NASDAQ:INOD Q2 2026 Earnings Call Transcript

Generated on 8/9/2026

Operator | Conference Operator:

Thank you for standing by. At this time, I would like to welcome everyone to the InnoData second quarter 2026 earnings call. All lines have been placed on mute to prevent any background noise. After the speaker's remarks, there will be a question and answer session. If you would like to ask a question during this time, simply press star followed by the number one on your telephone keypad. If you would like to withdraw your question, press star one again. Thank you. I would now like to turn the call over to Amy Agress. You may begin.

Amy Agress | Investor Relations:

Thank you. Good afternoon, everyone. Thank you for joining us today. Our speakers today are Jack Abuhoff, Chairman and CEO of InnoData, Rahul Singhal, President and Chief Revenue Officer, and Jayant Chauhan, Chief Financial Officer. Also on the call today is Marissa Espineli, Chief Accounting Officer, and Amish Pantarkar, Senior Vice President, Finance and Corporate Development. We'll hear from Jack and Rahul first, who will provide perspective about the business and then Jayant will provide a review of our results for the second 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 or other statements about future events. These statements are based on current expectations, assumptions and estimates and are subject to risks and uncertainties. Actual results could differ materially from those contemplated by these forward-looking statements. Factors that could cause these results to differ materially are set forth in today's earnings press release in the risk factor section of our Form 10-K, Form 10-Q, and other reports and filings with the Securities and Exchange Commission. We undertake no obligation to update forward-looking information. 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 Jack.

Jack Abuhoff | Chairman and CEO:

Thank you, Amy, and good afternoon, everyone. Q2 is another record quarter for InnoData. Revenue, adjusted gross profit, adjusted EBITDA, and cash all reached new highs, and we exceeded analyst consensus on all key metrics. Revenue is 92.1 million, up 58% year-over-year, exceeding analyst consensus by approximately 5.8 million, or 7%, and making Q2 are 12th consecutive quarter of year-over-year growth. To put that in perspective, in Q2, as in Q1, our quarterly revenue exceeded our annual revenue of just three years ago. Our adjusted gross margin, meanwhile, was 49%, up two points sequentially and nine points above our 40% publicly stated target. Adjusted EBITDA was 25.4 million, up 92% year-over-year, exceeding analyst consensus by approximately 8.5 million, or 50%. Fully diluted earnings per share were 41 cents per share, nearly double analyst consensus, 21 cents per share. Again this quarter, we delivered growth, margin expansion, and cash generation together. while investing in innovation that converts to revenue within quarters, not years. That is the business model working as designed. Last quarter, we told you to expect our largest customer to represent a smaller percentage of total revenue. In Q2, our largest customer represented 37% of revenue, down from 56% of revenue in Q1, while the big tech customer we announced last quarter scaled from 17% of revenue to 34% of revenue, becoming our second-largest customer. While our largest customer contributed less revenue in Q2 than in Q1 as a result of a change in the quarter to program structure and service mix, we continue to expect it to grow year over year for the full year. We also landed an important new customer in the quarter, one of the fastest-failing frontier labs. The upshot is our base continues to broaden in both customers and customer programs. Now, before turning to guidance, I want to share an important announcement about InnoData's leadership. Effective September 30, Rahul Singhal will become president and chief executive officer of InnoData and will join our board. And I will transition into the role of executive chairman. This is a planned transition made from a position of strength and for me, it is also a personal one. Many of you know Rahul from these calls, from investor conferences and from the work he has led over the past several years as a principal architect of metadata transformation and for strategic partner to the world's leading AI builders. He knows our customers, he knows our technology and he knows our people. Rahul has been central to every element of the strategy behind the results you've seen quarter after quarter. The board and I didn't have to look far for the right leader. Rahul earned this role, taking on expanding responsibility year after year and delivering every time. This is how we build this company. We grow our capabilities and we promote our own people. As executive chairman, I will remain deeply engaged, focused on partnering with Rahul to build capabilities enabled by our research team. Bringing these capabilities to the federal government and to the enterprise, I believe, is where I can best contribute to creating significant shareholder value. And as one of the company's largest shareholders, that is exactly what I want to be doing. Our work with the MAG7 and leading AI labs is on a firm path to greater heights and greater diversification. Our enterprise AI and federal strategy built on the differentiated technology we developed for the Frontier Labs represent opportunities for potentially driving high-quality recurring revenue that results in significant value creation. We are building InnoData to be a generational company. And with that same aspiration in mind, we were pleased to have announced recently that Jayant Chauhan joined InnoData as Chief Financial Officer. Jant's abilities round out an already strong finance team with Marissa Espineli stepping into the role of Chief Accounting Officer. Beyond the traditional CFO mandate, Jant will work strategically on capital allocation and capital markets, customer partnerships, M&A that can accelerate our strategy, and investor communications while scaling the financial infrastructure of the company we are becoming. Before I turn the call over to Rahul, let me address guidance. We are reiterating our guidance of 40% or more year-over-year revenue growth. We have some large new potential engagements in our pipeline with both existing and new customers that we believe are likely wins, but we have not yet factored them at all into our forecast at this point. As a matter of prudence, we will only factor them into our forecast when we know they're 100% won and we can forecast the timing of revenue recognition. I will now turn the call over to Rahul to discuss the market, our strategy, and the execution milestones that we believe prove the strategy is winning.

