



**GSI Technology, Inc.**

**Second Quarter Fiscal 2024 Results Conference Call**

**October 26, 2023**

## C O R P O R A T E P A R T I C I P A N T S

**Lee-Lean Shu**, *Chairman, President and CEO*

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

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

## C O N F E R E N C E C A L L P A R T I C I P A N T S

**Brett Reiss**, *Janney Montgomery Scott*

**Orin Hirschman**, *AIGH Investment Partners*

**Luke Bohn**, *Private Investor*

**George Gasper**, *Private Investor*

## P R E S E N T A T I O N

### Operator

Welcome to GSI Technology's Second Quarter Fiscal 2024 Results Conference Call.

(Operator Instructions)

Before we begin today's call, the Company has requested that I read the following Safe Harbor statement. The matters discussed in this conference call 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 call is being recorded today, October 26, 2023, at 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 conference over to Mr. Shu. Please go ahead, sir.

### Lee-Lean Shu

Good afternoon, and thank you for joining us to review our second quarter fiscal 2024 financial results.

I am pleased to share that we have achieved two significant milestones since we reported Q1 2024 earnings. First was the successful launch of the alpha version of our Copperhead compiler suite, a Python-

based tool that can harness the capabilities of the Gemini APU. We anticipate releasing the general release of the Copperhead compiler suite in early 2024. A full production version will be available later in 2024. Currently, we have partners learning, using, and developing algorithms with our compiler, including SHREC, a consortium of Universities and companies developing next-generation tech for space, and academic research groups exploring the future of computing architecture like MIT and the University of California Riverside.

The second milestone was the completion of the Gemini-II tape-out, which we announced last week. As a result, we are on track to have the chip back in our hands early next calendar year and expect to begin sampling the device in the second half of 2024. We are targeting Gemini-II partners and customers in low-power data center expansion and enabling data center functions at the edge. Examples of edge applications would include advanced driver assistance systems and HPC in delivery drones, autonomous robots, unmanned aerial vehicles, and satellites.

Turning to our financial results for the second quarter, revenue of \$5.7 million was at the high end of our guidance. Also, it's worth noting that our second-quarter fiscal year 2023 revenue was boosted by inventory builds with several key customers, which presents a challenging year-over-year comparison.

In addition to advancing the tape-out of Gemini-II, a significant area of recent focus has centered around our ongoing engagements with a key hyperscale partner. I am delighted to report that these discussions are making notable progress. Through our constructive dialogues with this leading cloud computing provider, we have gleaned invaluable insights into the precise design specifications required for Gemini-III to align with their requirements.

This collaborative effort has enabled us to chart a roadmap while identifying potential partners who can bring the essential financial and engineering resources to the table for the successful development, manufacturing, and launch of Gemini-III. This evolution will leverage the incorporation of High bandwidth memory into the APU architecture, thereby harnessing the full potential of in-memory compute advantages.

Now, I'll hand the call over to Didier, who will discuss our business performance further.

Please go ahead, Didier.

**Didier Lasserre**

Thank you, Lee-Lean.

Following up on the projects I mentioned last quarter, throughout the second quarter of fiscal 2024, our team continued to pursue opportunities with Gemini-I to advance our customer engagement. Currently, GSI has a Fast Vector Search plug-in available that allows cloud vector search users to seamlessly add APU accelerated search to their major cloud-provided hosted workloads with minimal latency from GSI hosted data centers. This FVS plug-in provides accelerated approximate nearest neighbor search response times, enabling access to a large enterprise service at lower power, that also meets price points for small and medium-sized businesses, opening a larger market for the service.

Switching to our SAR opportunities, due to customer feedback, we have decided to launch our low-power and high-efficiency SAR processing as a SAAS offering, along with an on-premise version. Moreover, we favorably completed benchmarking on a customer's data set. We are now engaged to bring that service to market.

One target we are engaged with, a startup building satellites that can provide computational capabilities on satellites in space, has identified the Gemini APU as its preferred provider for computing solutions for space. The big differentiator, in addition to low power performance, is the product's radiation tolerance feature.

