Everspin Technologies, Inc. (MRAM)
Stifel 2017 Technology, Internet & Media Conference
June 5, 2017

Note: This transcript has been posted on Everspin's website for the reader's convenience and prepared by a third party. Readers should refer to the audio replay, when available, on this site for clarification and accuracy. The company also posted slides for this conference and those can be accessed by clicking here.

This presentation contains "forward-looking statements" that involve risks, uncertainties and assumptions. If the risks or uncertainties materialize or the assumptions prove incorrect, our results may differ materially from those expressed or implied by such forward-looking statements. All statements other than statements of historical fact could be deemed forward-looking statements, including, but not limited to: any estimates of addressable market size and our ability to capture that market, market trends and market opportunities, customer growth, product availability, technology developments, or other future events; any statements about historical results that may suggest future trends for our business; any statements regarding our plans, strategies or objectives with respect to future operations or business performance; any statements regarding future economic conditions; and any statements of assumptions underlying any of the foregoing. These statements are based on estimates and information available to us at the time of this presentation and are not guarantees of future performance. Actual results could differ materially from our current expectations as a result of many factors, including, but not limited to: market adoption of our products; our limited operating history; our ability to raise capital; our history of losses; our rate of growth; our ability to predict customer demand for our existing and future products; our ability to hire, retain and motivate employees; the effects of competition, including price competition; technological, regulatory and legal developments; and developments in the economy and financial markets.

We assume no obligation, and do not intend, to update these forward-looking statements, except as required by law.

For more information about Everspin, please visit our website at www.everspin.com. .

<<Kevin Cassidy, Analyst, Stifel, Nicolaus & Co., Inc.>>

Good morning. And welcome to the Stifel 2017 Technology Conference. My name is Kevin Cassidy. I'm one of the semiconductor analysts at Stifel and it's my pleasure to introduce from Everspin, Phill LoPresti, the CEO and Jeff Winzeler, the CFO. Phill is going to start off with a Company presentation and then we'll open it up for Q&A. Phill?

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

Thank you, Kevin. Good morning everyone. I'd like to – okay, with that – properly there we go. All right of course, we have our first page here [Safe Harbor Statement].

So I'd like to introduce Everspin Technologies to you this morning. [Slide 3] Everspin is an emerging memory technology company. We spun-out of Freescale Semiconductor back in 2008. We produce what's called magnetoresistive random-access memory or MRAM. And what's unique about MRAM is that it provides the persistence of the non-volatile memory with the speed and endurance of a regular, high performance RAM or

SRAM. And this becomes extremely interesting for many applications today as cloud computing expands automated driving in vehicles, IoT. You're looking for high performance and non-volatility in the same type of memory.

This slide [Slide 4] kind of represents what most people are familiar with, up in the top left you have SRAM and DRAM. These are high performance memories, they write very fast, they have high endurance or their durability is very good. And then on the bottom right is what we call our storage, if you look even further to the right you would have hard drives but today NAND memory is considered mostly a storage memory.

Now you overlay this with our technology, the bright red spot is our Generation 1 product. You can see that it has virtually the same performance as SRAM and it is also highly durable and it writes very fast. Now you'll see DRAM there and right below it you'll see our Generation 3 ST-MRAM. Again very similar in performance of the DRAM but it has the high durability along with the very fast write speed.

We have also overlay, there's a lot of emerging memory technologies that are out there today. They are either phase-change memory, 3D XPoint, you may be familiar with, resistive memory, phase-change memory et cetera. These are again if you look at where they're positioned, they're much closer to being a storage memory as opposed to a working memory or performance memory. This is what uniquely positions Everspin in the market that we're focusing in which is high performance write caches and write buffers for storage systems with our Generation 3 products.

So why MRAM now [Slide 5], obviously, we've spun-out of Freescale 2008, MRAM was in development since the mid-late 90s at Motorola transferred over to Freescale. And of course Everspin has been working and developing products since we spun-out in 2008. We are the only company that's gone in production with MRAM. We've shipped over 70 million units of our MRAM products in for Generation 1, we service about 45% goes into RAID storage systems, 35% goes into industrial applications and the 20% into automotive.