Rahul Singhal | President and Chief Revenue Officer:

Thank you, Jack, and good afternoon, everyone. Before I begin, a personal note. I'm truly honored by the confidence both Jack and the board have placed in me, and I intend to repeat the results InnoData has extraordinary momentum, an extraordinary team, and an extraordinary opportunity in front of it. I intend to build on all three. One of the most significant developments of the past 18 months is the increasingly pivotal role that research and innovation are playing at InnoData. It is not overstating the case to say that research has become a growth engine and the means by which we increasingly differentiate expand existing partnerships and forge new customer relationships. Our growth is increasingly driven by research and innovation across the full model training lifecycle, from pre-training and post-training to model evaluation and benchmarking. Our innovation is producing intellectual property and differentiation that is generating demand. Several quarters ago, we talked about how people benchmarking frontier model performance. isolating weaknesses, building remediation datasets to address those weaknesses, and proving the efficacy of those datasets by training small models that were architecturally similar to the big ones. Today, we are doing much more than that. I'd like to share a few examples of what we are doing now because the work is fascinating in its own right, and because it gives you a sense of where we intend to take InnoData over the next several years. Through our research efforts, We established an early position in agentic reinforcement learning, one of the most important frontiers in AI development. With a large lab, we run a significant new program covering personalization of long horizon agents, which is now scaling. And we have also been awarded a second program covering reinforcement learning environments for desktop computer use agentic tasks. In the enterprise, We see companies quick to develop AI agents, but struggling to deploy them in production with confidence. We believe combining our trusted observability platform and our innovatively architected reinforcement learning gyms enable us to position ourselves as the AI deployment assurance layer. We see this as opening a huge opportunity, and this is what Jack alluded to a few minutes ago. In the quarter, We deepened delivery of these capabilities with one big tech customer and began delivery with another. This innovation has also opened active insurance and banking conversations that we expect to convert to pilots. Frontier model builders have also become intensely focused on dynamic, long horizon, agentic evaluation. In the quarter, we released two public benchmarks. including one that tests how well models perform on multi-turn, long-context, and multi-model interactions. A benchmark is an assembly of expert author prompts, rubric constraints, and L&M judges configured to test frontier models. Both are designed to surface failure modes that standard leaderboards miss, things like grounding drift and instruction forgetting. Besides these failure modes, frontier labs are working to improve. Each benchmark engagement results in a data strategy recommendation and sets us up to deliver scaled data generation to improve the model. In the quarter, we also expanded our capabilities in generating training data that extends the reasoning capabilities of the state-of-the-art models, delivering across five frontier labs and five domains. As AI moves from digital tasks to embodied intelligence, we are building the required data and measurement layers. This quarter, we signed two research agreements with a leading university and committed to a motion capture lab that we expect to come online in the next few months, capable of collecting sub-millimeter precision data for training robots and physical AI foundation models. In the quarter, we ran successfully ecocentric data collection pilots with leading robotics companies and that data connection practice shifted from individual pilots to scoping enterprise-scale multimodal programs, including a multilingual speech program spanning seven languages and a roughly 2 million hour ecocentric program that we hope to be awarded based on successful pilot results. Data, data engineering and data science are central to improving AI and to making it safe and trustworthy. That centrality is what enables our research to deliver capabilities across many different spheres. Data engineering innovations can solve big AI challenges, including in domains where you might not expect to find us. We mentioned one such domain in our Q4 call. How we had developed an AI model for drone and other small object detection that exceeds prior state-of-the-art benchmarks by 6.45% and how in the field their progress is often measured in fractions of a percentage point. A 6.45% improvement is a material advance. We are now working on demonstrating that capability to the government. Another example, as we announced earlier this week, we released the full stage of what we are calling our AI Cyber Training Suite. Well, data sets and evaluation systems that train AI coding agents to write secure code and to repair vulnerabilities in the company's existing software. When we tested leading open-weight models on their ability to repair verified flaws, the repair rate more than doubled after a single round of fine-tuning of just a portion of our data. Given that AI now writes a growing share of the world's code, the inability to trust that code without the security team reviewing everything it produces is the real blocker to enterprise adoption. We believe our suite has the potential to remove that blocker. Across frontier labs, federal and the enterprise, the pattern is the same. Research and innovation are creating differentiated capabilities that win programs and compound into durable customer relationships. We couldn't be more excited about the opportunity ahead of us. Jack, back to you.