On that note, we will be conducting full radiation tolerant tests on Gemini-I next month. By this, I mean the full range of tests required for customers considering using Gemini-I in space. Some of the applications that a radiation tolerant Gemini-I would be ideally suited for are SAR, ATR, which is Automatic Target Recognition, and Computing in Space, to name a few.

This summer, we announced that GSI was awarded an SBIR contract to perform a feasibility study to adapt Gemini-II to perform computing at the edge in collaboration with the United States Air Force and Space Force.

We are currently working on a second SBIR based on the Gemini-II software development, which is very promising and potentially a similar financial award as the first win. We continue to file more SBIRs as they bring two key benefits: a source of revenue, and they create use cases within the government for future APU opportunities.

Let me switch now to customer and product breakdowns for the second quarter. In the second quarter of fiscal 2024, sales to Nokia were \$1.2 million, or 20.3% of net revenues, compared to \$1.2 million, or 13.6% of net revenues, in the same period a year ago and \$1.9 million, or 33.5% of net revenues in the prior quarter. Military/defense sales were 34.8% of second quarter shipments compared to 22.4% of shipments in the comparable period a year ago and 33.8% of shipments in the prior quarter. SigmaQuad sales were 55.8% of second quarter shipments compared to 58.1% in the second quarter of fiscal 2023 and 58.6% in the prior quarter.

I'd now like to hand the call over to Doug...please go ahead Doug.

### **Douglas Schirle**

Thank you Didier.

We reported a net loss of \$(4.1 million), or \$(0.16) per diluted share, on net revenues of \$5.7 million for the second quarter of fiscal 2024, compared to a net loss of \$(3.2 million), or \$(0.13) per diluted share, on net revenues of \$9.0 million for the second quarter of fiscal 2023 and a net loss of \$(5.1 million), or \$(0.21) per diluted share, on net revenues of \$5.6 million for the first quarter of fiscal 2024. Gross margin was 54.7% compared to 62.6% in the prior year period and 54.9% in the preceding first quarter. The changes in gross margin were primarily due to changes in product mix sold in the three periods.

Total operating expenses in the second quarter of fiscal 2024 were \$7.2 million, compared to \$8.8 million in the second quarter of fiscal 2023 and \$8.2 million in the prior quarter. Research and development expenses were \$4.7 million, compared to \$6.4 million in the prior year period and \$5.2 million in the prior quarter. Selling, general and administrative expenses were \$2.5 million in the quarter ended September 30, 2023, compared to \$2.4 million in the prior year quarter, and \$3.0 million in the previous quarter.

Second quarter fiscal 2024 operating loss was \$(4.1 million) compared to \$(3.2 million) in the prior year period and \$(5.1 million) in the prior quarter. Second quarter fiscal 2024 net loss included interest and other income, net of \$71,000 and a tax provision of \$33,000, compared to interest and other income, net of \$14,000 and a tax provision of \$37,000, for the same period a year ago. In the preceding first quarter, net loss included interest and other income, net of \$80,000 and a tax provision of \$51,000.

Total second quarter pre-tax stock-based compensation expense was \$676,000 compared to \$661,000 in the comparable quarter a year ago and \$820,000 in the prior quarter.

At September 30, 2023, the Company had \$25.3 million in cash, cash equivalents and short-term investments, compared to \$30.6 million in cash, cash equivalents, and short-term investments at March 31, 2023. Working capital was \$28.8 million as of September 30, 2023, versus \$34.7 million at March 31, 2023,

with no debt. Stockholders' equity as of September 30, 2023, was \$45.4 million compared to \$51.4 million as of the fiscal year ended March 31, 2023.

Given the current global economic environment, our current expectations for the upcoming third quarter are net revenues in a range of \$5.4 million to \$6.2 million, with gross margin of approximately 55% to 57%.

Operator, at this point, we would like to open the call to Q&A.

**Operator**

(Operator Instructions)

Our first question comes from Brett Reiss with Janney Montgomery Scott. Please go ahead.

**Brett Reiss**

Yes, hi. Thanks for taking the question. Until corporate initiatives gain traction, how—can you just comment on cash burn rates? How long do you think the cash will last without you having to do some sort of financing?

