NASDAQ:INOD Q3 2026 Earnings Call Transcript

Generated on 11/25/2026

Operator | Conference Call Operator:

Hello everyone. Thank you for your patience and thank you for joining us. Welcome to the InnoData Q3 2026 earnings call. After today's prepared remarks, we will host a question and answer session. If you would like to ask a question, please press star 1 to raise your hand. To withdraw your question, press star 1 again. I will now hand the conference over to Amy Agress, General Counsel. Amy, please go ahead.

Amy Agress | General Counsel:

Thank you, operator. Good afternoon, everyone. Thank you for joining us today. Our speakers today are Rahul Singhal, President and Chief Executive Officer of InnoData, and Jayant Chauhan, Chief Financial Officer. Also on the call today is Anish Pantarkar, Senior Vice President, Finance and Corporate Development. We'll hear from Rahul first who will provide perspective about the business and then Jayant will provide a review of our results for the third quarter. We'll then take questions from analysts. Before we get started, I'd like to remind everyone that during this call, we will be making forward-looking statements, which are predictions, projections, and other statements about future events. These statements reflect management's views as of today and are based on current expectations, assumptions, and estimates as of today, and as such, the statements are subject to risks and uncertainties. Actual results could differ materially from those contemplated by these forward-looking statements. These statements include but are not limited to statements regarding our future financial and operational performance, business strategies, Mishra Mishra. and many more. Thank you. Thank you. In addition, during this call, we may discuss certain non-GAAP financial measures. In our earnings release filed with the SEC today, as well as in our other SEC filings, which are posted on our website, you will find additional disclosures regarding these non-GAAP financial measures, including reconciliations of these measures with comparable GAAP measures. Thank you. I will now turn the call over to Rahul.

Rahul Singhal | President and Chief Executive Officer:

Thank you, Amy, and good afternoon, everyone. Before I turn to the quarter, I want to take a moment to acknowledge the occasion. This is my first earnings call as InnoData CEO. I'm honored by the confidence that Jack and the board have placed in me, and I'm mindful of what I have inherited, a company with extraordinary momentum, an extraordinary team, and an extraordinary opportunity in front of it. I could not be more excited to discuss our results and strategy. I also want to acknowledge Jack. As executive chairman, he remains deeply engaged, partnering with me and our research team to build differentiated capabilities and to bring them to the federal government and the enterprise. Our strategy is shared, and the research-led engine that has been driving our results is the one I intend to build on. I said last quarter that I would repay the confidence placed in me with results, and I'm pleased to be starting with a quarter that lets me do exactly that. Q3 was another record revenue quarter for InnoData. Revenue was $95.6 million, up 53% year-over-year, exceeding analyst consensus of $83.3 million by approximately $12.3 million, or 15%. This was a 13th consecutive quarter of year-over-year growth. We exceeded analyst consensus on other key metrics as well. Adjusted gross margin was 40%, in line with the publicly stated 40% target. Adjusted EBITDA was $17.7 million, or 19% of revenue, exceeding consensus by approximately $2.3 million. Net income was $10.3 million, up 24% year-over-year, and fully diluted earnings per share were $0.29 compared with consensus of $0.17. Much as we hoped, Q3 and Q4 to date have been strong in new business wins. The customer we referred to last quarter as our second largest customer became our largest customer in Q3 and

