VIAVID
communications simplified.

GSI Technology, Inc.

Fourth Quarter and Fiscal 2023 Results Conference Call

May 16, 2023

CORPORATE PARTICIPANTS

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

Didier Lasserre, *Vice President, Sales*

Douglas Schirle, *Chief Financial Officer*

CONFERENCE CALL PARTICIPANTS

Nick Doyle, *Needham & Company*

Jeffrey Bernstein, *Cowen & Company*

Orin Hirschman, *AIGH Investment Partners*

Luke, *Private Investor*

PRESENTATION

Operator

Welcome to the GSI Technology's Fourth Quarter and Fiscal 2023 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 & Exchange Commission.

Additionally, I have also been asked to advise that this conference call is being recorded today, May 16, 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 day, everyone, and welcome to our fiscal fourth quarter and full year 2023 financial results earnings call.

The 2023 fiscal year was filled with many positive developments, new partnerships, and progress toward achieving our goals. We also experienced setbacks and unforeseen delays on several fronts with APU. We learned a lot during the year about the addressable market Gemini-I can reasonably pursue with our team, given our limited resources. However, we recently have made significant strides in leveraging third-party resources to help identify users, resellers, and OEMs. These resources are proving valuable in helping us identify opportunities for capturing revenue and increasing awareness of the APU's tremendous capabilities.

We have also sharpened our focus for Gemini-I to leverage our resources and prioritize near-term opportunities, such as synthetic aperture radar, or SAR, and satellites, where we have a superior solution. We understand these markets and know whom we can support and help with our offering.

Another focus application for Gemini-I is vector search engines, where our APU plug-in has demonstrated enhanced performance. To this end, we have dedicated more resources and prioritized the target customers that have expressed interest in leveraging our solution. Our data science team has been busy working on a SaaS search project with one leading provider, and we plan to pivot to other players in the space once we have met our deliverables with the first partner.

Looking ahead on our roadmap, we will build upon the work we are doing today in future APU versions to address large language model or LLM for natural language processing. Vector search engines are a fundamental part of ChatGPT architecture and essentially function as the memory for ChatGPT. Large language models use deep neural networks, such as transformers, to learn billions or trillions of words and produce text. This is another reason that vector search is an appropriate focus application with the APU.

Additionally, we are improving our Searchium AI SaaS platform to support our go-to-market strategy for search. We intend to use this tool to develop more potential partnerships like an Open AI plugin integration that we recently launched, and with other open source, decentralized search engines that use machine learning algorithms and vector search engines. The increasing size and complexity of enterprise data sets and the proliferation of AI in all aspects of business, are driving rapid growth in these search engines. Encouraged by the positive reception of our APU plug-in by several key players, we are optimistic about generating modest revenue from this market in the fiscal year 2024.

For both of the Gemini-I focus applications I have just mentioned, SAR and Fast Vector Search, we have set specific revenue goals that we aim to achieve this fiscal year. Our L-Python compiler stack has progressed in the past quarter. Our L-Python compiler stack is designed to offer Python's development advantages while delivering C's high performance without compromising either. Although our current focus applications do not require a compiler, we have a beta version in use currently and are on track to release a production-ready version later this year. L-Python will demystify the APU for any Python or C-developer.

I am excited to announce that we are on track to complete the tape-out for Gemini-II by this summer and evaluate the first silicon chip by the end of calendar year 2023. We aim to bring this solution to market in the second half of 2024. Gemini-II's design will provide significant performance enhancements with reduced power consumption and latency. These features will expand the future addressable market for the APU to larger markets such as Edge applications, Fast Vector Search, LLM, and advanced driver assistance systems, or ADAS, the last one being a vertical we would go after with a strategic partner rather than directly.

Gemini-II is built with TSMC 16 nanometer process. The chip contains six megabytes of associative memory connected to 100 megabytes distributed SRAM with 45 terabytes per second bandwidth or 15 times the memory bandwidth of the state-of-the-art parallel processor for AI. This is more than four times the processing power and eight times of memory density compared to Gemini-I. The Gemini APU is built

with bit processing, which allows fully flexible data format operation, an inherent advantage versus other parallel processors.

