# GSI Technology, Inc.

# Second Quarter Fiscal 2026 Results Conference Call

# October 30, 2025

## CORPORATE PARTICIPANTS

**Lee-Lean Shu,** *Chairman, President and Chief Executive Officer*

**Didier Lasserre,** *Vice President, Sales*

**Douglas Schirle,** *Chief Financial Officer*

## CONFERENCE CALL PARTICIPANTS

**Robert Christian,** *Private Investor*

**Mark Gadner,** *Private Investor*

**David Zalkowitz,** *ISQ*

**Christian Rug,** *CER Holdings*

**Michael Cooper,** *Private Investor*

**Michael Roberts,** *Roberts Capital*

**Marco Petroni,** *MG Capital*

**Mohammed Alsousi,** *Scale*

## PRESENTATION

**Operator**

Welcome to GSI Technologies Second Quarter Fiscal 2026 Results Conference Call.

At this time, all participants are in a listen-only mode. Later, we will conduct a question-and-answer session. At that time, we will provide instructions for those interested in entering the queue for Q&A.

Before we begin today's call, the Company has requested that I read the following Safe Harbor statement. The matters discussed in this conference may include forward-looking statements regarding future events and the future performance of GSI Technology that involve risks and uncertainties that could cause actual results to differ materially from those anticipated. These risks and uncertainties are described in the Company's Form 10-K filed with the Securities and Exchange Commission.

Additionally, I have also been asked to advise you that this conference is being recorded today, October 30, 2025, as the request of GSI Technology.

Hosting the call today is Lee-Lean Shu, the Company's Chairman, President and Chief Executive Officer. With him are Douglas Schirle, Chief Financial Officer; and Didier Lasserre, Vice President of Sales.

I would now like to turn the call over to Mr. Shu. Please go ahead, sir.

**Lee-Lean Shu**

Good afternoon, everyone, and thank you for joining us today. Let me start by highlighting two recent and important events for GSI.

First, we announced the research paper published by Cornell University in mid-October. The paper verified that our Gemini-I chip performs on par with NVIDIA's A6000 on certain AI tasks, while consuming roughly 98% less energy. This paper validates the disruptive potential of our compute-in-memory design, particularly for the near-term commercialization of Gemini-II. With 8 times the memory and 10 times the performance of Gemini-I, Gemini-II is positioned to deliver superior processing at a fraction of the power when compared to existing solutions.

This brings me to my second point. The market quickly recognized the significance of our compute-in-memory validation with the Cornell paper. Building on the momentum from the paper's findings, we closed a $50 million equity financing. We are now deploying that capital to accelerate execution across our hardware and software buildouts, making this a pivotal period for GSI's growth.

Post-funding, we are working on two initiatives in parallel:
First, we have begun the work to acquire the necessary IP for Plato, which will allow us to start hardware development. This IP provides critical connection integration to support broader system interfaces and prototyping for future customer applications. To accelerate Plato's time-to-market and capture market opportunities sooner, we are building out our existing Gemini-II team.
Second, to expedite the build-out of our Gemini-II software solutions and applications, we are investing in all the layers that make the platform more accessible and flexible for developers. These software tools are essential for customers integrating Gemini-II hardware into AI and signal processing workflows, particularly in edge and defense applications where efficiency and low power provide a competitive advantage.

Looking ahead, our initiatives for calendar year 2026 are centered on converting proof-of-concept (or POC) projects into commercial customers and expanding those relationships into larger production programs. Didier will provide an update on where those efforts stand today.

To sum up, our post-funding initiatives will be targeted and disciplined. We are rapidly moving forward with the Plato hardware design and software development, ramping up our Gemini software ecosystem, and strengthening ties with key defense and government partners in our POC and Small Business Innovation Research (or SBIR) programs. These actions position GSI to turn technical progress into commercial success in high-value edge and defense applications, such as drones, military vehicles, satellites, and other use cases, and ride the wave of AI compute innovation.

Now I hand the call over to Didi, who will provide more details on this topic and discuss our business development and the sales activities. Please go ahead, Didi.

**Didier Lasserre**

Thank you, Lee-Lean.