continues to expand with us. In the quarter, we signed a new program with this customer for pre-training data with a contract value of $29 million. So far in Q4, we have signed an additional program with this customer valued at approximately $18 million, and we have a robust pipeline with them going into 2027. We also have won five new programs at another customer, a large hyperscaler. Just to give you a sense for the momentum and success we have already achieved with this customer, in the trailing 12 months, we did approximately $3.1 million of revenue with this customer. But based on these new programs, we are now forecasting this customer may contribute more than $20 million in revenue in 2027. And beyond that, there's significant additional pipeline and discussions. We also added five new logos in Q3, including model builders and an enterprise data platform. And in the last few weeks, we signed an enterprise-focused foundation model company. Beyond that, we are in promising conversations with the leading frontier model developer for a major expansion, and we just opened discussions with the leading social media company. The business we are winning reflects how well we continue to align with the needs of frontier model builders, even as those needs grow more sophisticated across virtually every dimension. In just the last 18 months, model training emphasis has shifted from human-based evaluations to agentic traces within sandbox replicas of real-world environments, commonly called RL gems, and from synthetic data to a renewed appetite for real-world data. Each of these shifts has expanded both what is possible and what is required, and each has created new opportunities for us. Our significant new wins are also enabling us to meaningfully drive customer diversification. One dynamic worth noting, several times researchers we have worked with at the leading labs have moved to other labs or to well-funded startups, and when they land, they call us. They knew our capabilities and our people, so they brought the relationship with them. This has helped us, and we believe it will continue to help us grow our pipeline at a rate that outpaces SG&A growth. Today, we are raising our guidance from 40% to 42% for 2026. Given the magnitude of the business we are winning, our ways could have been more substantial, but for the fact that in mid-October, our traditionally largest customer decided to end one of its programs with us. The good news is that we are expecting this revenue to be replaced by new programs at this customer by Q1. But we expect it to cost us about $12 million in Q4 revenue. I'd like to emphasize again, as we have in the past, that within our large customers at any point in time, there are programs ending, other programs starting, and still other programs in discussion. This reflects the dynamic nature of our business and represents our business running as designed. I'd now like to take a few minutes to drill down on six specific growth drivers that we believe will propel our growth over the next 12 months. I'll start with driver number one, physical AI. Capital is pouring into physical AI faster than into any category since the LLM boom began. Physical AI is sometimes referred to as embodied intelligence or simply AI robotics. In the first half of 2026, global venture funding in the space totaled $47.4 billion across 521 deals. This exceeds the amount invested in physical AI from 2020 through 2024 combined. The bets are concentrated in robot foundation models and companies looking to build consumer and industrial robotics. Recent research confirms that robotics foundation models improve with the data The same way LLMs do. But because no internet-scale corpus of physical interaction exists, the data must be manufactured through teleoperation, egocentric human video, and simulation. All work we've begun doing. And as models move from lab demos to deployment in factories, warehouses, and homes, we anticipate demand expanding from one-off pre-training corpora to continuous task-specific data for fine-tuning and evaluation. As a data provider, we become recurring infrastructure. We've been working hard to capture the opportunity and have made significant strides across robotic data collection, including egocentric data, OMI gripper data, and motion capture data. The creation of large-scale AI-assisted action annotations for VLM training and robot safety testing and evaluation. Our physical AI data business is scaling materially. We operate large ecocentric data collection programs, stereo camera headset capture in home environments across the United States and Latin America, and six camera headset capture of complex industrial tasks across hundreds of businesses. Our customers are leading humanoid and frontier robotic labs with big ambitions. We are seeing just one of our recently acquired humanoid customers representing a potentially multi-million dollar opportunity for us next year. In September, we launched our state-of-the-art Vicon Motion Capture Lab to capture high-fidelity 3D human motion for world model training in humanoid robot training. And we have demonstrated an end-to-end real-to-SIM-to-real pipeline bringing captured human motion into i6 SIM, retargeting it to humanoid kinematics and executing our own Unitary G1. Beyond collection, we are quickly becoming a trusted annotation partner to physical AI labs, delivering more than 620,000 robot executable labeling jobs this year