And what we're seeing right now as I highlighted earlier is that there is a increasing demand for having high performance memory with persistence. And so it's timed very well with our ability to bring products to market now but there is actually a real big need for exactly the solution that MRAM provides. That's one reason. The second reason is, to this state there hasn't been a high volume supplier or a large manufacturing capability for MRAM.

Everspin's had a 200-millimeter fab that we've shipped over 70 million units already with, but to get to this large market segment in the storage space. We needed to move to 300-millimeter, we partnered with GLOBALFOUNDRIES. And GLOBALFOUNDRIES now has a 300-millimeter production line currently in Singapore that can obviously provide us with a very high quality and reputable supply of products. So that we can pursue this new storage market segment, which we believe is about $1.2 billion in 2018.

Secondly, everyone's familiar with memory that the higher the density the memory is the larger the application space or use cases are. [Slide 6] So in the past our first generation product was limited to about 16 megabits in the SRAM space. And now we've launched a 256 megabit DDR3 part, we announced that was in production starting in early March. And we are now getting ready to start sampling our 1 gigabit DDR4 product. So with those larger density products, we can now capture and target a larger market segment.

The other key element here is to get a product out and into production with customers you have to put in infrastructure and ecosystem. And so Everspin has been working to pursue this market for quite some time, we started working with DRAM controller companies for IP. We started working with SSD controllers and RAID controller companies to make sure that when they produce their next latest and greatest solid-state drive controller or RAID chip that they had an MRAM ready option in their DRAM controller.

So now that that device that let's say Microsemi or Marvell or Broadcom would sell has the ability perhaps to be able to do both DRAM and MRAM depending on what the customer wants to do and so by accumulating with those four items. The fifth item allowed us to now drive a design win pipeline, we've been sampling our 256 megabit product since August of 2016 and we've had a lot of customers interact and start working and looking for ways to use our technology.

So with all of that in place today, we're actually positioned now to be able to take off and challenge this larger market segment beyond our Generation 1 product. This gives you a pretty good summary of the three different generations. We use Generation 2 in between mostly as a ecosystem builder, we did have several customers go on production, we shipped some product there. But the primary goal of that product line was to allow us to enable ecosystem, the flash controllers, the RAID controllers and so on that we've done.

The third generation product on the bottom, which is our proprietary Perpendicular Spin Torque technology now allows us to go address this much larger market segment, which adds about $1.2 billion to our Generation 1 market segment, which was about $500 million. [Slide 7] So in total we service about $1.7 billion market opportunity in the 2018 time frame as depicted on this slide.

[Slide 8] This is a very busy slide but what I wanted to emphasize with this is that there are numerous applications for our Generation 1 product, 45% goes into enterprise storage, 20% into automotive, 35% into industrial. But we're in a slew of different applications as you can see here under each one of these categories and the primary benefit of our product are the key attributes I highlighted earlier, our durability or endurance, our fast write speed and the fact that we're nonvolatile.

So to kind of size up the opportunity for the Generation 3 product or Spin Torque 256 megabit and 1 gigabit products, this slide [Slide 9] kind of shows you solutions along the bottom that are presently used. We are primarily focusing and targeting this product that enterprise storage class products and you can see in the bottom you have on the left

bottom, you have a PCIe SSD, then you have a 2.5-inch form factor SSD and right next to it on the right is a RAID card.

One of the things that you see in common on all of these because they service enterprise, they are required to have a power fail-safe system. So that the contents of data that's being written through the DRAM write buffer is protected in the event that there's some kind of power fail. As you know the final destination is the storage media, it could be SSD, it could be hard drive. The problem is that writes very slow and so if you're trying to do a high performance system, you want to be able to write the data into a high speed buffer or cache.

And then allow the flash controller or hard disk controller to decide when to put it to the storage media. If power were to go out and it didn't complete that process then there's data sitting in the DRAM cache that would get lost. Enterprise customers don't like to lose data. So these power fail-safe systems are put in place. On the bottom left that's an additional card that slapped on top of the PCIe SSD similar to the photo above, that's a card that's full of capacitors.

And the only thing that those capacitors are doing are providing energy so that DRAM that's in that cache or buffer basically behaves like an MRAM. Okay, and if you go to the right that's a 2.5-inch form factor, you could see those small components there consuming virtually two-thirds of that board of just capacitors again doing the same thing. And on the next one, the RAID controller, this one's actually a little bit more interesting. This is a really small card, there's no place to put capacitors on the card or on the backside. So they tether with an umbilical cord either a battery pack or a capacitor pack.