Let me expand on the topics that Lee-Lean just highlighted. We continue to advance our ongoing projects, including our SBIR and POC engagements with potential customers.

Recently, Gemini-II has been approved for prototyping by the offshore defense contractor to whom we shipped a board and software to a few months back. This POC focuses on synthetic aperture radar, or SAR, applications for drones and other edge systems. What's exciting here is that our solution delivers the required performance while maintaining an extremely low power profile—around 15 watts—making it ideal for compact, energy-constrained environments. For added context on just how competitive this solution is, an incandescent lightbulb uses about four times more wattage than our solution.

We are also involved in a joint POC, involving two defense organizations and a drone integration partner. This Gemini-II project combines the Yolo model we developed with multimodal large language model processing at the edge, specifically targeting time-to-first-token (TTFT), a key performance metric for drones. Along with our partner, we successfully demonstrated the end-to-end application to the end customer. Gemini-II outperformed the competing solution, particularly in how quickly the model produces its first response. We are now optimizing the algorithm and expect to publish initial benchmark results before year-end, with a fully optimized version available in the first half of 2026. This algorithm would be for defense applications such as drones, satellites, and other military vehicles. Gemini-II is a central part of our near-term commercialization roadmap, and we are encouraged by the customer engagement and technical validation that it is receiving.

Turning to our Plato program, we are embarking on the journey toward a major milestone—the tape-out of the Plato chip in early calendar 2027. Over the next year or so, we plan to actively engage several strategic partners for Plato, who could provide funding and collaborate on testing and prototyping early versions of the chip. Their involvement would also support the development of software libraries and APIs, ensuring that Plato becomes a versatile, scalable solution across multiple markets, starting with defense. In military and defense applications, the APU's high-performance and low power capabilities provide unique advantages, and Plato will further enhance critical functions such as SAR imaging, object recognition, GPS-denied navigation, and data fusion for drones and military vehicles—delivering real-time tactical capabilities in compact mobile systems.

Plato's design builds directly on the foundation of Gemini-II. To accelerate time-to-market, we are acquiring "building block" IP that allows us to focus on differentiation rather than reinventing core components. Strategic partners would play a critical role not just in meeting our ambitious timeline but in shaping the chip's capabilities, validating its performance in real-world applications, and guiding future enhancements. Their technical collaboration and early adoption would position us to deliver a highly optimized, field-tested solution, strengthening our long-term leadership in specialized AI compute architectures well beyond the immediate financial support.

And lastly, a comment on our SBIR work. We recently received an extension of $751K to one of our Space Development Agency contracts, which includes additional funding for radiation-hardened beam testing of Gemini-II. The goal of this testing is to evaluate the robustness of the current Gemini-II commercial chip for possible use in satellite [or][and?] aerospace applications. While it's too early to confirm specific designations, we see this as a significant opportunity.

Let me now switch to the second quarter's customer and product breakdown. By revenue, I am referring to net revenue in the following comments:

In the second quarter of fiscal 2026, sales to KYEC were $802,000, or 12.5% of revenues, compared to $650,000, or 14.3% of revenues, in the same period a year ago and $267,000, or 4.3% of revenues, in the prior quarter. Sales to Nokia were $200,000, or 3.1% of revenues compared to $812,000, or 17.8% in the same period a year ago and $536,000, or 8.5% of revenues in the prior quarter. Sales to Cadence Design Systems were $1.4 million, or 21.6% of revenues, compared to zero in the same period last year, and $1.5 million, or 23.9% of revenues, in the prior quarter.

Military/defense sales were 28.9% of second quarter shipments compared to 40.2% of shipments in the comparable period a year ago and 19.1% of shipments in the prior quarter. SigmaQuad sales were 50.1% of second quarter shipments in fiscal 2026 compared to 38.6% in the second quarter of fiscal 2025 and 62.5% in the prior quarter.
I'd now like to hand the call over to Doug. Please go ahead, Doug.

**Douglas Schirle**

Thank you, Didier.

The Company reported net revenues of $6.4 million for the second quarter of fiscal 2026, compared to $4.6 million for the second quarter of fiscal 2025 and $6.3 million for the first quarter of fiscal 2026. Revenue growth in the quarter was driven by strong market momentum for leading SRAM solutions.