so far. This is a new focus area for us, which I expect to become a big growth driver for 2027 as robots move from labs to production. We are launching an annotation pilot with a leading robotics lab, which if scaled could be a sizable opportunity. We also want to pilot with the leading Frontier Lab. If the pilot is successful, we anticipate it could yield up to $7 million of run rate revenue in 2027. We're also working with a major technology company on safety red teaming data for robot foundation models. I'll next talk about driver number two, RL Environment for Personal and Knowledge Work Agents. I mentioned earlier the shift towards training in RL gyms. The reason for this shift is simple. An agent learns by acting and seeing the consequences. And labs can't let agents learn on live systems for safety and privacy reasons and because many actions can't be undone. So the agent trains inside a sandbox replica where it can explore, make mistakes, and reset thousands of times. InnoData is now producing at scale the kind of high quality RL environments that we believe will dominate model training for the foreseeable future. We anticipate this movement and we have built strong research team with deep expertise in environment design and reinforcement learning. We are particularly focused on two use cases, personal agents and knowledge work agents, and we intend to lead in these areas. On the personal agent side, we've been working for some time with one of our largest customers on personalization for its consumer agents. If you follow the consumer agent launches this year, you saw what personalization at that level looks like in the hands of hundreds of millions of users. We were part of the work behind that. Our data, our environments, and our methodologies are embedded in how agents of that kind learn what a person actually wants. We believe personal agents are the next competitive battleground and the winners will be the companies whose agents truly know you and can tailor the experiences to you. The sophisticated RL gyms that capture that kind of depth are a key differentiator in that race and our customers have told us the methodologies and tooling we have developed for extracting the data behind our personal agents are unlike anything they have seen. We expect to be building more of these personalization agent gyms over the next several quarters. On the enterprise side, the battleground among model builders has shifted from coding agents to agents that perform everyday knowledge work across desktop and browser-based workflows that span multiple systems, start from unstructured inputs and require judgment at every step. Models are now being asked to do the expert work that rules-based automation never could. We are building sophisticated RLGMs to train these models, and we are building what goes inside them as well. We are in the early innings of this work, and we believe our team and our technology position as well, and we intend to lead in training these new generation of agents. I'll next talk about driver number three, trust and safety. And I'll start with alignment, by which we mean the work of making AI systems behave safely and in accordance with human intent and policy. We believe the need for this work is increasing for a couple of very straightforward reasons. First, as AI moves from one-shot answers to multi-step reasoning and from chatbot assistance to autonomous agents, the cost of alignment failure increases and evaluating safety becomes more demanding. Secondly, models are being deployed in sensitive domains such as health and finance, where failure carries real consequences. As we have all seen in recent months, as models become more capable, the attendant risk becomes greater. The labs are responding by shifting more efforts toward alignment post-training and specifically humans-in-the-loop data. This presents a significant opportunity for inner data. Our alignment work today spans several high-stakes areas. For example, we are finalizing a new statement of work valued at approximately $6 million of potential revenue to perform chemical, biological, radiological, nuclear, and explosive, or CBRNI, risk evaluation agent safety work. This is with a large hyperscaler that we announced on our first quarter call had selected us as a global trust and safety partner. In addition, the customer has verbally committed to a program with us under which we will test agentic safety of some of their platforms. We believe this program has the potential to yield approximately $2.5 million of revenue next year. We are also in pilots on several mission-critical safety programs in health, cybersecurity, and robotics. We believe these pilots will be successful and that the programs will scale significantly. Several of these pilots use what we call temptation gyms. Environments built to reveal whether a model will pursue a legitimate goal by illegitimate means or cut corners to get there. I'll say more about these in a moment. Across trust and safety, as across our agentic and RL work, we are gaining ground because of how our customers see us as a trusted advisor, a serious applied research organization, and a reliable delivery partner. That combination makes us sticky within these accounts, and it lets us help shape demand, not just respond to it. I'll now turn to driver number four. End-to-end training pipelines that span pre-training to agentic AI. Here is the context. The open web is running dry. EPOC AI estimates that labs will exhaust the supply of