Gemini-II is the complete package that includes a DDR4 controller and external interfaces for PCIe Gen4 by 16 and PCIe Gen4 by 4. This integrated solution allows Gemini-II to be used in affordable edge applications while still providing significant processing capabilities. In simpler terms, Gemini-II combines different components together, allowing it to be used in less expensive devices, while still being powerful enough to handle demanding tasks at the edge of a network. Put another way, Gemini-II brings data center capabilities to the edge. This means that computationally intense applications can be done locally. For example, ADAS, delivery drones, autonomous robots, and UAV or unmanned aerial vehicles, and satellites.

Another application for Gemini-II would be IoT edge applications in critical infrastructure or processes requiring a reliable and efficient operation, for example, wind farms, to mitigate failure modes that can lead to significant financial losses or operational disruptions. Gemini-II's combination of high processing power, large built-in memory with tremendous bandwidth, and low-cost solution, provides a best-in-class solution for AI applications like Fast Vector Search, a growing market driven by the proliferation of big data and the need for fast and accurate processing.

Recently, we were granted a new patent for Gemini-II's in-memory full adder, which is a basic building block to allow Gemini-II to perform high processing power. We are thrilled to announce that we are currently in very early-stage discussions with a top cloud service provider to explore how Gemini-II's foundational architecture could deliver performance advantages. Just this year, we have seen the disruptive impact of large language models that understand and generate human-like language, like ChatGPT, Microsoft Bing and Google's Bard.

As the boundaries of Natural Language Processing continue to be pushed, we envision abundant opportunities in this market for Gemini-II and future versions of the APU. We believe that we have merely scratched the surface of the potential of large language models and the transformative impact they can have across numerous fields. Large language models' attention memory requires very large built-in memory and very large memory bandwidth on-chip.

The state-of-the-art GPU solutions have built-in 3D memory to address the high-capacity memory requirement but has poor memory bandwidth for adequate memory access. The limitation is going to get worse as large language models are progressing. Gemini chip architecture has inherently large memory bandwidth, is a natural migration to add 3D memory for the next generation Gemini chip to address the large memory requirement. This substantial improvement potentially translates into orders of magnitude better performance.

As a result, we would be strongly positioned to compete effectively in the rapidly expanding AI market, standing ahead of the industry's leading competitors. Our resources and teams are focused on applications where we have a high probability of generating revenue to capitalize on Gemini-I's capabilities. As we bring Gemini-II to market, we will be more experienced in approaching target customers and creating new revenue streams. We are formulating our roadmap for the APU, which holds tremendous potential. With future versions, the APU has the capability to cater to much larger markets, and the potential opportunities are quite promising.

In parallel with our Board of Directors, we are actively exploring various options to create Shareholder value. I remain fully committed to driving sustained growth and innovation in the year ahead. Thank you for your support and for joining us today. We look forward to updating you on our progress in the coming quarters.

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.

As Lee-Lean stated, we have sharpened our focus on a few near-term APU revenue opportunities. In addition, we have strengthened our team with a top data science contractor whose primary job is to accelerate the development of our plugin solution for the high-performance search engine platforms that Lee-Lean mentioned. We have also begun working with a company that offers custom, embedded AI solutions for high-speed computing using Gemini-I and Gemini-II.

Another critical development to improve our market access for the APU has been adding distributors. We are pleased to announce that we have added a new distributor for our radiation hard and tolerant SRAM, but also our hardened APU, for the European market. In addition to our partnerships and focus on near-term opportunities, we plan to build a platform to enable us to pursue licensing opportunities. This is in the very early stages, and we have work to do before we formally approach potential strategic partners.

That said, we have had a few preliminary conversations on determining what is required to integrate Gemini into another platform. This would allow us to identify the specific performance benefits for a partner's applications to ensure effective communication of the problem we solve in their system or solution.

We recently demoed Gemini-I for a private company specializing in SAR satellite technology. They provide high-resolution Earth observation imagery to government and commercial customers for disaster response, infrastructure monitoring, and national security applications. The satellites are designed to provide flexible, on-demand imaging capabilities that customers can access worldwide. They recently provided the datasets to conduct comparison benchmarks on the Gemini-I, and we are commencing the process of running those benchmarks. SAR is one market we anticipate that we can generate modest revenue with Gemini-I this fiscal year.

