

Everspin Technologies, Inc. (MRAM)
37th Annual Canaccord Genuity Growth Conference
August 10, 2017

Note: The transcripts are posted on Everspin's website for the reader's convenience and prepared by third parties. Readers should refer to the audio replays, when available, on this site for clarification and accuracy.

<<Matt Ramsay, Analyst, Canaccord Genuity Group Inc.>>

All right, I guess, so we'll go ahead and get started here, I know people will continue to trickle in the room for the next couple of minutes as the sessions get filtered through but. Good afternoon, my name is Matt Ramsay. I am the semiconductor analyst here at Canaccord Genuity.

Thank you very much for coming to this session here and for your attention, we're really excited to have Phill LoPresti and Jeff Winzeler from Everspin. Everspin is an exciting emerging memory company, that's developing its unique MRAM [Magnetoresistive Random Access Memory] technology for a host of applications. These guys can give you an update. They just reported their earnings this morning. So there's a lot of fresh things to talk about, but I think the long-term roadmap from an exciting new memory technology is right at the cusp of taking off.

And I'm sure the guys here will be able to tell you more about the business. But we'll save a few questions for the end as well and just refer you to the conference program there for disclosures about Canaccord's business with Everspin and make sure keep us all out of trouble.

So with that, thank you very guys for being here and enjoy.

<<Phillip LoPresti, President, Chief Executive Officer and Director>>

Thanks Matt, appreciate it.

Can everybody hear me okay? Excellent, great. [So just go and step in through the slide here, this is our Page 2 Safe Harbor statement. Sure you're all familiar with that. And so then go into Page 3, wanted to introduce if you're not familiar with Everspin Technologies.](#)

We went public in October last year and our primary mission out there is to provide MRAM products. These are memory solutions that offer the persistence that you normally would get in a non-volatile memory, but provides speed and endurance of a RAM or something similar to a working memory.

And so this has got a lot of meaningfulness especially in the enterprise storage space, which is what we selectively have targeted our next-generation of products, which we'll

talk about today. So memory, our MRAM memory is similar in the context of the non-volatile memory because it's persistent.

However, you'll notice some very big differences between the non-volatile memories that you're familiar with like NAND. And if you look at this chart [slide 4], you'll see in the bottom right that you have NAND and this is just a chart that's taking right speed and right cycles and charting against each other.

And you can see where the NAND traditional storage memory lays basically down in the slower rate speeds and the lower right cycles meaning less endurance or less durability. And then in the top left, you see the traditional working memories SRAM and DRAM. These write very fast, most likely symmetrical read and write speeds and they have a very high level of endurance or durability meaning they don't wear out over time as you would see with NAND.

And then you have MRAM and the emerging memories popping in here on this slide and you can see the MRAM products are first generation MRAM products are shown up there in the top very close in performance to SRAM and we call that our persistent SRAM technology. Those range in density from as small as 128 kbits up to 16 megabits.

And then just below the DRAM bubble on this chart, you see our ST-MRAM products. This is our generation three products it uses our perpendicular MTJ, magnetic tunnel junction. We've announced our 256-megabit DDR3 part in production. We'll talk a little bit about on one of the later slides about some of the highlights including the design wins there.

And then we also have our one gigabit DDR4 part, which we just started sampling recently to select customers also fitting in that bubble. So you can see how we fit more into the higher speed and higher durability memories like working memory whereas some of the other competing emerging memory technologies that are out there that you may be familiar with like 3D XPoint or Resistive RAM or Phase-Change RAM that bubble kind of lays out closer towards what NAND looks like.

In other words maybe a little faster than NAND and perhaps a little bit more durable than NAND. So why MRAM now, while in the industry MRAM has been around for quite some time. It was started getting developed into the mid '90s to late '90s. And our legacy comes out of Motorola, Freescale. We spun out of Freescale in 2008 and we've been advancing MRAM development for quite some time.