And so in this case, in any of these cases it all varies quite differently on the level of reliability and what kind of assurances you want to give the end product. But you can be adding anywhere between $20 or $30 to $70 to $90 worth of other stuff to make the write cache that's built at a DRAM today nonvolatile. Now you can swap all of that for MRAM and eliminate the complexity, eliminate the reliability risks of this. And of course it gives us an opportunity to sell a value-added product because we're not now selling a part that competes with the DRAM price. But we're selling a part that competes with the bill of materials solution required to make DRAM plus the other stuff persistent.

We want to highlight another reason why there is a compelling thrust towards using MRAM in solid-state drives. This is actually from a real customer opportunity, we've talked about the batteries and the capacitors and so on, the reliability issues. One other complexity that you have when you use batteries and capacitors is thermal considerations. If you have a big lithium battery in your system, besides the complexity you have in your ship it's separately with the server or storage device. You've got to keep it cool because what happens with lithium-ion batteries is, sometimes they flame out.

[Slide 10] And so there's a lot of concerns about that so you get to remove all those issues but you actually have another huge advantage and this is why SSD manufacturers are very interested in the technology. There is an issue called write amplification when you

use NAND that means when you want to write something to NAND, you actually have to open a whole page. And then you transfer what's in the cache to the NAND. When you do that you're actually wearing out all the bits even though you're not really rewriting something new.

At that point that's called write amplification to avoid that problem, SSD manufacturers have to overprovision NAND. That means if you want to guarantee a 10 terabyte drive for three years or for five years depending on the life that you want. You can calculate how many writes are going to be done to the bits and you know they're going to wear out. So you put extra bits in there, so that you can have the controller use the extra bits the overprovision bits to allow the SSD to last to three or five years.

In a typical application, this is a very recent modern design brand-new real customer. They're trying to produce a 10 terabyte drive and you can see in this example that they basically have to put 30% more NAND in the drive. So that means a 10 terabyte drive if you're guaranteeing it for three years or five years whatever their spec is, you have at least 30% more NAND. In this case you put 13 plus terabytes in there. And you have to pay for that but the customer only gets to use 10, you only get value for the 10 because that's what you're marketing it as.

With MRAM because you can eliminate or reduce this write amplification issue, you can actually reduce the over-provisioning by at least 50% some of the SSD manufacturers believe it could actually be more than that. But if you just take a 50% example and you use a nominal price per gigabyte of NAND of $0.25. You can see the cost differential between a drive with MRAM and a drive without MRAM. In this example it's about a $400 to $500 savings in NAND.

Today NAND is in short supply so obviously if you don't want to put more NAND in a drive and you can use it for more drives like in this case every 10 drives, you build with MRAM you get one free as far as NAND. So this is another driving and compelling event that's occurring in the industry for moving towards ST-MRAM.

[Slide 11] Now getting our product into market, always it is a tricky thing. You have to put the ecosystem in place like we've done have to have controllers but you also designing with architects to get them to include your solution on their board or in their system. And when you do that obviously, there is a ramp time to do that design. There is also a qualification window to do those type of projects. What we found out was that by introducing our own card, in this case this is called nvNITRO Accelerator Card that's basically populated with all MRAM and an FPGA. In this case, this is a protoboard that we are already started sampling.

We were able to perform at the highest speeds of any SSD, 1.5 million IOP sustained and that's really important because most SSDs don't maintain that level over the life of the drive. Ours is a sustained performance, we have a very low latency of about 6 microseconds end-to-end. And we can provide this product in various densities starting at 1 gigabyte going up to 16 gigabytes when our 1 gigabit chip is available. These products

will also not only be launched in a PCIe form factor. As you see there in the middle but also to the left on the bottom, there is a U.2 form factor drive. These are pretty common SSD form factors and typically go into flash array products.

And on the bottom, you'll see the nvDIMM space, which is also an evolving space that you're hearing a lot about our product and obviously solve that problem in many cases. And then on the far right is an upcoming M.2 form factor for storage systems. We will be producing these nvNITRO products eventually in all of those form factors.

