciate the magnitude of the impact Starship will have on the future of civilization. So right now, we deliver, call it, roughly 2,500 tons a year to orbit via Falcon. This is a big number by normal standards. The rest of the world, I think delivers around 300 tons. So, we deliver 80% to 90% of total earth mass to orbit per year with Falcon. Now with Starship, our aspirations, and I think we will achieve these aspirations, are to deliver well over 1 million tons to orbit per year and probably ultimately 10 million tons per year.

If you had a chart where you plotted things, you would need a very, very big table to put the chart on to even see any of our competitors on that chart and to make it so our next biggest competitor was one pixel high on the bar chart. You would still need a very big bar chart to represent the SpaceX Starship vehicle.

And it's worth noting that we actually encourage our competitors to copy us. We do nothing to slow them down. We actually launched competing satellite constellations at fair prices, the same price we charge everyone. So, we do encourage our competitors to copy us, but it's quite difficult to copy something like Starship. But we don't use patents as sort of a landmine exercise to stop people from copying us. We actively encourage it. When you think about it, we go from 2,500 tons a year to 1 million, potentially 10 million tons per year. It's really a ridiculously profound difference. But that is our plan, and I think we will achieve that plan.

Regarding Starlink, it's also worth just noting that the Starlink V3 satellite is about an order of magnitude more capable than the Starlink V2 satellite which was already a very capable satellite. We expect to launch about an order of magnitude more Starlink V3 satellites. That would mean roughly two orders of magnitude increase in the delivered bandwidth. Even if our monetization per bit dropped by a factor of 10, that would still mean a 10x increase in the revenue of Starlink. I think people are really underestimating Starlink here. This is a big deal. It's kind of hard for people to wrap their minds around this, but it's not out of the question that at some point, Starlink will deliver a majority of the world's Internet, at least in countries where we're allowed to operate, which is the vast majority of countries. This is important to bear in mind, and it's not in the infinity future. It's less than 10 years.

With the advent of AI and humanoid robotics and vehicle robotics and just a massive number of robots, the appetite of bandwidth will be much greater than it has been in the past. You think of a human consuming perhaps, or producing perhaps, a few hundred bits per second of output. Actually, the average human is outputting less than one bit per second over the course of a day which is 86,400 seconds. Now computers easily do billions of bits per second continuously over a 24-hour period. So, this is why I would expect the appetite for bandwidth to grow dramatically with the growth of AI and robotics. And really, I think Starlink is the only internet that can actually service that bandwidth.

Regarding AI, we're making rapid progress on Grok. Grok 4.5 was a huge improvement. We have Grok 4.6 coming out probably next week. Then, Grok 4.7 is about three or four weeks from today. We expect the cadence of AI development to improve dramatically. With Grok 5, which should be out before the end of this year, we will be incorporating the entire corpus of SpaceX data. Basically, all the data that SpaceX has ever produced, which is a tremendous amount over the course of a quarter century, will be incorporated into Grok training. We think this will make Grok by far the best engineer. These are all like very big things.

In addition, of course we are providing compute to others, and we are building and deploying compute, I think faster, or our rate of growth certainly is faster than anyone else. Our efficiency of compute deployment, I think is also the highest. We expect to end this year with over 2 gigawatts of compute. And probably our cumulative compute online by the end of next year will be several times higher. It may, let's say, be closer to 10 gigawatts of compute than 5 gigawatts of compute. That's a way to think about it. And going forward, we've decided to build exclusively on NVIDIA because we think the Vera Rubin architecture is the best architecture. We think it's the best AI computer, and we greatly value our close cooperation and partnership on many levels with NVIDIA. So, we're exclusive to NVIDIA.

This new capacity coming online will obviously further benefit the development of Grok, as well as our Cursor acquisition, which should close quite soon. I think we're through almost all of the regulatory hurdles.

Then with respect to the Starlink AI satellite, which will be essentially an optimized Vera Rubin NVL72 computer, this is not some sort of far-future, distant thing; we expect to start launching these next year. We think the design of the NVL72 VR computer is a much better design than, say, having a standard rack-style design, and so we expect to actually deploy this on the ground, as well as in orbit because we think it's going to be a radical simplification of the normal NVL72 rack. It will cost less. It will be more effective. If we're going to put it in space, why not want to put it on the ground? I think that's going to be pretty cool.