public human written text sometime between 2026 and 2032. Every stage of model development now depends on data that must be deliberately created, such as specialized pre-training corpora and RL environments and evaluations. I believe InnoData is one of the few companies in the world doing this work end-to-end and bringing deep domain expertise to all of it. On the pre-training side, we are going beyond standard web-based sourcing, instead securing unique, differentiated datasets. We see demand for this in several areas. Our largest opportunity is a potential program we are discussing with the leading frontier lab. It is intended to span a wide variety of unique data types at a run rate in the millions of dollars per month. Video is a good example of data that can't simply be scraped at the quality and with the rights that labs need. We are in conversation with one customer on a high-end cinematic video opportunity valued at approximately $2 million. Charts are another area we are focused on. Our teams are building complex charts and graph data sets to strengthen visual reasoning in models, and we have secured several new customers for these data sets. We are also expanding the real-world business data licensing practice. Frontier Labs want raw, real-world data as well as agentic trajectories with verifiable outcomes required to train their models. Because we work so closely with our customers, engineering teams, we have an advanced understanding of which datasets will deliver the performance lift they're looking for. And that same knowledge makes us a natural partner for the owners of that data who rely on us to connect it to the problems it can solve. Human evaluation continues to progress from generalist crowds to expert pools. We recently won a new contract value at approximately $1.2 million with a leading global provider of legal and regulatory solutions. Under this agreement, our legal experts will provide human-in-the-loop validation and the expert feedback used to fine-tune the customer's specialized legal AI models. The win highlights a strength in high-stakes workflows and establishes a scalable blueprint for high-margin, multi-year growth in enterprise AI. We're applying the same expert model to agentic AI, using knowledge workers to build agentic gyms, environments where agents perform real computer use and knowledge work tasks against realistic data. As agents take on more autonomous work, labs need to know not only whether an agent can complete a task, but whether it completes it honestly. Our temptation gyms are exactly built for that. These environments are deliberately designed with opportunities to cut corners. Gaming a reward, fabricating a result, misrepresenting what was actually done, and they are instrumented so labs can see whether an agent takes them. That lets lab catch behavior that in real deployment could amount to a misrepresentation of fraud. This kind of integrity testing is becoming essential and requires both domain expertise and environment building capability. Temptation gyms are a new category and we believe being early gives us capability that will differentiate us. Next, I'll address driver number five, enterprise AI. As we have said previously, implementation of generative AI models in the enterprise is a very large opportunity for us. We have been building a practice for the last three quarters and we continue to gain momentum. Just last quarter, we added approximately $5 million of revenue run rate of agentic implementation for a large customer. We also initiated two new enterprise relationships with a global technology hardware company and a major financial services firm. And we expanded our work with a social media customer. We have served since last year to about $2.5 million a year. We believe these accomplishments give us a strong foundation to accelerate from in 2027. We have built a three-tier partnership strategy that enables us to help our enterprise customers at the foundation model layer, the data layer, and the infrastructure layer. Our partnership with Anthropic, a leading foundation model provider, helps us drive enterprise agentic solutions into production. Our partnerships with Databricks and Snowflake, leading data platforms, help us build the AI-ready data foundations that every successful deployment depends on. And our partnership with Google Cloud enables us to deliver and scale solutions on enterprise-grade cloud infrastructure. Moreover, these channel partnerships gives us additional growth paths and joint opportunities. We believe our experience in training frontier models puts us in a unique position to understand how to create golden datasets, benchmark models, build custom rubrics and deploy end-to-end data engineering at scale to create robust, production-ready solutions for enterprises, and gives us a unique advantage compared to traditional systems integrators. Lastly, I'll talk about our sixth driver, our federal practice. We launched our federal practice about a year ago. Since then, it has built a position in government AI programs that compound each award. In its first year, that produced two programs and a third committed through our prime partners. The opportunity we described as a finalist in May, an AI development environment for the Department of War, is in phase one delivery with a production decision expected in the coming months. We were awarded advanced data labeling on a U.S. Navy autonomy program, the first phase of a competitive

down select selected on the computer vision capability we demonstrated to the government this year. And we have been identified for a subcontract role on a second multi-year Department of War program. Each of these is staged with a larger award decided after the current phase performs, which is where we have chosen to compete. We have a verbal win for one contract of 5 million per year for five years with signing expected by end of the fourth quarter. We hope to continue momentum across all programs. We are ending the year with real momentum, a clear strategy, and tailwinds from each of the six growth drivers I've just walked you through. I've never been more confident about where InnoData is positioned. The opportunities in front of us are both large and they are close at hand. We are investing to capture them and we expect those investments to pay back quickly and to build durable advantage. The next 12 to 18 months will be about execution and that is where Jayant and I are focused. Let me turn over the call to Jayant, our CFO, to share his perspectives about the business and our Q3 financial results.