**Douglas Schirle**

Yes. We've taken a look at that. We typically take a look at it every quarter. If nothing improves, which we don't expect, we have cash that should last at least a couple of years at this point.

**Brett Reiss**

If I may just follow up, what do you think your burn will be for this following year?

**Douglas Schirle**

This year we'll be around \$13 million to \$14 million. Yes, we have one extraordinary expense coming up this fiscal year that requires about \$2.4 million cash outlay. That's for the mask set for the Gemini-II product that just taped out that Lee-Lean previously mentioned today. That's somewhat of an irregular occurrence that only happens every few years.

**Brett Reiss**

Okay. Thank you.

**Operator**

Thank you.

(Operator Instructions)

Our next question comes from the line of Orin Hirschman with AIGH Investment Partners. Please go ahead.

**Orin Hirschman**

Hi, how are you? Can you give us a little bit more color on—is the data center customer interested in Gemini-II for any purpose or only Gemini-III? Have they played with Gemini-I to understand it better? What is attracting them to the Gemini architecture? I have one follow-up.

**Didier Lasserre**

Okay, so let me start with that. With the hyperscaler, we're engaged with. We've had when—we've had ongoing discussions with and we've since opened up other discussions with other hyperscalers. Certainly, as Lee-Lean mentioned in his earlier comments, the information we gathered from our first engagement was instrumental to be able to shape our story and our offering for future hyperscale conversations. What they're—interested right now would be a next-generation, specifically, something we'll call Gemini-III for now. Because that one is going to really be geared towards GenAI or the LLM models.

With that said, as you know, these hyperscalers are very large companies. Certainly, Gemini-II is certainly not out of play with them, it just wouldn't be the solution specifically for the GenAI conversations we're having.

**Orin Hirschman**

In terms of other hyperscalers and data center customers, it sounds like there's some level of progress. How would you characterize it? Have people played with benchmarks yet for any of them? I think would be an important point to note if it's true if you've gotten that far.

**Didier Lasserre**

Correct. What we've done is we've shown obviously the technology with Gemini-I because silicon is here. Software is here. We can run true benchmarks. We've then obviously run expected benchmarks off of Gemini-II. Going back, and I just remember what part of your earlier question was, what's interesting about our technology to them? There's certainly a few areas. Obviously, the performance is important to them. The low power has certainly grabbed their eye. Also other features like the fact that we're a bit processor. This is very important to a lot of these folks because right now, if they're using a GPU, they're really locked into depending on the GPU 8-bit or 16-bit, 32-bit or 64-bit traffic patterns.

We're a bit processor. We don't care what your traffic looks like. It can be 4 bit. It can be 64 bit. It can be—make up a number of 1,000 bit and we can adapt to that on the fly and cycle to cycle. That flexibility is also very important to them.

**Orin Hirschman**

In real life, there's been a lot of talk about a very bit type of super processor in the literature. Obviously, no one's really had one commercially where it's programmable in that fashion. Are you programmable today in that fashion, number one? Number two is what are some of the applications where that becomes so enticing?

**Didier Lasserre**

You're talking about the bit processing?

**Orin Hirschman**

Yes.

**Didier Lasserre**

Yes. There's a lot of research going on right now. They're finding out that some efficiencies aren't coming in like the expected 16-bit and 32-bit. As I mentioned, 5-bit was one that came up recently. Candidly, I don't remember the exact application or the research on that but 5 bit was one of them that was important. One of the benefits with this bit processing is that as our customers told us, we're future-proof. What I mean is if there's an application that comes in that needs 5 bit, we're there. If there's one that needs 128-bit, 256 bit, we're there. We don't have to redesign our part to be able to address those needs because of the fact that we have 2 million bit processors on our Gemini-I. They can work, as I mentioned, as 1 bit or make up a number between 1 million and 2 million. We're future-proof from that respect which is obviously eye opening to them.

**Lee-Lean Shu**

Yes. Just to comment. Right now, the GenAI, the large language model is a very big model. It's a very, very memory-intensive application. They will lead to a very big memory. They will need a very fast memory. The effort—one of the effort is they try to reduce the data format, if you have less number of bit to calculate, then you can less memory, you require less memory throughout. Okay. We have (inaudible) processing capability which you can just mention 5 bit, 4 bit, 3 bit, 2 bit, we will be there. With this, we have more capability to keep up with the software innovation and that's what some of the customer likes about this issue.