In conclusion, the SpaceX team is solving some of the hardest engineering problems in the history of humanity. I think the team is succeeding because I think we've got in a lot of ways, the smartest, most dedicated team of humans that has ever existed. I'm incredibly proud to work with such a team and I'd like to thank the team for their incredible hard work. And also, thank you to investors who have supported us along the way.

Gwynne Shotwell: Thanks, Elon.

We had an exceptional second quarter. Today I'll focus on the operational and commercial highlights as well as the initiatives we prioritized going forward.

Starting with our Space business. We remain the leading launch provider globally, with 78 total launches, and 1,041 tons of mass to orbit delivered in the first half of this year, primarily allocated to our own internal Starlink missions, but with great customer missions as well.

While our launch capabilities are key to supporting our Connectivity business, we continue to have a very robust commercial manifest and a very robust and growing U.S. government mission demand which drives revenue for the segment. We're currently launching at our highest Falcon cadence and are on the precipice of operationalizing Starship.

As Elon said, the first two flights of our V3 Starship vehicles put us in an excellent position to achieve our near-term goals of reaching orbit, catching and reusing the ships and boosters, and deploying our V3 Starlink broadband satellites.

Speaking of which, Starlink had a standout quarter, driven by continued growth across our consumer, enterprise, and government businesses. During the second quarter, we added more than 1.7 million Starlink subscribers globally, that's consumer. This reflects our best quarter of new customers to date, an increase from the 1.4 million we added during the first quarter. We achieved this while holding ARPU, which is the average revenue per user, stable at \$66 per month compared to the first quarter. We ended the second quarter with service availability in 167 markets and have continued to activate new markets since then with some very exciting announcements coming soon, maybe even as early as today.

As of June 30, our constellation of operational Starlink broadband and mobile satellites in orbit grew to roughly 10,200 with our 9,600 broadband satellites delivering roughly 800 terabits per second of total downlink capacity. Following a successful speed run test of the V3 Starlink satellites on the most recent Starship launch, during which we connected every satellite through every laser link, we intend to deploy V3 satellites into the constellation for operational use on the upcoming Starship missions. The significant amount of capacity we're able to add to the Starlink constellation from the V3 satellites will enable us to continue providing even better service. I think it's pretty great already – but to do so while serving more and more customers over the world. In fact, in the years ahead, we expect Starlink will represent a significant portion of global Internet traffic, which Elon also talked about.

We also see significant growth opportunities for Starlink on the enterprise and government side of the business. We believe revenue from these markets have the potential to reach a scale at least comparable, but likely exceed our consumer business as we continue to gain share. In the second quarter, we signed a major agreement with American Airlines as well as activated our service with new partners including Southwest, Virgin Atlantic, Iberia, and Aer Lingus. We have become the provider of choice among major airlines with outstanding customer feedback, but still have significant room for growth. In fact, we heard from one of our airline customers, shockingly, customers are flying shorter hop flights instead of direct so that they could ensure that they're on a Starlink-activated flight. They've never seen this in the business before.

On the government side, we won more than \$6 billion in U.S. contracts in Q2, supporting major Space Force programs that offer our nation mission-critical communications and sensing capabilities, and we see even more room for growth in this sector in this coming year.

In Q2, we also launched the new Starlink mobile partnerships with international carriers including SoftBank, NTT Docomo, and Spark New Zealand. Our near-term priority is launching mobile V2 satellites on Starship ahead of integrating the 65 MHz of EchoStar spectrum later next year. This spectrum transfer was recently approved by the FCC and represents a foundational competitive advantage for Starlink mobile.

In AI, the industry is hurdling towards ever more compute demands. We're rapidly expanding our compute capacity to meet our own needs as Grok expands, as well as those for other leading companies such as Google and Anthropic.

In terms of our own Grok models, the enterprise feedback on rollout of Grok 4.5 has been very positive so far. Token consumption tripled out of the gate after the July release and we continue to see strong signals on usage and monetization for the model. We're looking forward to welcoming the Cursor team to SpaceX to integrate our engineering teams and begin to benefit from a combined sales capability. In review, our momentum as a business is extraordinary.