Jayant Chauhan | Chief Financial Officer:

Thank you Rahul and good afternoon everyone. Let me walk through our third quarter results and also share our thinking on investments and capital allocation. Revenue for Q3 2026 was $95.6 million, up 53% year-over-year and 4% sequentially, our 13th consecutive quarter of year-over-year growth. This exceeded analyst consensus by $12.3 million, or 15%. Adjusted gross profit was $38.3 million representing adjusted gross margin of 40%. This is in line with our 40% target. In any quarter, our gross margin is driven by program mix. Adjusted EBITDA was $17.7 million, or 19% of revenue, up 9% year over year. This was ahead of consensus of $15.4 million by approximately 15%. Net income was $10.3 million, up 24% from the $8.3 million we reported in Q3 last year. Fully diluted earnings per share was $0.29 compared with consensus of $0.17, a beat of approximately 70%. Our year-to-date effective tax rate was approximately 12% compared to our long-term target range of 23% to 25%. The lower tax rate was driven by tax benefits recognized this quarter. Turning to the balance sheet. We ended the quarter with approximately $164 million in cash and cash equivalents. Excluding approximately $10.3 million of customer prepayments, which are a pass-through, our cash was $153.7 million. This is up $19.3 million from $134.4 million on the same basis at the end of Q2. We have no debt. Our $50 million Wells Fargo credit facility remains undrawn. We did not use our ATM program during the quarter. With the full quarter behind me, let me share some perspectives on investments and capital allocation. Given the opportunity set Rahul described, we will continue to invest in areas we believe will become meaningful growth drivers in 2027 and beyond. We intend to fund these investments from operating cash flow, and we expect them to pay back quickly. On capital allocation, with approximately $154 million of cash, net of customer prepayments, and no debt, we have real flexibility. Our first priority is organic investment in the growth drivers. Beyond that, we will be selective about M&A that we believe can accelerate our strategy. And we intend to make these investments while maintaining a strong balance sheet. To sum up, Q3 delivered record revenue at just a bit ahead of expectations at our 40% gross margin target and $19.3 million of underlying cash generation. While we continue to invest in the drivers, we believe will enable our next phase of growth. With that, I'll turn it back to the operator. Operator, we are ready for questions.

Operator | Conference Call Operator:

We will now begin the question and answer session. If you would like to ask a question, please press star one to raise your hand. To withdraw your question, press star one again. We ask that you pick up your handset when asking a question to allow for optimum sound quality. If you are muted locally, please remember to unmute your device. Please stand by while we compile the Q&A roster. Your first question comes from the line of George Sutton with Craig Hallam. Your line is open. Please go ahead.

George Sutton | Analyst, Craig-Hallum:

Thank you guys and thanks for all the information. So can you explain the process of your labor pool and how that works? So you have a program that ends and this happens all the time. I think typically you would like to deploy those folks into other jobs immediately. You mentioned Q1 could be a similar kind of a deployment. Do these require different skill sets Do you then rehire from that pool? Can you just give us a picture there?

Rahul Singhal | President and Chief Executive Officer:

Hi, George. Thank you for your question. So yeah, I think the way the programs start and stop is dependent on the types of tasks that these frontier model builders are working on. Sometimes the skills that they're working on could be redeployed, and sometimes those skills cannot be redeployed. The other thing to note is, When these different frontier model builders are in different cycles of their training and they are in different training, training the different tasks. So I think it becomes a matter of what's available, what's usable or reusable at that time from a Liverpool perspective. And I think when we talk about Q1, I think it's different programs that we are expecting us to start.

George Sutton | Analyst, Craig-Hallum:

I understand. And you made an interesting comment when you were talking about the physical AI, I think specifically to robotics work, that you look to become recurring infrastructure, which is quite different than the normal programs that you've traditionally worked on. Can you just explain that?