We've been shipping our first-generation product [Toggle] back from 2006, 2007 timeframe. We shipped over 70 million units to-date and it's been in this persistent space or lower density space. But MRAMs, why is it taking off now what's the differences. While first of all, there is more and more system requirements that require this high performance with persistence.

Also for us to be able to get into volume production with a production fab, it's hard to get customer traction if you don't have a viable production line and a stable supply of product. And so what we've been able to do is not only put a production line in place where we ship the 70 million units. It is a 200-millimeter MRAM line, but we also put in place a partnership with GLOBALFOUNDRIES for our 300 millimeter spin-torque part. So that box has been checked, very key to expanding the market and the use cases is the density of the part.

We've now launched 256-megabit and one-gigabit parts by advancing the density of the part you advance the application use cases and we've targeted those parts specifically at write caches and write buffers in solid-state drives, flash array products and also RAID type systems.

To be able to be positioned to attack a market you have to have good customer relations and having some contact with them to understand the type of product they want. We've had an outstanding relationship with all our first-generation customers, in which we sold quite a bit of product into the storage industry, which helped us put in place a very powerful ecosystem that allows customers to use these new MRAM chips that we've deployed. And such that we have MRAM ready SSD controllers, MRAM ready FPGAs and we have MRAM ready RAID controllers.

So when you put all those together and start rolling out these products you now have the ability to fill in a pipeline of customer opportunities, and give us the revenue potential that the Company has today.

On this slide [6], I'm highlighting the three various generations of products that we have, we're the only Company to-date that has put any MRAM product in production. We started that with our Generation 1, I have talked a little bit about that, we sell about 45% of our product of Generation 1 into the RAID system market. It's where we got our relationship with a lot of the storage companies. We also sell quite a bit into the industrial and automotive space. And we then targeted a product line specifically for storage applications and that was our [Spin Torque] Generation 2 and Generation 3.

We used our Generation 2, as the ecosystem, pipe cleaner or the ecosystem catalyst to get the controllers and to get FPGAs teed-up to be able to become MRAM ready and really what we wanted to tee everything up for was for a third-generation perpendicular Spin-Torque, because this has the most legs out of the technology and that it's scalable.

We've now produced a 64-megabit chip, on 90-nanometer. We moved to GLOBALFOUNDRIES on their 300-millimeter plant. And we've produced 256-megabit and we have been able to quickly parlay that into a one gigabit 28-nanometer chip, which we're now sampling and that shows again the ability for this technology to scale.

I highlighted earlier, that the higher the density that you can offer, it gives you more market opportunity [slide 7] to pursue and that's clearly the case here. We believe when we have our one gigabit and 256-megabit in production it actually gives us the

opportunity to address up to 1.2 billion for the spin-torque portion of our business. And we still have our first generation persistent SRAM products that address about \$400 or \$500 million opportunity.

This is a good slide [slide 8] to kind of show all the places where MRAM is being used today. And also where it's expected to be used in the future, so 70 million units, we have over 600 customers, over 100 different products SKUs that have been out in the industry for quite some time, we're in a wide variety of applications in the industrial space.

We've been providing product to companies like Siemens and companies like Schneider and Mitsubishi for quite some time. In the automotive industry, we've had good success in the infotainment area, with our first-generation product. And we have a lot of use cases that were being looked at, that include many of the bullets that you see underneath here.

We also have a lot of interest now coming out of the assisted driving and ADAS area where they are actually not only interested in our first-generation products. But they have use cases where they need higher density high-bandwidth memories that would actually take our spin-torque technology into some of those use cases.

And then there is the enterprise storage area, where our first generation has been a journal memory for the RAID systems for quite some time with the two big industry leaders that are there. They now Broadcom and Microsemi but those started out as Adaptec PMC-Sierra and LSI and we are now going further into the storage systems by getting into the write caches and write buffers by introducing our spin-torque products.

And you can see on slide number 9, where we are targeting our use cases for our spin-torque parts. Our 256-megabit, we feel the entry point for that is in the SSD market including the PCIe based cards and the U.2 form factor, 2.5-inch form factor drives.