GSI was recently awarded a Phase 1 Small Business Innovation Research, also known as SBIR. SBIR is a United States government program that supports small business R&D projects that could be commercialized for specific government needs. For this contract, we will collaborate with the Air and Space Force, to address the problem of edge computing in space with Gemini-I. Gemini I is already radiation-tolerant, making it particularly well-suited for space force missions. This contract is a milestone for GSI Technology, as it will showcase the APU's capabilities for the military and other government agencies and provide great references for similar applications. We have submitted other proposals for a direct-to-Phase 2 project and other SBIR proposals are in the pipeline.

On that note, we received verbal confirmation just this morning that we have been awarded a Research and Development contract, which could be worth up to $1.25 million, to integrate GSI's next-generation Gemini-II for Air and Space Force Missions applications. This revenue would be recognized as milestones are achieved, and a typical timeframe is 18 months to two years. Once the agreement has been finalized and executed, we will issue a press release with the full details.

Let me switch now to the customer and product breakdown for the fourth quarter. In the fourth quarter of fiscal 2023, sales to Nokia were $1.2 million, or 21.8% of net revenues, compared to $2 million, or 23.1% of net revenues, in the same period a year ago, and $1.3 million, or 20% of net revenues, in the prior quarter.

Military defense sales were 44.2% of fourth quarter shipments compared to 22.3% of shipments in the comparable period a year ago, and 26.2% of shipments in the prior quarter. SigmaQuad sales were 46.3% of fourth quarter shipments compared to 47.6% in the fourth quarter of fiscal 2022, and 45.2% in the prior quarter.

I'd now like to hand the call over to Doug. Go ahead, Doug.

Douglas Schirle

Thank you, Didier.

I will start with the fourth quarter results summary, followed by a review of the full year fiscal 2023 results. GSI reported a net loss of $4 million, or $0.16 per diluted share, on net revenues of $5.4 million for the fourth quarter of fiscal 2023, compared to a net loss of $3 million, or $0.12 per diluted share, on net revenues of $8.7 million for the fourth quarter of fiscal 2022, and a net loss of $4.8 million, or $0.20 per diluted share, on net revenues of $6.4 million for the third quarter of fiscal 2023.

Gross margin was 55.9% in the fourth quarter of fiscal 2023 compared to 58.6% in the prior-year period, and 57.5% in the preceding third quarter. The decrease in gross margin in the fourth quarter of 2023 was primarily due to the effect of lower revenue on the fixed costs in our cost of goods sold. Total operating expenses in the fourth quarter of fiscal 2023 were $6.9 million, compared to $8.1 million in the fourth quarter of fiscal 2022, and $8.5 million in the prior quarter.

Research and development expenses were $5 million, compared to $6.5 million in the prior-year period, and $5.5 million in the prior quarter. Selling, general, and administrative expenses were $1.9 million in the quarter ended March 31, 2023, compared to $1.5 million in the prior-year quarter, and $3 million in the previous quarter.

Fourth quarter fiscal 2023 operating loss was $3.9 million compared to an operating loss of $2.9 million in the prior year period, and an operating loss of $4.8 million in the prior quarter. Fourth quarter fiscal 2023 net loss included interest and other income of $101,000 and a tax provision of $191,000, compared to $47,000 in interest and other expense and a tax provision of $21,000 for the same period a year ago. In the preceding third quarter, net loss included interest and other income of $61,000 and a tax provision of $84,000.

Total fourth quarter pretax stock-based compensation expense was $515,000 compared to $714,000 in the comparable quarter a year ago, and $654,000 in the prior quarter. For the fiscal year ended March 31, 2023, the Company reported a net loss of $16 million, or $0.65 per diluted share, on net revenues of $29.7 million, compared to a net loss of $16.4 million, or $0.67 per diluted share, on net revenues of $33.4 million in the fiscal year ended March 31, 2022.

Gross margin for fiscal 2023 was 59.6%, compared to 55.5% in the prior year. The increase in gross margin was primarily due to product mix. Total operating expenses were $33.5 million in fiscal 2023, compared to $34.9 million in fiscal 2022. Research and development expenses were $23.6 million, compared to $24.7 million in the prior fiscal year. Selling, general, and administrative expenses were $9.9 million, compared to $10.2 million in fiscal 2022. The decline in research and development expenses was primarily due to the cost reduction measures announced by the Company in November 2022.