Rahul Singhal | President and Chief Executive Officer:

Yeah, absolutely. I think we are very excited about physical AI. So if you think about the investment that is going in physical AI, it is, I believe, going to outpace the investment that happened in LLM. Now, one of the most interesting things about physical AI or embodied intelligence is the fact that you can't scrape data from the web as the LLM started to. So the pre-training data or the pre-training corpora doesn't exist. The way these models are getting trained is it starts by doing something called as egocentric data, where you are using cameras to collect data for different tasks. So a task could be you're collecting data for loading a dishwasher, or you could be collecting the data for Holding Your Laundry, right? Those are egocentric data. Then you have something called as UMI datasets, and then you have something called as motion capture datasets. Now, these are different modalities, different ways, different technologies and techniques that are being used, and each one progressively becomes a lot more interesting, a lot more complex, and the models start to become better. This is where the world models start to play in. because there is not a single layer of physical corpora that people can use, you have to continuously be collecting, creating, annotating the data so that these robots can actually execute around that. So I think it's a very interesting and exciting opportunity for us.

George Sutton | Analyst, Craig-Hallum:

Super. And then just one for Jayant. So relative to the guide for Q4, Historically, what you've done is your guide includes only what you have as booked business. So if you get additional business during the quarter, that creates some upside potential. Can you just talk to the thought process with the Q4 guide and then also reference the pipeline of opportunities and kind of speed of deal cycles that could result in more revenues in Q4?

Jayant Chauhan | Chief Financial Officer:

Hi, George. Thank you for the question. In terms of how we think about guidance, that approach is still consistent with in the past what we have done. We continue to be prudent about it. And you're right. We generally factor in only revenue from active programs, new awards, or smaller pipeline opportunities where we actually have a clear line of sight. We do not factor large pipeline opportunities until they are awarded and until we know the timing of that revenue recognition. So continue to be consistent with that approach as of now and going forward as well. And on to your second question on, of course, and on your second question of what we're seeing in the pipeline, I believe Rahul talked about a lot of those opportunities which we are seeing today in terms of guiding for 2027. As you're aware, we'll announce that along, or rather we intend to announce that along with our Q4 results. So you'll hear more about that over the next quarter.

Rahul Singhal | President and Chief Executive Officer:

Thank you.

Operator | Conference Call Operator:

Your next question comes from the line of Alan Klee with Maxim Group, LLC. Your line is open. Please go ahead.

Allen Klee | Analyst, Maxim Group:

Yes, hi. You spoke about how your adjusted gross margin hit your target of 40% plus for the quarter. In the prior two quarters, it kind of over-earned. It was above. And could you just give us a little more color about What happened with the mix and is there a reason why you think it might go in a certain direction going forward? Thank you.

Rahul Singhal | President and Chief Executive Officer:

Hi, Alan. So I think as we have consistently said that our target gross margin is 40%. And in any quarter, the program makes the margin mix on those programs change. and we also tend to do a lot of programs where we believe that they could be big growth drivers for us in following quarters or they're building capabilities which we tend to take at lower margin. So we have consistently guided to 40% gross margin and we will continue to guide to 40% gross margin because I think this business requires us to ensure that we are investing in different programs, building our capabilities and investing, going ahead in a way that we can continue to lead the market.

Allen Klee | Analyst, Maxim Group:

Okay, thank you. And then For the project that was shut down in October, and then I think you said that's going to have a $12 million impact in the fourth quarter, but the same former largest customer is going to add that amount in the first quarter. Is the margins of the business shut down? Is there a way to think? Is that materially different than the... The margins of the new business?

Jayant Chauhan | Chief Financial Officer:

Oh, hi. Let me take that, Rahul. Well, look, on the new business, as you know, we don't really guide or we don't talk about any project level margins. So I would just point to the fact that we'll talk about 2027 and what that looks like during our Q4 results.

Allen Klee | Analyst, Maxim Group:

Okay, and then the last thing about capital allocation, where you were talking about organic growth, were you trying to imply that the amount of money you might spend on organic growth to fund that will increase relative to where it had been, or what were you trying to get across to us? Thank you.