Gross margin was 54.8% in the second quarter of fiscal 2026 compared to 38.6% in the year-ago second quarter of fiscal 2025 and 58.1% in the preceding first quarter of fiscal 2026. The decrease in gross margin in the second quarter of 2026 was primarily due to a change in the product mix.

Total operating expenses in the second quarter of fiscal 2026 were $6.7 million, compared to $7.3 million in the second quarter of fiscal 2025 and $5.8 million in the prior quarter.

- Research and development expenses were $3.8 million, compared to $4.8 million in the prior-year period and $3.1 million in the prior quarter. The increase in research and development spending compared to the prior quarter is primarily due to changes in the levels of stock-based compensation expense and amounts of government funding received under SBIRs in each quarter recorded as an offset to research and development expense.
- Selling, general and administrative expenses were $3.0 million in the quarter ended September 30, 2025, compared to $2.6 million in the prior-year quarter and $2.7 million in the previous quarter.

Second quarter fiscal 2026 operating loss was $(3.2) million compared to an operating loss of $(5.6) million in the prior-year period and an operating loss of $(2.2) million in the prior quarter.

Second quarter fiscal 2026 net loss included interest and other income of $43,000 and a tax provision of $41,000, compared to $149,000 in interest and other income and a tax provision of $23,000 for the same period a year ago. In the preceding first quarter, net loss included interest and other income of $13,000 and a tax provision of $54,000.

Net loss in the second quarter of fiscal 2026 was $(3.2) million, or $(0.11) per diluted share, compared to net loss of $(2.2) million, or $(0.08) per diluted share, for the first quarter of fiscal 2026. For the prior year second fiscal quarter of 2025, net loss was $(5.5) million, or $(0.21) per diluted share.

Total second quarter pre-tax stock-based compensation expense was $856,000 compared to $663,000 in the comparable quarter a year ago and $341,000 in the prior quarter.

At September 30, 2025, the Company had $25.3 million in cash and cash equivalents compared to $13.4 million at March 31, 2025.

- Working capital was $26.8 million as of September 30, 2025 versus $16.4 million at March 31, 2025.
- Stockholders' equity as of September 30, 2025 was $38.6 million, compared to $28.2 million as of the fiscal year ended March 31, 2025.

Lastly, for the third quarter of fiscal year 2026, we expect net revenues in a range of $6.0 to $6.8 million and gross margin from 54% to 56%. We remain focused on disciplined execution to bring Gemini-II to market, advance our roadmap for Plato, and drive long-term shareholder value.

Operator, at this point, we'll open the call to Q&A.

**Operator**

Thank you. We will now be conducting a question-and-answer session. If you would like to ask a question, please press star, one on your telephone keypad. A confirmation tone will indicate your line is in the question queue. You may press star, two if you would like to withdraw your question from the queue. For participants using speaker equipment, it may be necessary to pick up your handset before pressing the star keys.

Our first question comes from Robert Christian, Private Investor. You may proceed with your question.

**Robert Christian**

First of all, I'd like to congratulate you on the Cornell verification. But I'd also like to know, have you done any work with the auto industry on autonomous vehicles?

**Didier Lasserre**

We have not yet. As we've talked about in past calls, we have limited resources, and that takes a tremendous effort for that market space. We're currently starting in the military defense arena, but we certainly believe our technology will adapt well in those areas. It will be a focus for us in the future, but not yet.

**Robert Christian**

Okay. Thank you.

**Operator**

Our next question comes from Mark Gadner, Private Investor. You may proceed with your question.

**Mark Gadner**

Hey, guys. Good afternoon. I had a question on the $50 million placement you recently did, was that with a strategic investor? What sort of investor? Was there a holding period to that stock?

**Douglas Schirle**

No, it was just someone that was interested in the Company, wasn't strategic in any way. There is no required holding period for the shares.

**Mark Gadner**

Got it. Okay. Then just a follow-up to that. Have there been any inquiries from more strategic investors since the report came out from Cornell?