Our engineering teams are scaling Starship and our next-gen satellites while our sales teams expand Starlink, aviation mobility, and global markets. At the same time, we're bringing on additional compute capacity and expanding and deepening our AI partnerships across a broader set of customers.

While I've been in this business for 24 years, I could not be more excited about the months ahead with all the exciting and amazing things going on. It really feels like we're just getting started all over again. Now, I'll turn the call over to Bret.

Bret Johnsen: Thank you, Gwynne.

We demonstrated strong results in Q2. Revenue growth accelerated across every segment. We narrowed our net losses by nearly half and on a year-over-year basis, and we grew adjusted EBITDA meaningfully faster than revenue. Importantly, we significantly strengthened the balance sheet through our IPO and inaugural investment-grade bond offering. The combination of our global launch leadership, recurring Starlink revenue and rapid growth in AI infrastructure continues to drive improved operating leverage. Now I'll walk through the financials for each segment.

For the Space segment, revenue grew 55% sequentially and 29% year-over-year to \$962 million, driven by a higher number of larger customer launches and a favorable customer mix. Segment costs and expenses rose \$389 million year-over-year, and we continued to accelerate R&D investments in our Starship program. Starship aims to quadruple payload capacity and reduce launch costs by 10 times compared to our Falcon 9 rocket, unlocking significant capabilities across all of our business segments. We continue to make progress building the infrastructure required to support thousands of Starship launches per year including accelerating Raptor and launch vehicle production, the build-out of our Gigabays and making significant progress towards activating multiple launch pads at Starbase and in Cape Canaveral at Pad 39A and Pad 37. Space Adjusted EBITDA was a loss of \$205 million for the quarter, primarily reflecting these higher R&D expenses.

Moving to the Connectivity segment. Revenue was up in the Connectivity segment was up to \$4.3 billion, up 32% sequentially and 66% year-over-year. Growth was driven by record net additions of Starlink subscribers and a sharp increase in enterprise and government revenue.

Starlink subscriber ARPU was \$66, unchanged compared to Q1. As we continue our global expansion, we're taking a localized go-to-market approach that ensures our product and service offerings fit local needs. While our geographic expansion may drive down blended ARPU over time, we anticipate continued subscriber momentum to deliver strong revenue growth. Enterprise and government revenue grew 108% year-over-year and represents a durable source of revenue, contributing to strong segment margins overall.

As Gwynne mentioned, we see enormous opportunities for growth in this segment as we remain, for example, less than 10% penetrated in one of the largest segments being aviation.

Total costs and expenses for the Connectivity segment increased \$970 million, or approximately 58% year-over-year. The increase was driven by higher spend to support our revenue growth, including growth in our satellite constellation, increased R&D investment for our next-gen V3 satellites, and incremental marketing spend. As Elon mentioned, our V3 satellites are expected to deliver 10 times the amount of broadband capacity and data density from our current generations, which I think we believe are already pretty amazing, and which are essential to supporting continued subscriber growth and higher value service tiers.

Income from operations from the Connectivity segment increased 79% year-over-year to \$1.7 billion, outpacing revenue growth, and driving nearly three points of operating margin expansion. Connectivity segment Adjusted EBITDA rose 64% year-over-year to \$2.6 billion.

AI segment revenue was \$2.6 billion, up 213% sequentially and 247% year-over-year. The increase was driven primarily by new cloud services agreements, as well as growth in Grok and X subscription revenue. Advertising revenue grew 7% sequentially, as we overhauled our advertising platform technology to enable AI functionality and enhanced self-service for all customer types.

In Q2, we entered into Cloud Services Agreements providing access to compute capacity at our Colossus and Colossus II sites. The initial ramp from these agreements contributed \$1.6 billion of incremental AI infrastructure revenue in the quarter.

Total cost and expenses for the AI segment rose \$1.6 billion year-over-year, reflecting increased R&D primarily due to higher infrastructure spend as we accelerated compute deployment. We meaningfully narrowed our AI segment net operating loss to \$1.3 billion and turned Adjusted EBITDA positive for the segment in the second quarter at \$1.1 billion. The incremental revenue from new hosting deals generated high incremental EBITDA margins as we monetized available compute capacity.