Jack Abuhoff | Chairman and CEO:

Thanks, Rahul. I also want to take a few minutes to connect this quarter's results to the structural economics of our business and to spend a few minutes talking about the broader market dynamics. First, operating leverage. Revenue grew 58% year over year, while adjusted EBITDA grew 92%, roughly 1.6 times faster. Each incremental program builds on the same core operating infrastructure, So the marginal cost of the next program is meaningfully lower than the cost of building that capability from scratch. Second, margin quality. Adjusted gross margin of 49% is nine points above our publicly stated target. The expansion is driven by mix and bolstered by the high-value pre-training programs and off-the-shelf data sets. where we retain IP and monetize the same asset across multiple customers. These are the software-leveraged economics we have been deliberately building toward. Now turning to the broader market dynamics, there are debates about whether we are at a peak AI capex, whether competition will commoditize models, and what the recent security incidents mean for the industry. Now these debates play out against extraordinary numbers. Hyperscaler capital spending is guided to roughly $700 billion this year, nearly double last year, with estimates revised upward throughout the year. But we believe each of these debates resolves in favor of the data, evaluation, and assurance layer we provide. Let me explain. If CapEx comes under pressure, monetization pressure rises, and monetization runs on deployment, Fine Tuning, and Assurance, our business. If inference commoditizes, two things follow. Labs engineer for use case specific differentiation, which requires specialized data, and AI becomes more accessible to the enterprise, which requires more assurance, not less. Again, our business. If security incidents multiply, they prove the need for exactly the engineering we announced this week, yet again, our business. However the market moves, we believe it moves toward the work that we do. I will now turn the call over to Jayant, our new chief financial officer, to walk through the financials.

Jayant Chauhan | Chief Financial Officer:

Thank you, Jack, and good afternoon, everyone. I'm Jayant Chauhan, InnoData's chief financial officer. and as you know, this is my first earnings call since coming on board in July. Marissa has transitioned into the role of Chief Accounting Officer and I'm thankful to her for her partnership in getting me up to speed quickly. I've spent the past several weeks getting to know the business and meeting our teams here in the US and around the world. I'm energized by what I've found. I look forward to getting to know many of you on this call and afterwards. With that, let me walk through our second quarter results. Revenue for Q2 2026 was $92.1 million, up 58% year-over-year and 2% sequentially, our 12th consecutive quarter of year-over-year growth. This exceeded analyst consensus by $5.8 million, or 7%. Adjusted gross profit was $45.4 million, representing adjusted gross margin of 49%. that was 2 percentage points higher than Q1 and 9 percentage points above our externally communicated 40% target. The improvement was driven by the mixed shift towards higher margin programs. Adjusted EBITDA was $25.4 million, or 27.5% of revenue, up 92% year-over-year. This exceeded analyst consensus of $16.8 million by approximately 50%. Net income for the quarter was $14.4 million, double the $7.2 million we reported in Q2 last year. Fully diluted earnings per share was 41 cents, exceeding the consensus estimate of 21 cents by approximately 95%. Our effective tax rate for the quarter was approximately 18% compared to our long-term target range of 43 to 25%. The lower tax rate was driven by tax benefits recognized this quarter. Turning to the balance sheet, we ended the quarter with $250.4 million in cash and short-term investments. Excluding customer prepayments, which are a pass-through, our cash and short-term investments position was approximately $134 million, up $37 million sequentially. We remain undrawn against our Wells Fargo credit facility. Lastly, after market close today, we will find a prospective supplement establishing an at-the-market equity program with Goldman Sachs as lead agent alongside a broader syndicate. The program provides an efficient supplemental capital markets tool that we can use selectively and opportunistically. Our balance sheet is strong with cash and short-term investments of approximately $134 million net of customer prepayments and has no debt outstanding at end of future. The program preserves optionality to support future growth initiatives, potential strategic opportunities, and continued balance sheet trends as we scale. With that, let me close. This was a good quarter for me to step into, and I'm looking forward to building on the growth and financial discipline this team has already established. With that, I turn it back to the operator. Operator, we are ready for questions.