**Douglas Schirle**

Please remember that the Cornell paper was just announced last Monday, October 23. There are multiple opportunities we're looking at and parties that we're engaged with, and those conversations continue to be productive. As Didier said, we have limited resources, so we have been careful in targeting our resources to promising markets where we can deliver a differentiated solution. I think we have some very significant opportunities that he's already mentioned.

**Mark Gadner**

Got you. Okay. Thank you.

**Operator**

Our next question comes from David Zalkowitz with ISQ. You may proceed with your question.

**David Zalkowitz**

Yes. Is there any plan to have Cornell or another third party validate the Gemini II information, a different technology? I know the Cornell report was Gemini I. Is there a plan to do that similar type of analysis for Gemini II?

**Didier Lasserre**

Yes. You're absolutely correct. Cornell received this Gemini I board many years ago, and they've written a few other papers. This was a continuation of that original board. We are talking to them about getting a Gemini II board to them and also other researchers as well.

**David Zalkowitz**

Okay. Then you're working with the military. I didn't see anybody on the Board or senior management team that has real military defense experience. Is there any plans to beef up that area of the Management or the Board of Directors in order to target those applications?

**Douglas Schirle**

Yes. No, that hasn't come up as a discussion or topic on the Board. At this point, there are no plans to revise the Board. It doesn't mean that we won't in the future, if it makes sense, though.

**David Zalkowitz**

Okay. Then I saw you're developing your own large language model, which you're going to release some information on at the end of the year. Just curious why you wouldn't just use the plethora of large language models that are already out in the market and why spend resources developing your own?

**Lee-Lean Shu**

No, we are not developing our own large language model. We are working on the open source large language model like Gemma 3 12B.

**David Zalkowitz**

I'm just reading a press release. The press release says currently developing a multi-modal LLM that targets edge applications.

**Didier Lasserre**

Correct. Yes. Gemma-3 12B that's the model, and we're developing our algorithms to work with that model.

**David Zalkowitz**

Okay. Why would you do that as opposed to utilizing other LLMs are already developed?

**Didier Lasserre**

In this case, it was the definition from the POC that we're working on. As we talked about, there are two government entities that have approached us and a partner to do a POC, and that is the model that they requested.

**Lee-Lean Shu**

Yes. Also, there are certain aspects of the model, which support multi-modal well. They can support the image very well and in addition to the text. That's why they pick on this one.

**David Zalkowitz**

Okay. Thank you. Appreciate it.

**Operator**

Our next question comes from Christian Rug from CER Holdings. You may proceed with your question.

**Christian Rug**

Hi. I was wondering, how are you differentiating your APU versus GPU competitors in terms of power, latency and cost efficiency?

**Didier Lasserre**

That's a pretty broad question. If you look at the Cornell paper, that certainly hits on the power. The comparison was to an NVIDIA GPU and the use case they use, the performances were on par, but we were 98% less power. That certainly shows that. With the SAR algorithm that we've been talking about, certainly, our image creation time is faster at a lower power point as well. We've done benchmarking on certain use cases based off of input from customers on what they'd like to see. There are times where we beat them strictly on power. There's times we beat them strictly on performance. Well, I shouldn't say that, we've never lost to them on power. But there certainly are times that we have the advantage on both performance and lower power.

**Christian Rug**

Okay. Then my second question is, given the performance claims and potential of Gemini II APU, have you had any engagement or partnership discussions with larger semiconductor or AI-focused companies?

**Didier Lasserre**

Right now, we're focused on the customers at this point. We haven't had any discussions at least recently with other semiconductor companies.

**Christian Rug**

Okay. Then my last question is, how does the power factors play into building AI data centers at a large scale?

**Didier Lasserre**

We're focused on the edge right now. Everything we've talked about now is for the edge. Certainly, the data centers have a real power issue as well. There's no secret there. But what we've been focused on right now with Gemini II and certainly with the next-generation chip Plato will be at the edge. As we discussed with Gemini II, this project we did with this offshore defense contractor, we limited our chip to one of the four cores that are on the chip to get it down to 15 watts. If you look at Plato, depending on how it's used, can be as little as 4 watts and maximum 12 watts to 15 watts. We're really focused at the edge, not in the data center. But let’s not discount the edge. It will be a very large market, with many use cases.