The operating loss for fiscal 2023 was $15.8 million compared to an operating loss of $16.4 million in the prior year. The fiscal 2023 net loss included interest and other income of $202,000 and a tax provision of $372,000, compared to $60,000 interest and other expense and a tax benefit of $45,000 a year ago.

At March 31, 2023, the Company had $30.6 million in cash, cash equivalents, and short-term investments with no long-term investments, compared to $44 million in cash, cash equivalents, and short-term investments of $3.3 million in long-term investments at March 31, 2022.

Working capital was $34.7 million as of March 31, 2023, versus $45.8 million at March 31, 2022, with no debt. Stockholders' equity as of March 31, 2023, was $51.4 million compared to $64.5 million as of the fiscal year ended March 31, 2022.

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

Operator

Thank you. (Operator Instructions) The first question comes from the line of Raji Gill with Needham. Please proceed with your question.

Nick Doyle

Hi. This is Nick Doyle on for Raji Gill. Two questions on Gemini-II. Are all the costs related to the tape-out and then the testing volume production contemplated in your current outlook? Then could you expand on what kind of applications you're seeing traction in with that Gemini-II, specifically anything in ADAS and then using the large language models? Thanks.

Douglas Schirle

Yes. In terms of R&D spending, yes, most of what we're spending today is on Gemini-II. We have the hardware team here in Sunnyvale and the software team in Israel. There will be a tape-out in the first half of fiscal 2024 for Gemini-II. It will run probably about $2.5 million. Other than that, the R&D expenses should be similar to what we've seen in the most recent quarter.

Didier Lasserre

Regarding the applications, you cut out, were you talking Gemini-I or Gemini-II?

Nick Doyle

Gemini-II, please.

Didier Lasserre

Yes. Gemini-II, ADAS, as we discussed in the conversation before, is something we want to address. We most likely will use a partner to do that. As far as the large language models, as we discussed, we certainly feel that the Gemini technology, the advantage in the technology certainly will be applicable there. Whether it's a start with Gemini-II or if it's also customized with the Gemini-III is to be determined.

Nick Doyle

Okay. That makes sense. Then just a quick one. Did you say if there is a timeline? Is there a timeline for the rad hard road map for the product you mentioned in the EU?

Didier Lasserre

The Rad-Hard and Rad-Tolerant SRAMs are available today. We have done some testing; it's been at least a year and a half. We did the testing on the APU, Gemini-I specifically, came back very favorable.

The beam was a little bit off that da, it was limited the test we could do. We are actually going to do the full complement of radiation testing in the second half of this year so we have all the data requirements for the folks that will be sending it into space. Officially, the APU will be Rad-Tolerant sometime by the end of this year.

Nick Doyle

Excellent. Thank you.

Operator

The next question comes from the line of Jeff Bernstein with TD Cowen. Please proceed with your question.

Jeffrey Bernstein

Hi, guys. Just a couple of questions for me. Just wanted to make sure I heard right. You brought on a consultant that's helping target applications for Gemini-I. Is that right?

Didier Lasserre

They're specifically helping us write the interfaces for some of the fast vector search platforms that are out there.

Jeffrey Bernstein

Got you. Okay. Then you said there's a custom embedded AI solutions supplier, and that guy is going to now integrate Gemini-I into some high-performance compute solutions for clients. Am I getting that right?

Didier Lasserre

Partially. It's not limited to Gemini-I. It's Gemini-I and Gemini-II. They have a multitude of different potential applications ranging from SAR to sell-in applications to marine, search and rescue. There's a lot of different applications that they're looking at it for. Some of the cases, they'll be able to use essentially our Leda boards, but in many cases, they will be developing their own ultra-small boards for some of these applications that our boards are considered a little too big for those applications. It's a multitude of different applications, and it will be for both Gemini-I and Gemini-II.