Operator | Conference Operator:

At this time, I would like to remind everyone in order to ask a question, press star, then the number one on your telephone keypad. And your first question comes from the line of George Sutton with Craig Hallam. Please go ahead.

George Sutton | Analyst, Craig Hallam:

Thank you. First, congrats to Rahul and welcome to Jayon. Jack, I still hope to harass you with questions regularly. So I'm curious if We can talk about the things that are not in your guidance. You mentioned some opportunities that aren't necessarily 100% booked yet, thus not in guidance. Can you give us any bigger picture in terms of what some of those opportunities look like? And will you give us more regular updates, perhaps, than just the quarterly announcements?

Jack Abuhoff | Chairman and CEO:

Sure, George. Thank you. And needless to say, I look forward to your questions as often as you'd like to bring them to me. Yeah, no, we were thrilled with the quarter. Really, I think there were a lot of proof points laid down in the quarter and some of the things that we're learning as we go forward are as important to us as the financial signal that you're seeing today. The innovation that we're accomplishing, that we're producing is teaching us is laying out the direction for us. It's showing us that reliability in agentic enterprise AI can be engineered. It's showing us that the kinds of innovations that we're capable of creating and the difference that we can make by operating at the data engineering layer in kind of very random things, drone detection, cybersecurity. These are just two examples. When I look at the set of opportunities we have, they run the gamut. And now I'm responding to your question about the things that are significant, some quite large things that are not in our guidance today. They run across our capabilities. There are things that are on the government side and the enterprise side. There are things that are on the frontier model side. A lot of the capabilities that we're demonstrating now in agentic AI, both deployment and training, are prominent in our pipeline. We're excited about it, but from a methodological perspective, we maintain the discipline to count our chickens only once they're hatched. So we're looking forward to sharing more as we proceed forward. through the second half of the year, we think it's going to be exciting.

George Sutton | Analyst, Craig Hallam:

So the security incidents that we're starting to see in AI are obviously concerning and seem to have created a very nice new opportunity for you. I wondered if you can just walk through that and obviously if you can bring to bear the press release from a couple of days ago with some of your capabilities. What does that mean in terms of opportunity for you?

Jack Abuhoff | Chairman and CEO:

Yeah, good question. So, you know, I think when we look at the problems that the enterprise is having, you know, they want to embrace the genetic AI, but, you know, can they trust it? You know, what are the reasons that they may not be able to trust it? And certainly when, you know, they're reading about, you know, models escaping their sandboxes or, you you know, elite cyber capabilities when they escape containment and things like this, that becomes a real concern. Now, one of the reasons that that concern exists is a lot of the frontier models that are capable of these cyber security disruptions were themselves built on training data that contained unpatched code. It's fascinating. So basically, if you can identify the things that went into their training data mix and you can build an agent that can detect those code aberrations, can detect the code that's been introduced even when patches were subsequently introduced. And then from that, if you can enable that AI to generalize to new novel threats, things that it hasn't seen, and identify threats that are in the existing software, you've got a very capable set of technologies that enable the enterprise to more safely adopt AI. So we're having some interesting discussions about that. We think it's another example of the kinds of innovation that we're increasingly capable of.

George Sutton | Analyst, Craig Hallam:

Gotcha. There's one other question. Obviously, we're seeing more federal government testing of models before they are released, a lot of it through red teaming. Can you just give us a sense of your involvement in the broader federal area?

Jack Abuhoff | Chairman and CEO:

There are a couple of things there. I think we're having a lot of interesting discussions with with players in the government about how we can partner with them and where we can cooperate with them. We're also discussing things with agencies. The ability to be represented in the Tradewinds marketplace is an accepted opportunity. Solution for, you know, different things is a huge, you know, opportunity and a huge advantage that we now have. I think from a perspective of, you know, what will be the federal government's, you know, relationship with AI, there are two things there. There's first, you know, they're very much accelerating variability to procure AI solutions and, you know, get the best. The other thing that we're seeing is the frontier model companies are inviting the government proactively to help them regulate the agency. So when you look at what the eventual need will be for things like benchmarks and evaluations and red teaming, we released two benchmarks this quarter that we think are very novel and very useful. to deliver those kinds of things and evaluation work on behalf of the government is an opportunity that we're tracking.

George Sutton | Analyst, Craig Hallam:

Super. Thanks for answering the questions.

Operator | Conference Operator:

Thank you. Your next question comes from the line of Alan Klee with Maxim Group. Please go ahead.

Alan Klee | Analyst, Maxim Group:

Yes. Hello. You mentioned one of the positives this quarter was a higher mix of higher margin projects. I was wondering, should we think of this as a trend towards that or maybe that was just the mix this quarter and then they revert back to kind of where it's historically been?