We ended the second quarter with 1.4 GW of nameplate compute, up from 1 GW in Q1 and 400 MW a year earlier. We expect to end this year at over 2 GW of compute capacity.

Total company capital expenditures in the second quarter were approximately \$18.4 billion, of which roughly \$15.8 billion supported AI compute infrastructure. The remainder funded ongoing Starship and launch infrastructure, satellite production, and global ground station expansion. Additionally, we paid out \$856 million during the quarter under our Spectrum Credit Agreement related to our pending EchoStar transaction.

Our capital allocation priorities remain clear and disciplined. We plan to continue to invest aggressively in the three areas we believe generate the highest long-term returns: Starship development and production scale, next-generation Starlink broadband and mobile constellations, and AI compute infrastructure. Across all three of these investments, we remain focused on capital efficiency, return on investment, and preservation of a strong investment-grade balance sheet.

The second quarter was transformative from a capital markets perspective. We completed our initial public offering, raising approximately \$85.7 billion in net proceeds, and followed that with a \$25 billion inaugural investment-grade senior notes offering that was partially used to repay our \$20 billion bridge loan. The senior notes were issued across five tranches with a weighted average interest rate of 5.855%, and average maturity of 11.7 years.

We ended the quarter with \$100 billion of cash, cash equivalents and marketable securities and \$47.5 billion in backlog.

Looking ahead, we continue to see robust demand in all three of our business segments, but in particular, in our cloud services arrangements. We see increasingly favorable economics with each agreement we signed. And as Elon mentioned, we expect the supply-demand imbalance in the compute market to continue. The current economics have translated into a less than one year payback on our new capital deployments for compute. For example, in the first few weeks of the third quarter, we've already contracted an additional \$6.7 billion of Cloud Services revenue over a six-month period that begins ramping starting in October of this year. We believe this puts us on a trajectory, including contribution from Cursor, to reach \$100 billion of ARR, or annualized revenue run rate, by the end of this year, based on our expected revenue in the month of December of this year.

With that, I'll turn the call back to Andrea.

Question and Answer Section

Andrea Williams: Thank you, Elon, Gwynne, and Bret. This concludes our prepared remarks. For today's Q&A session, we will take our first questions from the equity analysts.

Eric Sheridan, Goldman Sachs: Thanks so much for taking the question. Thanks for all the details in the prepared remarks. Maybe just one on the Connectivity segment. Can you discuss the backlog of enterprise and government contracts in that segment of the business and how we think about that backlog, converting it to revenue, and what that might mean for the pace and cadence of growth in the next couple of years?

Gwynne Shotwell: Yes. We are quite bullish on our enterprise activities. I mentioned a loan that we received from the government, over \$6 billion in contracts. You can look at those as tranches and beginnings of additional capability that we will bring online, some of which we will have to compete. But fundamentally, we have the contracts in place to fulfill many phenomenologies, both the ones we've recently announced as well as additional ones that we'll be working on going forward.

Enterprise revenue is quite sticky. We have never lost an enterprise customer. They're quite happy with the Starlink capability that we have right now. Bret mentioned we are only 10% penetrated in the aviation industry, which definitely represents a lot of headroom for revenue going forward.

Maritime – Starlink, I think, opens up the TAM dramatically. Starlink offers the capability that ships don't really have much connectivity right now because they can't afford VSATs, we can basically outfit those ships with Starlink Maritime.

So, I'm quite bullish on enterprise, very bullish on government and quite happy with the stickiness that we've seen in those markets.

Elon Musk: Yes. I'll just elaborate a little bit on that, which is that obviously in order to secure a large amount of enterprise revenue, you have to demonstrate a very high reliability, very high uptime.

In the early days, Starlink did not have high uptime. There's a lot of enterprise customers, most enterprise customers, that still have that maybe early Starlink experience, where it was a little patchy. We're building up a large enterprise sales team just to go out there, meet with enterprise, basically corporate and government customers, show them the incredible uptime and low latency of the Starlink system today. They realize that they can actually treat Starlink as a primary provider, not sort of as a backup provider. This requires just going out there, talking to people, educating them about the current state of the Starlink system and where it's going. I would expect enterprise revenue to substantially exceed consumer revenue.