Jeffrey Bernstein

Got you, okay. Then as far as the large language model kind of applications, I think there's two potentially. Correct me if I'm wrong. One is just to run queries as opposed to train and just run queries of these large matrices quickly and at low power. I guess the other one has to do with making training more efficient by being able to not redo matrices over and over again, as you do new learning. Is that right? Which are we talking about here today having some light of the day?

Lee-Lean Shu

Yes. Our primary target will be to the search, which is increase path. We are not on the training path, but if you can do search efficiently, you can help the training. We can do dual shop training or single shot training, which means you don't need even to train the data set. If you have a first query come in, we don't recognize, we can store into our memory chip. The second time, similar item come in that you can recognize it right away. That's very different from traditional training.

Traditional training, you have to run the whole model, whole data set, over again. That's very, very time consuming. If you can do dual shot training, you have the capability to do that, then you can save the training part tremendously.

Jeffrey Bernstein

That's great, thank you.

Operator

The next question comes from the line of Orin Hirschman with AIGH Investment Partners. Please proceed with your question.

Orin Hirschman

Hi how are you?

Lee-Lean Shu

Good.

Orin Hirschman

One of the things that the Gemini architecture, in memory processing architecture is very good at, which really wasn't a tremendous interest when you first introduced Gemini, was this natural language processing. Over this time, the whole world has changed, and you've got things like ChatGPT and other similar types of NLP situations where it actually exactly fits in to what you do best.

It sounded like from one of the prior comments on the last question, that you're actually having code and drivers written to be able to optimize the use of Gemini-I and certainly Gemini-II for this application. I would think that one of the simple applications where you could sell a lot of boards is just simply on the acceleration where everybody is having difficulty using GPUs, because this is not where a GPU signed on an AI side in terms of NLP in order to accelerate something like ChatGPT.

Lee-Lean Shu

What's the question again?

Didier Lasserre

Yes. What question?

Orin Hirschman

The question is in fact, is that a priority in terms of what you're working on to be able to introduce your own acceleration boards to do it with partners or is in fact, a great application? It sounds like certainly, so far on the call, that it's a great application for the Gemini APU.

Lee-Lean Shu

Okay. I think I discussed this on my statement. The biggest challenge for the large language model, are two-pronged. First one, you need a very large memory. The second one, you need a very high-bandwidth

memory. Those are two very difficult things to achieve. I think today in the market that nobody has this solution. It's a good solution.

Just as I mentioned, we do have very exciting discussion with a—we call it a larger cloud service provider, and to see how we can help from our Gemini foundational architecture to see how we can help to move this thing forward. We already have a very, very good memory bandwidth. That's why I mentioned in my statement, I say we are 15 times memory bandwidth over today's state-of-the-art GPU or payroll processers. That's our inherent architecture. If we can add this one to the high memory capacity, then (inaudible) which nobody in the market can provide. Now we're very excited. We try to explore this advantage we have and see where we can go from here.

Orin Hirschman

Any ideas when you will—the coding to the interface to be able to demo the type of acceleration gains that we're talking about with something like a ChatGPT or something like that, so customers can actually see some type of benchmarking even with Gemini-I and maybe a simulation until Gemini-II is ready. When will that code be ready? That's what you mentioned you're working on.

Lee-Lean Shu

Yes. The OpenAI, they have a plug-in. Basically, you can put your software to plug in to the main machine and then you can utilize the existing model and then do the plug in. Right now, we are working on it, Gemini-I presently and the Gemini-II to follow on. That at least—second, you can extrapolate from how well those are working and extrapolate to the future.

Orin Hirschman

Any idea when we might see some benchmark in coming months?

Lee-Lean Shu

How was it? Maybe a quarter or two, we will have something to tell you guys.

Orin Hirschman

Okay. Just a related question but even more futuristic. There's talk of doing something similar to, let's say, and there are a number of projects. In fact, even you had an early project with MUVE, M-U-V-E to be able to show off what you can do in terms of visuals as well? I guess my question is taking that same natural language processing and doing it on a visual level is beyond belief in terms of computationally intensive, but also well suited for what you guys do. Is anybody talking about doing anything like that?

Obviously, you did that early demo, which impressed a lot of people. Obviously, that's even a step beyond what almost what people have dreamed of today, but you can't do that using current architecture. Any thoughts on that from a futuristic perspective and will that need a Gemini-III, or can that be done in Gemini-II? Then one last follow-up question?