Jayant Chauhan | Chief Financial Officer:

Yeah, I'll take that, Rahul. Yeah, as you know, we've been talking to you guys in the past and to our investors as well. There are a lot of questions on how are we thinking about capital allocation given the balance sheet strength. So just basically clarifying that organic investments continue to be our priority. And that's where we'll focus as we invest behind the various growth vectors and opportunities Rahul laid out. So that continues to be the priority. And if there is opportunistically, we see many opportunities which can accelerate a growth strategy and create value for shareholders, we will continue to look at that as well. But priority is organic investments.

Allen Klee | Analyst, Maxim Group:

Thank you.

Operator | Conference Call Operator:

Your next question comes from the line of Hamid Khorsand with BWS Financial. Your line is open. Please go ahead.

Hamid Khorsand | Analyst, BWS Financial:

Hi. Just one question about this large customer moving out one project to another. Is that other project won by you, or do you have a confirmation that's happening in Q1?

Rahul Singhal | President and Chief Executive Officer:

Yeah. So, Hamid, we don't comment on comment on exact timing. We anticipate winning more work, and we are in late stages of several programs that we anticipate will help us fill the gap. So I think the way to think about it is programs start and stop. New programs are starting. We are working with multiple stakeholders, multiple business units within our largest customer. And this business moves really, really fast. So we anticipate filling the gap.

Hamid Khorsand | Analyst, BWS Financial:

And then it sounds like you have a new large customer. What's the percentage of revenue?

Jayant Chauhan | Chief Financial Officer:

Yeah, so you're talking about the third large customer Rahul talked in his comments earlier? Yeah. Yeah, so at this point in time, we're expecting to do about $20 million in revenue from that customer over the next year.

Hamid Khorsand | Analyst, BWS Financial:

Okay. And then as far as... What we're seeing from our largest customer is

Rahul Singhal | President and Chief Executive Officer:

This market is moving significantly faster and the investments go in a variety of different directions. The customers tend to invest in things like VLMs, which is vision learning models. They're investing in robotics. They're investing in LLMs. They're investing in Agents, right? So the business is moving fast and we are working with variety of different business units across the gamut within the largest customers in trying to address their needs and work with them. We are spending time with them on reinforcement learning gyms where we are creating a cutting edge gyms that allow the agents to be trained and giving them our expertise and our thought leadership in helping them trying to dominate that market. So I think different places is where we are. We continue to work with them in variety of different programs.

George Sutton | Analyst, Craig-Hallum:

Okay, thank you.

Operator | Conference Call Operator:

There are no further questions at this time. I will now turn the call back to Rahul Singhal for closing remarks.

Rahul Singhal | President and Chief Executive Officer:

Thank you, operator. To wrap up, Q3 2026 was another record revenue quarter for InnoData and another across-the-board beat. We delivered 53% revenue growth, our 13th consecutive quarter of year-over-year growth, adjusted gross margin of 40%, adjusted EBITDA ahead of consensus, and strong net income, all while continuing to invest in the research and innovation that is driving the business. Our diversification continued. A customer that represented less than 5% of our revenue a year ago is now our largest customer, giving us two customers of comparable scale. Another large hyperscaler customer, with which we did only $3.1 million in revenue in the trailing 12 months, is on track to be a $20 million-plus customer next year. We also have one notable new customers across Frontier Labs, enterprise data platforms, and robotics. We are raising our guidance to 42% or more growth for 2026. Given our wins, the raise would have been a lot more, but for an air gap in Q4 revenue from our historically largest customer that we forecast getting filled by Q1. Most importantly, we believe we are positioned at the epicenter of the most important transition in AI right now. The move to the next S-curve of agents, reinforcement learning and embodied intelligence, and at the center of the trust, safety and alignment work that will determine whether that transition succeeds. These are primarily data challenges and they're precisely in our wheelhouse. I believe 2026 will finish as a strong year for InnoData and 2027 will prove equally or even more exciting. It is a privilege to be leading this company. I

thank our team and customers around the world and all of you for continuing to be on this journey with us. We look forward to speaking with you next quarter. Thank you.

Operator | Conference Call Operator:

This concludes today's call. Thank you for attending. You may now disconnect