[Slide 12] Another key element if you're reading the paper – the press and listening to what's going on, Everspin started obviously quite some time ago working with GLOBALFOUNDRIES. To be our partner to mass produce our standalone products, we needed a reputable high quality, high capacity fab and GLOBALFOUNDRIES stepped up. But there was a reason that they were interested. The reason is, is that MRAM actually offers a real significant advantage for embedded memory cores for SoC designs.

So as you know GLOBALFOUNDRIES is the number two largest foundry. They produce a lot of product for various companies in the industry, many that you're familiar with. And the problem is the memories today that they're embedding are having difficulty scaling below 28-nanometer in particular embedded flash.

So we signed an agreement with GLOBALFOUNDRIES in late 2014, we started developing our chips but they licensed the technology to deploy and they announced in September 2016 that they're going to use embedded MRAM in their FDX process, which is a fully depleted SOI technology, very low power process, very good for low power microcontrollers, IoT controllers et cetera.

So this is a huge market, you're also seeing competition now for GLOBALFOUNDRIES, Samsung has announced that they're also going to be doing embedded MRAM. TSMC has alluded to them as well having embedded MRAM. So we actually feel pretty much like we validated our position, which we started back in 2014. This was going to be a very significant use case for MRAM in the future.

[Slide 13] So the Company has been established for quite some time, we have global operations, sales network across the globe. We have regional distributors and we also have global distributors many of the names you're familiar with. We run two different kinds of scenarios in one case for 300-millimeter the spin-torque products or fabless. GLOBALFOUNDRIES provides us the CMOS and they also do the processing of our custom bit at the end. And in the case of 200, we do a semi-fabless where we buy wafers from the foundries and then process and then in our 200-millimeter fab.

[Slide 14] So just to hit some highlights, Generation 1 Toggle, we've been seeing now these based on our recent earnings call, very strong bookings. We are getting a lot of traction with automotive. Our Gen 3, we announced slightly ahead of our schedule to offer our production version of the 256 megabit. We've announced three design wins, one from a major OEM for flash array product. We started sampling nvNITRO cards, we're

really targeting that initially at the high frequency traders. And so we believe that's going to be a big opportunity for us there because of the low latency and performance of the device.

We also expect to be sampling our 1 gigabit later this year. So we're pursuing and continuing to drive aggressively, we have many customers that are interested in getting that part as well. We see the storage class memory industry now evolving and solidifying. So we've joined several industry consortiums including Gen Z, NVMe Workgroup, OpenCAPI, which is related to the Open Power that we were previously in the – that's the IBM Open Power group.

And we've also started trying to strengthen and reinforce our organization, we've brought in a new Head of Sales, Annie Flaig, who has quite a bit of experience in the enterprise server storage area. We've also appointed Pat Patla to our Senior Marketing role and we've added two new board members that have a lot of expertise and experience in this industry. One previously CTO of SanDisk and also former Chairman and CEO of Virident are now both on our board.

And Jeff, you want to talk a little bit about the financial highlights?

<<Jeff Winzeler, Chief Financial Officer>>

[Slide 15] Yeah. Just from a financial highlight perspective, the company went public last year. We raised $37.2 million in that IPO, we also did a concurrent private placement with a strategic called GigaDevice that was interested in investing in the company. So we got some additional funds through that. That really provided capitalization for the company to be able to go out and win some of these design wins that are so critical for our forward success.

In terms of year-over-year, our total revenue only grew 2.1% from 2015 to 2016. However, gross profit dollars increased by $1.1 million, our gross margin was up to 54%. The Company has historically had over 50% gross margin. We reduced our OpEx by $1.3 million year-over-year and obviously through the balance sheet in the IPO, our balance sheet is much stronger than it was in 2015. For Q1 we had a really good quarter $7.9 million in revenue, which was up from the prior quarter.

Again gross margin of 54% in Q1. And then from a cash perspective, we took the opportunity to refinance our debt line, which gave us another $5.7 million of additional capital to utilize for working capital purposes and it reduced our debt service for the next year again freeing up some cash.

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

[Slide 16] Good. So just to kind of wrap up here. We are the only Company that's offered and has MRAM, many type of MRAM product in production today. We've actually done

three generations of the technology over 70 million of our Gen 1 already out, a thousand plus customers out in the industry.