Lee-Lean Shu

You're asking whether—how we want to do in the future generation. Is that…

Orin Hirschman

No. Specifically, the more important part to me is just in terms of incredible visual search capabilities, almost like NLP search capabilities on visuals. You did that impressive early demo with MUVE and then some other experimental projects. People all over the world are starting to do experimental projects on massive amount of visual data.

Any more thoughts as to—that's obviously very suitable or uniquely suitable for what you do versus just GPUs for that matter. Any other interesting projects like that MUVE project? I know it's a bit futuristic, but has anybody done more in terms of that massive type of visual search, comparative visual search using NLP for visual search, using Gemini?

Lee-Lean Shu

Yes. We look at—with our—just I mentioned with our partner, we extend the architectural advantage of Gemini architecture. We look at the one workload, if we have enough memory, we will be 10 times faster than any solution exists today. That's why we are saying, hey, we have this inherent advantage there, but the thing is we don't have enough built-in memory for that. If we can combine for the future road map, if we can put enough memory into it, that's why you are looking for all the magnetic performance better than existing solutions.

Orin Hirschman

On that note, a closing question. Just in terms of what nanometer geometry is being used for Gemini-I, Gemini-II and what you're thinking for Gemini-III. Obviously, that will affect what you just discussed in terms of the ability to pack in memory et cetera. If you can kind of tell us more about that? Then just one follow-up, and that's it from me? Thank you so much.

Lee-Lean Shu

Yes. Today, we—Gemini-I is 28 nanometers, and the Gemini-II is 30 nanometers. If we look at the future, today's sales, our GPU is a 4 nanometer. If we look at future and then we do 5 nanometer and then we build in the 3D memory in there, because the only way you can get a high-capacity memory with a reasonable framework is 3D memory. If we put in the 3D memory with a 5 nanometer, we'll be order of magnitude better.

Orin Hirschman

This is the follow-up question, with understanding that in terms of Gemini-III, but knowing that Gemini-II is going to be the platform coming up here shortly, the key platform, in terms of your ability to accelerate NLP again, non-visual, forget about that futuristic question. But here today in terms of accelerating NLP applications and ChatGPT, et cetera, is Gemini-II got enough in it so that your competitive flash even superior on that type of application to leading edge GPU, optimized GPU like Hopper-style GPU. Have you passed that with Gemini-II? The question only is, can you leapfrog it even further? That's my last question? Thank you so much.

Lee-Lean Shu

As I mentioned there are two things, big memory capacity or big memory bandwidth. We have one of them. If any workload can fit into our chip, we will be the best solution out there. There are many, many cases like that. Even the ChatGPT, it doesn't have to be a humongous dataset. Okay. It can be a smaller dataset. The dataset can fit into our chip. We will be the number one in the market.

Orin Hirschman

Okay great, thank you so much.

Operator

(Operator Instructions) The next question is from the line of Luke Boin (phon), Private Investor. Please proceed with your question.

Luke

Hi, good to be back. Hope you all are well. Very exciting announcement and development. Great to hear all the comprehensive layout there. Just for really kind of minor clarifications and going a little bit broader with the near-term potential. Wondering if your Amazon Web Services server offering is capable of fielding, say, just a broader range of companies, potential end use cases that could more or less play around with your service without having to go through more complex processes of embedding or other integration processes, just plug and play?

What you can do for their applications, especially thinking about vector search, but also rich data, like was mentioned, maybe for metaverse, maybe for advanced registration, things like that. Wondering how you're seeing the potential to expand Amazon Web Services or a similar offering on say, Azure or other clouds, and especially how that would relate to an earlier rollout of Gemini-II from your own facility, your own servers on the cloud?

Didier Lasserre

As we've discussed in the past, we've started the integration with Open Search. That's ongoing. We have already set up our own servers for that. We have some here in our Sunnyvale facility, some in our Israeli facility and then we also have some at an off-site facility that's directly across the street from AWS West, and it's directly connected. We've had that in place with the Gemini-I. Over time, obviously, we would migrate those to Gemini-II.