Adam Jonas, Morgan Stanley: Elon, first one for you. Just confirming, you said you expect to be closer to 10 GW of nameplate compute by end of '27, closer to 10 GW than 5 GW. Do you have line of sight? How confident is your line of sight on permitting, chips and turbines for that? Then I have a follow-up for Gwynne.

Elon Musk: So, we're actually aiming to far exceed that gigawatt number in terms of power online, power cooling and electrical equipment. Our tentative target is to actually have 20 GW at the power and cooling level online by the end of next year. Now I don't think we're going to achieve 20 GW, but we want to have a series of projects that cumulatively come to 20 GW by the end of next year. Some of them won't pan out exactly on time, but I would expect that we still probably have at the power plant level, something close to 15 GW. Assuming that maybe a quarter of the projects take longer than expected.

Our goal is to have far more power, cooling and electrical equipment than we have GPUs. That's the logical thing to do given the relative expense of GPUs versus balance of system. SpaceX as a company is incredibly good at hardware. In fact, I'd say probably outside of China, SpaceX and Tesla are the two best companies on earth at hardware. I think perhaps even including China. We're taking a small amount of the expertise that we use for rockets and satellites and applying that to scaling terrestrial data centers. We're finding that even a small amount of what we've learned building rockets, which are incredibly difficult, applied to data centers, yields tremendous benefits. We expect to have far in excess of the power cooling that's needed. Then our understanding with NVIDIA is that we will receive a very significant percent of their GPUs next year.

Adam Jonas, Morgan Stanley: Gwynne, just a follow-up. You were recently asked what is one of the more underappreciated parts of the SpaceX investment story and you said robotics. Can you elaborate on why you believe that? How do robots fit into the SpaceX TAM?

Gwynne Shotwell: Let me try to put something together for you. I think what I might have said is underappreciating the amount of connectivity that we will need going forward. I think robotics, especially robotics tied to AI, like for autos and humanoid robots, will dramatically increase the demand for connectivity. Dramatically. Elon talked about that in his remarks.

Elon Musk: I mean you can think about satellites as being robots. I mean they're autonomous, and they do not need servicing. In what may seem as the far future, but will come faster than you think, using robots on the Moon to scale up manufacturing on the Moon, which I know sounds like super sci-fi right now, but it's going to happen, will enable us to build the mass accelerator on the Moon. And if you have a mass accelerator on the Moon and you have solar and radiator production on the Moon, you can – I know this sounds totally nuts – but you can probably scale to 1,000x the economy of Earth in terms of intelligence launched to space. Probably maybe even one million times. So, we are going to land a lot of tonnage on the Moon. We're going to build factories on the Moon. The robots will be helpful with that. Then it's like, well, how much intelligence would you like as we grow towards being a Kardashev 2 scale civilization?

John Godyn, Citigroup: One of the hallmarks of SpaceX to date is how efficiently the organization has used capital so far in its journey. But it's also no secret that the capital needs are significant to support your vision from here. Obviously, the company is moving fast. So, I was hoping you could talk a bit more about the Capex outlook across the segments, touch upon any areas where buying versus building might accelerate growth. The Mobile strategy obviously comes to mind there. And any related thoughts on capital needs for next steps in extreme vertical integration like Terafab, anything that's needed to support future growth.

Bret Johnsen: I appreciate the question. I think we are very proud of how we have managed capital allocation over the years, for sure. I think we are continuing down that path, right? If you look at it, we have been very efficient to date, and I think we'll continue to be. If you look at the rest of this year, to answer specifically on Capex, I think you should probably think that the next two quarters are very similar to the current quarter from a Capex level perspective.

I do think as you look longer term, I think Capex, really, I think from our perspective all Capex is not the same, right?

What we're seeing right now is a situation where we're able to deploy – specifically on the AI compute side – capital in such a way that we're getting less than a one-year payback. It's almost like your capital is becoming an item that you would put otherwise in COGS because it's moving so quickly as far as monetization and moving into revenue.