**Christian Rug**

Okay. All right. Thank you.

**Operator**

Our next question comes from Michael Cooper, Private Investor. You may proceed with your question.

**Michael Cooper**

Good afternoon. Can you talk about the total addressable market that you're looking at over the next five years? Then how you expect that to ramp. I'm guessing you have a number of different scenarios, maybe a range of scenarios. You could give us a sense for how large this market is for these markets? I'm sure you're looking at various markets. Then what kind of price points your boards or chips go into products?

**Didier Lasserre**

Sure. It's a good question. Michael, I don't have the numbers in front of me. But certainly, there was a report very recently that was issued by one of the researcher analysts at Needham & Company that discussed the drone market specifically. I don't have it in front of me, but I want to say it was tens of billions at least market size. I want to say it might be larger than that when you include other edge processor systems. Certainly, it's a very, very large market. As we've discussed, we certainly feel with the power profile of our chips, along with some of the algorithm work that we're doing for, like Lee-Lean mentioned, the multimodal inputs, whether we take an image or text or voice in the future, along with the time to first token advantage that we have. We think that we're well positioned to address that market. That was question one.

The second question was—I think it was a two-part question you had. Pricing. Yes. Pricing, we'll give you generalities. But certainly, it's going to be priced differently by market, but it could be a few thousand dollars a board to $10,000 a board that contains the chip. Then the chip will sell, again, based on the market, but the chip could be $1,000 or more depending on the market and the volume.

**Michael Cooper**

You're working in gross margins in the 80-ish percent range?

**Didier Lasserre**

Yes, it will be above where we are corporately today. Again, it really depends on the market and how it's sold. It could be 60% to 80%. It really depends on how it's sold, whether it's in a board, in a server, whether it comes with software or not. There's a lot of different aspects that would move that margin needle.

**Michael Cooper**

Great. Thank you.

**Operator**

Our next question comes from Michael Roberts with Roberts Capital, who is rejoining us. You may proceed with your question.

**Michael Roberts**

Thank you. On capital deployment on the $50 million raise, can you give an idea of how that plans to be allocated, whether it's percentage or dollar amount, amongst the Gemini II completion, software development and the new Plato chip that you referenced?

**Lee-Lean Shu**

Yes. On the Plato because there's some fixed costs that we have to spend like IP costs and the mass tape-out costs. Those are fixed $50 million, $60 million, $70 million kind of range. The rest of them, I think they're probably pretty even between the Gemini II and the Plato, that's mostly engineering costs, the internal cost, and it will be distributed even inside the Company.

**Michael Roberts**

Evenly across. Okay. Thank you. In terms of then based on your cash runway now, what revenue or gross margin level do you expect to reach operating breakeven then?

**Douglas Schirle**

If you can assume, I don't know, 65% to 70% gross margin once we get into this. It's something that I need to take a look at. We're still putting our plans together in terms of hiring levels and so on. SoC teams or whoever we need for the chip development, additional software teams that we need for the software development. I don't have all those numbers yet to do a calculation.

**Michael Roberts**

Understood. But are there concrete milestones and dates then for the Gemini II in terms of expectation of pilot shipments or expected initial production orders?

**Didier Lasserre**

Yes. We will be doing some pilot shipments. We've done a couple already and we plan on doing more in the first half of 2026 calendar. This POC that hopefully, we'll be able to discuss a lot more in the upcoming months, depending on the schedules on that could give more substantial revenues in the back half of calendar 2026.

**Michael Roberts**

Noted. From the current evaluation customers now, has any purchase orders or letters of intent been provided yet?

**Didier Lasserre**

I'm sorry, could you repeat the question?

**Michael Roberts**

Yes. Have any of the evaluation customers provided any purchase orders or letters of intent yet against that production?

**Didier Lasserre**

They're still in their evaluation at this point. As we talked about, the Board that we sent along with the software to this offshore defense contractor, they have done a review and they've put us as what's called good acceptance in their system, which means it's passed and been accepted. Now we're going through the possible use cases. They have a couple of different divisions. Two of them, we think will be a good fit. One obviously is the SAR division. The other one is what they call their AI division. We're looking for practical applications that can then, like you say, turn into design wins and revenue. We're doing that with the customer today.