Jack Abuhoff | Chairman and CEO:

Yeah, so it's a very good question. I'm going to answer it in the following way. I think it's both. And now let me explain what I mean by that. We do bid on work that has a lower gross margin than the one that you're seeing today. Some of those projects could be large. We would intend to take those on. If we win those, I think the cash flow from them will be, we would anticipate to be quite compelling. Would that mean the gross margin on a weighted basis would decline somewhat? It would. On the other hand, from a strategic perspective, the things that we're working on, the things we're innovating, will likely have a higher revenue quality. And we measure, you know, revenue quality at where we think of revenue quality as a function both of gross margin and, you know, the recurring nature of that revenue. So I think over time, strategically, it's going to trend upward. I think on a quarter-by-quarter basis, it will depend on product mix.

Alan Klee | Analyst, Maxim Group:

That's helpful. Thank you. Ben, also you talked about using off-the-shelf data sets more often to do the training. I'm just trying to, could you explain a little of, like, Do you own the data or you get to use it and use it multiple times? And if you don't, kind of the relative, if you don't do that, how you're accessing the data?

Jack Abuhoff | Chairman and CEO:

Sure. So the off-the-shelf data sets up until now, and I'll come back as to why I said that, for the most part up until now, are data sets that we engineer. And we engineer them around model deficiencies that we detect in our benchmarking. So when we see that there's a deficiency or when we see that or when we identify a capability that the frontier models are looking to create and we can engineer a data set that helps them get there, rather than waiting for them to request that of us, we build that data set, we maintain or we retain the IP associated with that data set. We enabled them to use those data sets for training their models. It's good for everybody, right? It's good for our customers and it's good for us. And that's one of the contributors to higher margin profiles. There are also times when on behalf of someone else who owns a data set, we will represent them. We perhaps do some engineering to that data. We will configure it so that it's ready for models to be trained on it. And then we will invite our customer partners to utilize that data as well. But most of what you're seeing today is data that we've figured out how to assemble around particular model needs and frontier model capabilities.

Alan Klee | Analyst, Maxim Group:

Thank you. My last question is, in the most likely case scenario, is there any reason that it would be likely that there would be a sequential decline in revenues in the third or fourth quarter?

Jack Abuhoff | Chairman and CEO:

So it's, you know, within the constraints of our Business model, it's certainly possible. And if it were to occur, I don't know that I would particularly care. So what I care mostly about is where we're taking the company and where it's going, not quarter to quarter performance. The kinds of innovations that we're producing today, the track record we're getting, the new customers that we're waiting, I think over time, you know, we'll continue to in order to our benefit. And I think that we're going to continue to grow this company. in a very significant way over the next several years. If we were to win a very large one-time project that were delivered in two quarters and then there were an error gap after a third quarter, would I consider that a failure? Not at all. What I would consider a failure is if we're not maintaining the relevance that we are right now to our customers and if we weren't identifying huge market opportunities that I believe we'll be able to explore it over the next several years.

Alan Klee | Analyst, Maxim Group:

Okay, thank you very much. Appreciate it.

Operator | Conference Operator:

This concludes our question and answer session. I will now turn the call back over to Jack Abuhoff for closing remarks.

Jack Abuhoff | Chairman and CEO:

Thank you very much. To wrap up, Q2 2026 was another record quarter for InnoData. It was an across-the-board beat. We delivered 58% revenue growth, 49% adjusted gross margin, 92% adjusted EBITDA growth, and significant cash generation. It was our 12th consecutive quarter of year-over-year growth as well. We're seeing that diversification is happening in practice. Our largest customer declined to 37% of revenue while our overall business grew. And the customer that generated essentially no revenue a year ago is now our second largest customer. We announced a planned leadership transition. Rahul will become our president and CEO on September 30th. I'll become our executive chairman. I'll be focused on building long-term differentiating capabilities across our enterprise and federal markets. Meanwhile, Jayant Chauhan has joined as CFO, further strengthening our financial leadership and enabling me to do some of the things that I want to do. As one of the company's largest shareholders, I believe this is a tremendous path forward to very significant shareholder value creation. As we've discussed, our growth is increasingly research-driven and innovative, from novel benchmarks and reinforcement learning environments to capabilities and agenda deployment assurance, physical AI. I think we're at the very early stages of many of this. So we're very excited about what lies ahead. We're very confident that 2026 can be a tremendous year for InnoData, and I thank all of you for continuing to be on this journey with us.

Operator | Conference Operator:

Ladies and gentlemen, that concludes today's call. Thank you all for joining. You may now disconnect.