That's different than, I would say historically on the SpaceX level, where we're building assets that are going to benefit us from launch sites for the next couple of decades. So, we certainly manage for optimizing ROI on our capital all the time from that perspective.

Elon Musk: Yes. I'll add a little bit to that. As Bret mentioned, we are expecting to reach a \$100 billion-plus ARR in December of this year. It's probably also worth mentioning that our internal projections for reaching \$1 trillion in revenue, not ARR, but revenue, have moved up from 2031 to 2030.

So, prior to the IPO, the financial projections we had were reaching \$1 trillion in revenue in 2031. We now expect that to be in 2030. There's a non-zero chance of that being in 2029.

Ronald Epstein, Bank of America: Could you walk through what you learned on Starship launch 13? What was better than expected? What was worse? Then how do we think about this in the journey to rapid reusability, the heatshield and so on and so forth?

Elon Musk: I was a little worried about that Flight 13 because think about Apollo 13; things went wrong on that flight. But this flight actually went incredibly well. The heatshield appears to be very robust. The ship is still floating in the ocean, and we expect to recover the ship and bring it back for analysis. So, I would say things look very good.

That's why, assuming we receive regulatory approval to do so, we'll attempt to catch the Ship with the tower on the next flight, which is tentatively scheduled for the end of this month.

We expect the cadence of flights to be increasing rapidly, and probably a year from now we will be doing at least one flight a day, possibly more.

Doug Anmuth, JP Morgan: Elon, can you elaborate on the key unlocks from leveraging your rockets and satellite expertise just in accelerating the pace of compute build-out? Then, totally appreciate the \$6.7 billion in incremental compute deals – can you talk more about just how you expect the supply-demand environment to persist going forward and whether you'll be able to keep this premium pricing for these deals and sustainability?

Elon Musk: Sure. Well, as the saying goes, normally they would say: "this ain't rocket science". Well, actually, rocket science is literally our daily business. And rocket science is an idiomatic expression for extreme technological difficulty and there's a reason for it because, the thing is, let me tell you what rockets desperately want to do every flight... they desperately want to blow themselves into tiny pieces. Then the engineering struggle is to convince the rocket not to blow itself into tiny pieces and actually deliver payload to orbit. But you'll note that some of our competitors do, in fact, blow their rockets up. This does not mean they're incapable. In fact, I think our competitors are very smart people, but despite being very smart, their rockets still blow up.

So now if you apply a small amount of the engineering skill that goes into making gigantic rockets reach orbit on a regular basis reliably to terrestrial data centers, this is like, frankly, the New York Yankees going in and playing a little league team. It's ridiculous, frankly. Terrestrial data centers are a trivial problem compared to making gigantic reusable rockets which are launched frequently. And the Starlink satellites also are alien-level technology. There's still nothing that compares to our Starlink V2 satellites, and our V3 satellites are really more than an order of magnitude better than our V2 satellites. So, you take a small percentage of that engineering, apply it to terrestrial data centers and you get an amazing outcome. So that's what we're seeing.

Now from a monetization standpoint, if you say: "what's my best guess for monetization per watt of, say Rubins?" My guess is it would be somewhere between \$30 and \$50. It's just a guess. Take a guess as a guess. But what we're seeing with AI is that the intelligence per watt is increasing rapidly. That means the usefulness per watt is increasing rapidly. I think anyone who's a close follower of the AI arena, you start seeing smaller and smaller models do more and more amazing things. There's an

argument that we may get to the point where even an H-100 is able to deliver useful AI that's comparable to a human engineer. Based on current trends, I think that's not an impossible outcome. It's like how useful would that be? I think the usefulness is going to be extremely high.

I mean you think of where was AI 12 months ago? 12 months ago would have been August last year. That would have been before the advent of Claude 4.5, which I think should be considered one of the milestones. And a credit to Anthropic, it was, I would say, a shock to the system how good Claude 4.5 was, and that was 10 months ago basically. And if you say what about an AI two years ago? I think well, an AI two years ago feels like something that should be in a museum, okay? Where you're like, you're not going to believe this, but two years ago, this was AI, this relic. I think by the end of next year, it's not clear to me that there's anything that, digital at least, that AI won't be able to do based on the current rate of improvement.