This is a use case where 256-megabit is a good entry point for that and also we expect our one gigabit to get traction into those use cases as well. And then moving to the right, you can see the RAID opportunity, we will see continue to provide solutions for the general memory but we also now have the ability to solve the write cache. And the key element that you see across all three of these is, if you start with the RAID opportunity, write caches for RAID use DRAM but DRAM isn't persistent. So to make it persistent, you have to add a bill of materials to it.

And you can see in that case it's a dongle or tethered battery, which can add significant cost and impact the reliability and durability of that product. If you move over towards the your SSDs, you could see all these small components on one there, on the bottom left, you can see that those yellow components are all capacitors that's actually a board that is all capacitors both sides of it's populated with capacitors. And it's attached to a board that looks similar to the one above it. And the only thing that that board is doing is providing capacitance to hold a residual charge, so that power were to be lost, the data that's in flight that's traveling or not yet to the media has enough energy on the card in the system to complete the write to the media. So you don't lose the data.

And then all the way on the far right, we see future opportunities as our density scales to one gig and above where we can start addressing some of the NVDIMM opportunities that are growing out in the market.

So besides the obvious bill of materials opportunity that we have and I want to characterize that clearly because we're not competing against memory by itself. We're not trying to replace DRAM only in these caches and buffers. We're actually replacing DRAM plus the bill of materials that has to get strapped around it to make it look like MRAM and be persistent.

And in those cases, that affords us the ability to sell our part at a higher ASP, but when you start making these high-density form factor drives like you're going to start seeing more and more of like this actual case of a 10-terabyte drive there are other advantages. If you're trying to squeeze 10-terabytes of a high performance NAND into a small U.2 form factor you generate a lot of heat.

If you can find a way to squeeze enough capacitors in there you have to worry about the thermal considerations that are inside that form factor, which have a tendency of deteriorating the life of those components like batteries and super caps. But more importantly the compelling event here is that because our MRAM is instantly write-protected when you write to it, it allows SSD manufacturers to eliminate a portion of their over provisioning that they have to put into these drives.

And in this example [slide 10] about half of the over provisioning can come out and as you do this math calculation, I won't go through it but with a 10-terabyte drive you typically have to put a little bit over 13-terabytes of NAND in there, that's 3 terabytes that you get really no value for. They're really in there to provide the durability of the drive to last for three or five years.

There is a cost associated to that and you can see that if you are able to remove half of that like this customer has pointed out to us that they can do, you can eliminate about a terabyte and a half of NAND. And you can now save somewhere between \$400 to \$500 depending on the price of NAND that you're buying. So you could see not only you got a better form factor, better reliability you don't have to worry about thermal considerations, you also get a cost savings in the amount of NAND that you have to procure.

Moving to page 11, you can see our nvNitro series, we had a real exciting week this week in many ways at the Flash Memory Summit. We made quite a bit of announcements I'll highlight some of them on the next several pages. But these are this is just a single slide that shows you the picture of our nvNitro is squarely in the middle there you can see our half height, half length card, this is a one gigabyte or two gigabyte version of our card that's going to be built out, we're using our 256-megabit chips.

You could, see down to the bottom left is the U.2 form factor, we also have that on display and we were running demos for both of those at the Flash Memory Summit. And

we've also provided many of our customers various DIMMs for them to evaluate our technology. And we actually even have a customer design in and we've worked with customers on their M.2 form factor. The key thing of our nvNITRO, it is the fastest slowest latency drive that you can get out there, we call it an nvNITRO Accelerator and it's very exciting especially when you start building to densities even at 8- and 16-gigabyte was a huge, huge benefit to be able to do 1.5 million IOPS sustained in an very low latency of six microseconds end-to-end, it's performance that you don't see in any other drives out there today.