**Orin Hirschman**

Just going back to benchmarking, have you done benchmarking for some of the data center and hyperscaler customers?

**Didier Lasserre**

Candidly, most of the benchmarking we've done have been based essentially off of POCs, where customers have come in and said, "This is what I'm using. Here's my data set. What do your numbers look like?" We've done real life case POCs. We've done that for one of our Israeli customers that we talked about and then also one of our recent SAR customers.

As far as broad market benchmarking, we haven't done a lot but we know that's critical and that's something that's on our agenda to start doing. It's just been a resource issue for us to be able to make that happen. But we understand it's important and that's what we're next on the list.

**Orin Hirschman**

I see. Two more questions, if I may, and then I'll let other people ask. In terms of proof-of-concept using customer data, have any of the hyperscalers given you data sets to actually show them what you can and can't do with it?

**Didier Lasserre**

No, these are other applications. These have been SAR applications and other applications. The hyperscalers know. Because as I mentioned, the discussions early on have been revolving around this next-generation device. We've given them—we have given the benchmarking what we anticipate we could do, absolutely but they're not true benchmarks, they're calculated benchmarking.

**Orin Hirschman**

Okay. Last question, just on the—you mentioned as you're using the Software-as-a-Service for other people's data and you had some early tests going back a while ago and that, that needed to be revamped for some additional APIs. Where are you up to with any of the potential customers, offering it as a plug-in or anything like that?

**Didier Lasserre**

Yes. There's two areas that we're focusing the SaaS on, fast vector search and SAR. As I mentioned in my comments earlier, the SAR was after discussions with one of the folks that we're talking to in that area. We have an on-prem solution which for some of their customers is needed. Then obviously, one of the guys we're talking to wants to put this capability on a satellite so that would be a physical sale as well. But right now, most of their production they're doing is via SaaS over AWS using GPU instances. That's when we—after discussing with them, that's something that we looked at. That obviously makes sense for us to do.

We're opening up the SAR for that. The fast vector search we've been putting in place. We're still doing some benchmarking internally for some of the factor search and databases, guys like (inaudible) so that's an ongoing project for us.

**Orin Hirschman**

In terms of trying to get that fast factor search as a plug-in so that customers could actually play with it through AWS or something like that, any progress there?

**Didier Lasserre**

I'm trying to understand what I can say there. We've had some discussions recently with a very large data center company to basically integrate easier into their system right now. Those discussions have been happening over the last couple of weeks. Then certainly, once those get worked out, then we'll be able to have a smoother offering right now. Because as you know, right now, it's—we're basically—they're offloading the searches off to our data center that we have put into a facility called CoreSite and so we're trying to streamline that offering.

**Orin Hirschman**

Okay. Thanks so much.

**Didier Lasserre**

Thanks, Orin.

**Operator**

Thank you.

(Operator Instructions)

Our next question comes from the line Luke Bohn, an investor. Please go ahead.

**Luke Bohn**

Hi. To zoom out from parsing applications and the prospects for various partnerships, could we assume that when Gemini-I hits the scene and you all are able to offer that to be tooled by various users, will that



outperform a large swath of the current high-performance computer offerings that are on the market, including the leaders of NVIDIA, AMD, Samsung, even Google, custom chips? Is that safe to say?

**Didier Lasserre**

Okay. I'm trying to follow your question here. First of all, we do have Gemini-I already; so I'm assuming you were referring to Gemini-II.

**Orin Hirschman**

Gemini-II, correct.

**Didier Lasserre**

Yes. Gemini-II and then you also mentioned a custom chip. Gemini-II is not a custom chip. Gemini-II is going to be a standard offering that we're going to have. Yes, obviously, we've done some calculated benchmarkings which have looked very favorable. Once we have the chip in hand, as Lee-Lean said, we'll have it in hand early next year and it's probably going to take at least a (inaudible) to really get something we can do benchmarking. It will be around summertime before we can do true benchmarking. Yes, we anticipate them to be very favorable over what we have and what's in the market space.