**Michael Roberts**

All right. Very helpful. One last question, then I'll let others proceed. Can you elaborate on that software stack maturity then, the compiler SDK model porting tools and when developers outside of GSI will have access? Going towards the ecosystem adoption of what we have?

**Lee-Lean Shu**

No, for the Gemini II, right now, we are developing the library and algorithm. Now after that, we will move on to the tool and the compiler work. We are developing this with the partner and the customer we have.

**Michael Roberts**

All right. Understood. Thank you very much.

**Didier Lasserre**

Thanks, Michael.

**Operator**

Our next question comes from Robert Christian, Private Investor. You may proceed with your question.

**Robert Christian**

Yes. Can you help me understand why the Company is not going after data centers in view of the environment impact with energy consumption and cooling? It seems like we're leaving a lot of money on the table, even if it was just licensed so others could use the technology.

**Didier Lasserre**

I'm not sure how long you've been following the Company, but we had talked about another potential road map product at the time we were calling it Gemini III, and that was going to be geared towards the data center. That one needed a different kind of partner and it needed a lot more funding. It would have required

a very aggressive process node and would have been much more expensive. Again, it was targeted for the data center.

In the meantime, we were getting way too much positive feedback and interest on the edge, and we were getting SBIR dollars, and there are other dollars, research dollars that we've submitted for to try and get, and it's all for the edge. The decision was made. We couldn't do both. It was one or the other at this point. We remain focused on the edge. Not to say with more influx of cash, we can't beef up the team and go after the data center, but it strategically made sense for us to remain at the edge for now. I understand that data centers get much of the press today, but don't overlook the edge market. They will be large and vast in use cases.

**Robert Christian**

Okay. But there's not a possibility, say, of NVIDIA or a Micron to come in and develop the chip and we get a percentage of it?

**Didier Lasserre**

Yes. That's certainly very possible. I can't say those discussions are happening, but we had some discussions in the past where that was kind of the model we were looking at. The answer is yes, we could do that. There's nothing in the hopper right now.

**Robert Christian**

Okay. Thank you very much.

**Operator**

Our next question comes from Marco Petroni with MG Capital. You may proceed with your question.

**Marco Petroni**

Yes. You guys just recently raised $47 million net, and you had $13 million last quarter, and the balance sheet shows only $25 million now. I was wondering where that money went, number one. Number two, going forward, what type of capital allocations do you need to build out the software team to do all this other stuff that we've been talking about?

**Douglas Schirle**

Well, the first answer is that, that transaction closed after the balance sheet date. It was an October transaction and the $27 million that you see is as of September 30. Then in terms of capital allocation, I think we answered a previous question where we're looking at some IP and other stuff that we need to purchase for Plato, and then we expect to split funding between software development and the Plato development.

**Marco Petroni**

How much cash do you have on hand currently?

**Douglas Schirle**

Well, take $$25 million balance sheet date plus we got another $47 million. That should give you a reasonable estimate.

**Marco Petroni**

All right. Thank you.

**Operator**

Our next question comes from Mohammed Alsousi from Scale. You can proceed with your question.

**Mohammed Alsousi**

Yes. I just want to know if you have attracted any interest from any potential new customers after the Cornell study on APU performance.

**Didier Lasserre**

The customers we've been talking to, we've been talking about this low-power advantage for some time, and we've done benchmarks on several applications with some of our customers. They're aware of that. In that respect, it's not a surprise to our customers we've been talking to that we have this low-power advantage. This just illustrated it for the rest of the public as a third-party validation of what we've been saying.

**Operator**

This now concludes our question-and-answer session. I would like to turn the floor back over to Mr. Lee-Lean Shu for closing comments.

**Lee-Lean Shu**

We look forward to seeing you at upcoming events and to your participation in the third quarter fiscal 2026 earnings call. Thank you.

**Operator**

Ladies and gentlemen, this concludes our conference for today. Thank you for your participation. You may disconnect your lines and have a wonderful day.