Another part of our business model [slide 12] was to be able to build our manufacturing ecosystem. We aligned with the GLOBALFOUNDRIES as our manufacturing partner for 300 millimeters. In part of that relationship with them it included a licensing agreement where GLOBALFOUNDRIES has taken a license for our MRAM technology and they announced back in September 2016 that their first process that they're going to deploy it on is their 22-nanometer fully depleted SOI process, they call it 22FDX.

With that I'm happy to say recently they had an article on EE Times where they've said essentially put themselves open for business for MRAM, embedded MRAM type of projects, they are offering PDK kits now that are allowing customers to start playing and designing and I'm expecting in the near future to start seeing some design win opportunities come from that. Of course, Everspin in this case, in this business model it's a royalty business model. So when GLOBALFOUNDRIES succeeds in winning designs and shipping material to their customers we'll then get a royalty from that business in the future.

Global Operations [slide 13], we have sales network across the globe, leading distributors, regional distributors as well as global distributors. We work with all the big packaging and assembly house and test houses.

And then some of the recent highlights [slide 14] that I wanted to hit, our Generation 1 Toggle, bookings have been very strong, very good design activity. In our earnings call today, we expressed a 33% increase in design activity over last quarter.

Generation 3 [Spin Torque] as I highlighted, we launched four nvNITRO Storage Accelerator cards, they are based on 256-megabit parts of ours. We also announced a partnership with SMART Modular. They are launching their own card, the NVM Express card, that's also using our 256-megabit. This was one of the four design wins that we announced earlier in an earnings call.

Now moving into the production cycle, we began sampling our pre-production engineering samples of our one gigabit part DDR-4 to select customers. And we also have expanded our ecosystem, we've worked with Mellanox Passport Program demonstrating a high-speed memory transfer over fabric. We've also expanded our relationship with Xilinx FPGA capabilities, which we use on our nvNITRO card. And as I highlighted in the embedded MRAM space, GLOBALFOUNDRIES is progressing

really well we're really excited about their opportunities now going out there selling into the low power IoT space using their 22FDX.

With that, I'm going to let Jeff bring you up to date on our financial highlights [slide 15].

<<Jeff Winzeler, Chief Financial Officer>>

Yes. So a couple of things, we announced in our earnings call today that we had a very strong Q2. We announced revenue which was a Company record of \$8.9 million for the quarter that was at the high-end of our guidance, up from \$7.9 million in the previous quarter. So pretty good revenue growth but more importantly, it came with really good margin too, our gross in Q2 was 65%, which is also a Company record. We typically target our financial model of between 40% and 52% gross margin. So due to a couple of onetime event as well as some strong performance we had a lot of upside to that.

Other important things that we did in Q2 is the company had an existing debt facility, we refinanced that in the second quarter. And what that did is it added working capital to our balance sheet and just as importantly it's an interest-only loan for the first year, so we reduced our one-year debt service by about \$3 million. So that helps preserve cash and operating capital for the company.

Our 2016 year-over-year results were good, we grew our revenue 2.1% overall but our underlying MRAM technology grew more than that. Our gross profit dollars were up, our gross margin as a percentage was up and we actually brought our OpEx down in our year-over-year results. We went public in 2016 so our balance sheet was much stronger after that our October IPO, we raised \$40 million in the IPO and another \$5 million in a concurrent private placement at the same time. So very much strengthened the balance sheet to prepare ourselves to really bring out this third-generation technology that Phill has been talking about.

And so through that IPO and through some of the activities that we had last year, we're in a much better shape to take advantage of this Gen 3 MRAM technology that we're right in the middle of deploying.

<<Phillip LoPresti, President, Chief Executive Officer and Director>>

Thanks Jeff. So just let me conclude over here some of the highlights [slide 16]. Everspin's MRAM technology, only company so far to deliver commercially available MRAM solutions in three different generations over 70 million units in the first gen and now offering the highest density one gigabit product sampling and going into production next year.