Now—but I do want to add to the comment, you said custom chip. One of the things, and we've mentioned this in the past, that we're also looking to do is obviously offer IP sales as well because we do understand we have a unique technology. But we also understand that certainly, we'll start with the hyperscalers. They do a lot of their own custom silicon. One of the conversations that we're having with the hyperscalers besides talking about Gemini-III for the GenAI is we're also talking about IP sales as well for any custom basics that they're doing on their own. You asked a lot of questions in there. I'm hoping I hit all the topics.

**Douglas Schirle**

Yes, that does answer the two implications there. That Gemini-2 will be a compelling competitor and depending on how you're able to benchmark or how much tooling and rendering it requires to VIP or via that chip (phon) directly but that it will, yes, essentially, at 16-nanometers, just be competing toe-to-toe or clearly outcompeting a lot of these 5-nanometer leading high-performance chips.

**Didier Lasserre**

If you recall, Gemini-I, the benchmarking we've done in the past have been very positive against what's out there today. Now with Gemini-II, as we discussed in the past, it has 10x of performance and 8x the L1 memory on chip. It's certainly going to be something that's going to be very compelling for the market.

**Luke Bohn**

Excellent. All right. Yes, eager for it. Thank you. That's all I have.

**Didier Lasserre**

Thanks, Luke.

**Operator**

Thank you. Our next question comes from George Gasper, an Investor. Please go ahead.

**George Gasper**

Yes, thank you. Could you go into a little additional information for us on the military and defense applications and the interconnect that the Company has had in Israel over the years, in fact, initially the whole Gemini approach came out of there? Considering what's going on in space, is there broader applications for the advancements that you're making going into space? Is what's happening with Israel right now in terms of this situation in the Middle East, is there—are you seeing anything happen that could create a more military-related effort?

**Didier Lasserre**

Okay. I'm going to try and get all those questions now. Let's start with the military. Yes, there are a lot of different possible use cases with the military which is also why we've had some successes recently with these SBIRs that we've been filing. Some of the interest that we've seen have—from, we'll call from the (inaudible) area is certainly SAR for sure. Object detection is another. Automatic target recognition, another. We've also been, for imaging as well, for change detection, just—and also one of the customers that we're engaged with, in fact, they have two of our servers on loan that they're doing demonstrations with is they want to start putting data centers in space, just a standard data center application in space; so that's another one.

As far as—I'm not sure with the Israel sites. As you say, this original technology was an acquisition we made and the company was out of Israel. I'm not sure what the question was revolving around that acquisition.

**George Gasper**

Well, basically, just as a reference, I mentioned it. Considering what's going on in the Middle East right now and the connections that the Company has had with Israel, I would think that there would be an attempt to try to get some applications moving forward to give Israel more recognition and how they're trying to monitor things.

**Didier Lasserre**

Yes. We've had some discussions, obviously, with (inaudible) that we've talked about in the past and other entities there on several applications. I won't go into exact detail but there's certainly—I'm sure we'll—there will—certain areas that we've talked about with them will most likely get accelerated now because of obviously what's happening now, but right now, we're not having those discussions right now. There's more pressing issues in Israel right now.

**George Gasper**

I see. Okay. Well, it looks like you're about to make some generally better movements forward with Gemini-II going to III. Hopefully, this really starts to turn the Company around in terms of its revenue volume and would expect that by the end of this fiscal year there's some real momentum going forward. Can you comment on it?

**Didier Lasserre**

Yes. There's certainly momentum. As time goes on, we hope to announce more traction. By the end of the fiscal year, any revenue that comes will be coming from Gemini-I. As we've discussed, Gemini-II, we're not going to be even sampling it until summer of next year; so that falls into fiscal '25.

**George Gasper**

Got you. Okay, thank you.

**Didier Lasserre**

Thanks, George.

**Operator**

Thank you. As there are no further questions, I would now hand the conference over to the Management for their closing remarks.

**Lee-Lean Shu**

Thank you, all, for joining us. We look forward to speaking with you again when we report our third quarter fiscal 2024 results. Thanks.

**Operator**

Thank you. The conference of GSI Technology has now concluded. Thank you for your participation. You may now disconnect your lines.