Those are in place. We do have some SAR demos that people can run off of those remotely. It's not set up yet to be able to do load your own data. It's the data sets are already in there, which you can run. We're not at the point yet where you can enter your own data, at least not larger datasets. That is certainly the direction we're going. We're just not quite there yet.

Luke

Do you have a timeline on when you would be able to roll out those interactive features and capacities?

Didier Lasserre

We're shooting for this year. Some of the examples you brought up is going to be—we're going to get some help from this data science contractor that we have on board now. It's something we're trying to rollout second half of this year.

Luke

Excellent. All right. That's all I have. Pipeline flooded (phon). Appreciate it, yes.

Didier Lasserre

Thanks Luke.

Operator

The next question is a follow-up from the line of Jeff Bernstein with TD Cowen.

Jeffrey Bernstein

Hi, yes. I just wanted to see if you can give us an update on the Elta SAR application, and what's going on there?

Didier Lasserre

Yes. As you recall, we did the POC with them, and it was a very broad POC. It could be used for different vehicles or vessels. It could be used at a lot of multitudes of heights from 100 meters to much, much higher, obviously into space. The initial program they're looking at for us was just a single laptop, I guess you could call it that. They had already been using a GPU. They're using the GPU still for that program. There's a form program that they're looking at us for now. We're going through that process with them.

It won't be another POC, because we've already done one, but it will be a kind of a bit of a different project than what we are working on with them. It will still be under SAR, and it will still be the same algorithm. It should be a simple integration.

Jeffrey Bernstein

Okay. Then just wondering about waiting to hopefully get some space providence on the Rad-Hard SRAM and wondering if you guys have any visibility now on when that launch might happen or is it permanently scrubbed?

Didier Lasserre

No, it's not permanently scrubbed. We follow-up—I get your frustration because I'm with you on this one. It's not scrubbed. There were multiple programs that they—when I say they, there was a few as we talked, defense contractors were using it. There have been a couple of the programs that have been scrubbed, but the larger ones we're looking at have not been scrubbed. They're certainly still out there. They've just been pushing out the launch dates. We're just not getting a good feel for exactly when the next launch is going to be.

Originally, we know they were delayed, because they couldn't get some critical components. Now it's just a matter of getting them to actually do it. The answer is, we're still optimistic about it, it's just the timing is elusive for us on when it's actually going to happen.

Jeffrey Bernstein

Can the European distributor kind of do anything on the Rad-Hard piece or are they stuck with just doing Rad-Tolerant until you get that space province or is it a different approach in Europe?

Didier Lasserre

No, they're definitely going to be going after everything. The folks that we've already sent parts to that were looking to get heritage. It's really just a heritage part. The heritage just is the signal to the world that says your parts have been launched into space, and they work. It's an additional check mark and a box for a lot of these folks. It doesn't change the fact that our parts are already internally qualified to work up there.

We know they will work based off of the testing that we have done. This European distributor is going to be finding additional opportunities for us. The folks that we were looking to do the heritage for the short-term launches, those were U.S.-based companies. We have shipped some Rad-Tolerant and at least one Rad-Hard to a European customer, but they were not the ones we anticipated to get us the initial heritage.

Jeffrey Bernstein

Okay, all right, great. Then any update on some of the scientific applications as Weizmann Institute come back for more boards or any analogous type customers in pharma, med-tech, biotech, universities, et cetera?

Didier Lasserre

Universities, yes. We're candidly not spending a lot of time on that market. The revenue opportunities for the other markets we've discussed today are larger. We do have two universities that—let me think—yes, two. There are two different applications for two different universities that are looking at them for genomics. They will be running—they'll be essentially doing the algorithms and doing the write-up. Personally, we are not spending much efforts ourselves. We've already done a plug-in, specifically for the BIOVIA, Tanimoto.

It just doesn't make sense for us, based off of our limited resources, to spend more time developing more algorithms for more platforms. The revenue volumes there just aren't as great as they are in other markets we're addressing.

Jeffrey Bernstein

Make sense, make sense. Thanks.

Operator

There are no further questions at this time. I will now turn the presentation back to the hosts.

Lee-Lean Shu

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

Operator

That does conclude today's conference. We thank you for your participation and ask that you please disconnect your line.