Over 600 customers, over 70 million units shipped our application focuses is really key, is that we target these bill-of-material type of solutions so we're not playing in the commodity memory market. We have a strategic relationship with GLOBALFOUNDRIES, which is really key to allow us to accelerate development on the

leading-edge process nodes. But also to have an extremely viable production supply of product for our key customers in the enterprise space.

We have a substantial IP portfolio, which is important in this market. We have over 300 patents issued and over 150 patent applications pending. The ecosystem is pivotal for the market. We have great partners like Xilinx, like Microsemi, like Broadcom, these are MRAM ready type of controllers or FPGAs that can interface with our technology and allow customers to build SSDs and RAID systems. And we have a pipeline that continues to grow and expand and we're focused in and actually already doing business with many of the industry leaders in industrial automotive in the transportation space as well as storage.

All of this together, with this huge market opportunity that I highlighted in one of the earlier slides of about \$1.7 billion addressable market, provides an attractive long-term growth and margin profile as Jeff said in his previous slide. So I'll conclude there. There's still a couple of minutes left for questions, if there are any.

Q&A

<<Matt Ramsay, Analyst, Canaccord Genuity Group Inc.>>

So I might just start this, and thank you for the presentation. On your earnings call this morning, a number of things that were announced, I think a couple of them are important. One of which is our lead customer moving on to your third-generation 1 gigabit part more quickly because I think it's a great testament to the technology and second that roadmap moving towards production ready product earlier and then so that's one thing, can you talk about that a bit more of and then on the flip side there was a little bit of a push out in the expectations for how quickly the revenue right things left for this new technology. So maybe you can just talk through those new items a little bit more?

<A – Phillip LoPresti>: Sure, so let's address the one gigabit opportunity. Our timing of the one gigabit chip was we've always been saying that we would be sampling our products sometime this year. And we've made fairly good progress, we still have a little more work to go with that chip, but we were engaged with a customer that has an opportunity in this flash array product that they really needed to use an MRAM product to solve the problem that they had, which was they didn't have space for capacitors and they didn't want to put capacitors in because of thermal considerations. And that MRAM really offered them the best solution and because they were targeting a high-end, they really preferred the one gigabit solution. So they – we've worked real close with them to make them comfortable that we can achieve the timelines that are necessary for their project.

We're working very diligently, my teams talking to them multiple times per week and keeping very close. And so they've announced to us, we've had to get production capacity, get ready with our partner GLOBALFOUNDRIES that they're going to need material in Q4 of this year, which we're working to deliver to them which would allow

them to make pilot builds in Q1. And then their intention is to be in production with their system shipping it out to their customers in Q2 of next year.

Based on what our regional planning was, we didn't really anticipate a lot of revenue ramp for the one gigabit until probably the second half of 2018. So yes, as Matt's described it's a little bit earlier and it's an early design win. And it's kind of an atypical win, usually wins don't come until you've already delivered the production sample. And in this case, this customer has gone forward with an engineering sample, quality and we're moving ourselves along there.

On the other side of it, I think we've highlighted the 256-megabit did go in production on time in the March quarter at the end of Q1. And what we're seeing is the SMART Modular design is going in to production that's using 256-megabit. But what we highlighted is there's a possibility now that we're sampling the one gigabit chip. That other customers that may have been considering a 256 and we're not saying that anyone has decided that but there's a possibility now that time – based on the timing of things that some of these customers might migrate to a one gigabit solution.

And so all that would tend to do is if they were to migrate to a one gig solution that would put them as a production ready opportunity further on-time because we're not going to be production ready with the one gig until the end of Q1 next year, whereas the 256 is ready for production now.

So I think that might be the interpreted push or delay if that kind of situation occurs with any of the projects and yes, that would certainly move that production to be timed with the one gig as opposed with the 256.

<<Matt Ramsay, Analyst, Canaccord Genuity Group Inc.>>

Okay, well, thank you very much for your time.

<<Phillip LoPresti, President, Chief Executive Officer and Director>>

Okay. Thank you.

<<Jeff Winzeler, Chief Financial Officer>>

Thank you.