Then look at the rate at which logic and memory is being produced. One must always consider a limiting factor here. Limiting factor currently is memory. The memory output is increasing by around 20% per year. Now normally, that would be fantastically fast and amazing for any large, mature industry. But ask yourself, is the demand increasing by 20% a year? No. The demand is increasing by 200% a year, maybe higher. So, if you've got demand increasing much faster than supply, then Economics 101 would suggest that the price increases; it does not decrease.

Edison Yu, Deutsche Bank: I wanted to come back on the comment about the \$100 billion in ARR. Is there any way you can unpack that a little bit and perhaps what it takes to get there? Then in relation to that, Elon, I know you mentioned you pulled the revenue target forward by a year. Can you maybe call out maybe a few things that gave you this extra confidence?

Bret Johnsen: Well on the unpacking, I think the key there is, it's really growth in all three of the business segments. But the biggest piece is – and we already talked about additional business we've already closed in the first couple of weeks of July on the AI cloud services side – but certainly, as you know, we've already announced some previous deals with Google and Anthropic that start to ramp either later this quarter or in October. So, there's a lot of momentum from that side of the business as well as we're seeing strength in our own models, and Cursor coming into the fold as well. It's really all three segments that are driving that, but the Cloud Services piece being the largest contributor of that piece.

Elon Musk: To be clear, the \$100 billion ARR in December is not a question mark. That's what we'd achieve if we basically did nothing. So, I think it may be higher than that. It probably will be higher than that. In terms of the relative percentage of breakdown of what revenue generates what, I think the breakdown is probably similar to what we shared pre-IPO. It's just happening a year sooner.

And I will go back to what I said and what Gwynne and Bret mentioned earlier, which is that we do expect a very rapid increase in revenue from Starlink communications from the V3 satellite. Because, like I said, it's more than 10x the capability of the V2 satellite and we'll be launching 10x as many, and we have the direct-to-sell constellation going up as well. So, if you were looking at the historical growth of Starlink based on V2, but we have a step change, a massive step change happening with the V3 satellite.

John Hodulik, UBS: Great. Two things. First, Elon, just following up on that comment on Starlink Mobile. When do you expect to start launching the V2s? Can you talk about your appetite for additional spectrum after the EchoStar deal? Then just how big do you think the opportunity is in Mobile? And if you could compare that to the Broadband business, that would be great. Then secondly, just a quick follow-up on the compute. It looks like you're going to have a big ramp in compute coming in 2027. How should we think of that capacity being leased versus used to train Grok? I guess in other words, how much of that has to be withheld to train Grok versus that can be leased out to third parties?

Gwynne Shotwell: I'll go ahead and take the Starlink Mobile questions, John. We're going to start flying the next-generation Starlink Mobile satellite, which is confusing. It's actually the second version as opposed to the V3, which is the next-generation Broadband satellite. We're going to start flying those next year.

And just so you understand currently the system that we're operating uses about 5 MHz of bandwidth, and we operate that through local telco providers. We have 65 MHz of bandwidth available to us through the EchoStar spectrum, which gives massive increase in capability. Another way to look at it is we were also probably 10x the number of satellites.

Simplifying, you could look at the Starlink Mobile leveraging the next-generation satellite and the EchoStar spectrum as being 100x better, 10 times better than what we've got right now. It's a pretty decent service. You can do voice and video calls over signal and WhatsApp. So, we anticipate really strong uptake of that capability.

As far as revenue, I'm just going to take a very high-level top-down. The Big Three in the United States, AT&T, Verizon and T-Mobile, roughly between them \$600 billion a year, and I anticipate us to be able to acquire quite a few of their customers because I think our service will be better. We will eliminate dead zones, leveraging the satellites in orbit. It will be better during any sort of natural disaster because, surprisingly, even though space movies make space look super dangerous, it is a pretty quiescent environment.

I'm quite excited about Starlink Mobile going forward. We will start to fly the satellites next year, and we will start providing service end of next year.

Elon Musk: We obviously have more than one question being asked here. Obviously if somebody is trying to create a financial model, some of these questions are... you have to plug in some number here. For Grok training, we do expect the percentage of compute that we use internally for Grok training to decrease over time and to have far more of the compute used for inference, or for renting training and inference compute to others. So, I would expect maybe 10% of our compute to be used for Grok training, something like that.

Kutgun Maral, Evercore ISI: I had a follow-up on Starlink Mobile and your ambitions in the U.S. market in particular. I think a lot of us are trying to better understand what a U.S. direct-to-consumer service could look like. To become a true fourth U.S. carrier, would imply well over \$100 billion of mobile Capex over time, much of auction spectrum and greenfield infrastructure. So how are you weighing that potential build plan against the MVNO or acquisition path? And what's your internal base case on how far you'll take this ambition? Or is this all a little bit too premature at this point?

Gwynne Shotwell: The spectrum that we purchased from EchoStar does have terrestrial components. So, we definitely intend to build out the terrestrial component. You will have not only the capacity from the satellites themselves, but you will have a build-out of the terrestrial; basically the hardware and systems necessary to make a true mobile service exactly what you want it to be.

I'm not going to talk about the Capex for that. It's a little, well, I don't want to say it's immature. I'm just going to say I'm not going to talk about the Capex for that right now. But let me give a little hint, and that is you could put a base station – a cellular base station on the gear that holds a Starlink broadband dish. You can have these little Femtocells around the country, and you deploy that as you need it. You don't have to necessarily spend many billions of dollars in low-band spectrum all upfront before you deploy your system. That's why I don't want to talk about the Capex for building out the terrestrial because we have a lot of really great and new ideas for how we're going to do that, and I think it will be quite Capex efficient.

Elon Musk: That point is probably worth just a few more words, because this is quite an important point. Instead of having to deploy these very expensive and difficult-to-locate, large cellular base stations, we feel reasonably confident that we can deploy a large number of small stations. Essentially, they're really just Starlink dishes, that also provide connectivity in the mobile spectrum bands and have them be all over the place. The Starlink antennas are located on the roofs of houses and businesses. So, they've got a clear view for providing connectivity directly to cell phones on the ground. In fact, providing connectivity that is, we think, probably better and higher bandwidth than what is currently available from cellular providers. This is an important point.

Andrea Williams: Thank you. With the remaining time we would like to cover some of the questions from individual shareholders. I will read them aloud.

We got a lot of questions on the progress for Starship Human Landing system. Can we provide any updates?

Elon Musk: Obviously we need to make Starship super reliable for satellites before we put any people onboard. But we are expecting to have a very high launch rate, and thus achieve a level of safety needed for humans very quickly, probably by the end of next year.

Gwynne Shotwell: As far as the milestones go, propel and transfer in orbit is critical to both our internal SpaceX as well as our HLS ambitions. Artemis III mission is next year. That's where we will dock with the Orion spacecraft. We will follow that up with a direct-to-lunar cargo mission on Orion, and then we want to put boots on the ground, boots on the Moon in 2028.

Andrea Williams: For the Cursor acquisition, there were several questions asking about the combined product roadmap for the companies once it closes.

Elon Musk: We'll probably make some announcements or discussion of that outside of the earnings call. We don't want to jump the gun on regulators for the closing of the acquisition. We're trying to get the acquisition closed as quickly as possible. I think we're pretty close to that but we have to be wary of jumping the gun on regulatory closures.

Andrea Williams: Next question from the individual shareholders: once you start launching Broadband V3 satellites, how quickly will they translate into improved Starlink service?

Elon Musk: We need some critical mass of V3 satellites, probably at least on the order of about 1,000. Approximately second quarter next year would be, I think, where we'd get to that point.

Andrea Williams: And our last question, is the heatshield the hardest remaining problem for Starship?

Elon Musk: Look, I don't want to jinx it or anything, but I think I would consider the heatshield problem solved at this point. We want to take a close look at the ship that is currently floating in the ocean. But all indications from data and our visual inspection suggests that we have solved the heatshield problem, which is arguably the single biggest problem. That doesn't mean we won't make improvements to the heatshield. Of course, we'll continue to make improvements to the heatshield, but I would say that we do not see any technical obstacles at this point to achieving full and rapid reusability.

Andrea Williams: Thank you. This concludes today's call. You can reach us at ir.spaceax.com and we will see you next quarter. Cheers.

####