We really target application specific solutions, we're now trying to replace commodity DRAM or NAND. We're really trying to replace where those products are being used but another bill of materials is wrapped around them to make them appear persistent like our MRAM is natively.

We have a strategic relationship with GLOBALFOUNDRIES, which we believe will allow us to continue to accelerate development of new products like our 1 gigabit. But also out in time perhaps 2019 and beyond, we can start seeing revenue from the royalty stream through the embedded MRAM program.

We have a very strong IP portfolio, we are obviously believe in proliferating the technology and enabling other partners, including memory companies or other foundries in the future. We have a very solid ecosystem that we're continuously improving by joining those consortiums. It also allows us to be consistent with the way storage class memory is going to be used in future server and storage architectures. We are focusing in on our design win pipeline, we've got numerous customers that have been playing with our 256 megabit part since back in August of 2016. And so team in for sales organization and marketing group is really focused on converting all those opportunities or as many of them as possible into design wins and ultimately for production to allow our business to ramp.

And overall I believe the company provides a very attractive long-term growth and margin profile for investors to consider. And so that concludes the presentation portion and we have about seven minutes to talk about questions.

<<Kevin Cassidy, Analyst, Stifel, Nicolaus & Co., Inc.>>

Yeah, I'll open up to the audience first if there's any questions otherwise, I'll ask a few.

Q&A

<Q>: [Question Inaudible]

<Q – Kevin Cassidy>: I will repeat that question. Just for the webcast, I'll repeat the question it was around the competitive situation in embedded MRAM, just this morning in DigiTimes as a report of TSMC and Samsung, how did they suddenly now have MRAM?

<A – Phillip LoPresti>: So I joined the Company back in 2010 and I went out looking for a partner to do our 1 gigabit, 256 megabit, 300-millimeter stuff we talked to the Samsung, to TSMCs, UMC, GLOBALFOUNDRIES, Micron, Hynix. All of these companies have been working on MRAM for quite some time internally. And they're obviously everyone has a different deployment approach for it. It turns out that like we

believed and the reason our relationship with GLOBALFOUNDRIES is solidified is that there is a very strong advantage to deploy MRAM as an embedded memory in your foundry.

So the three very significant foundries out there TSMC, GLOBAL and Samsung obviously all are going to compete. Now GLOBALFOUNDRIES once they announced they were going to do embedded MRAM. There were internal programs at all these companies to try to figure out how to take advantage of the MRAM. And then they realized that embedded is the right path. So now they've obviously accelerated their programs to bring out that technology to their foundry customers as well.

We don't believe that that competes with our Generation 3 products right now because those are standalone products that we build. This is really in the foundry space servicing SoC customers that want to embed MRAM as a nonvolatile memory solution in their SoC.

<Q – Kevin Cassidy>: Thanks, Phill. Along those lines of competitions, what I think is the point that many investors miss is other MRAM solutions or SRAM facing or replacing SRAM where you have an architecture of DRAM in your Gen 3. Can you talk a little more about that what the Gen 1 addresses as far as the market and where Gen 3 is going?

<A – Phillip LoPresti>: Okay. So this actually dovetails into the embedded strengths as well. MRAM is a very versatile memory and we stay away from calling it a universal memory because it's not going to replace all the other memories that are out in the industry. There's going to be a place for the storage class type of devices like NAND and there's going to be places for where DRAM makes sense in 3D XPoint. They're all going to find some tier storage that makes sense. But in our technology MRAM has the versatility that you can actually configure it to look like an SRAM or you can configure it to look like a DRAM or you can configure it to look like Flash.

And in our Generation 1, we configured it to look like SRAM that's why our Generation 1 products, 128,000 bits to 16 megabits that's kind of in the serial interface and asynchronous SRAM interface space, those kind of parts. But to get to a larger market segment, we knew we had to increase the density. So that's why we went then to DRAM interfaces because most of the 256 meg and 1 gig, 2 gig and so on are in DRAM interface that need high performance like our device.

So the point I guess Kevin is that it's a very versatile memory, you can actually configure it to be optimal in either of those cases SRAM or DRAM.

<Q – Kevin Cassidy>: Okay. And then along those lines is when you talk about the automotive applications in semiconductors, automotive is the fastest growing segment. How is in the MRAM best suited for automotive applications?

<A – Phillip LoPresti>: So one of the things again I was highlighting earlier that there is an increasing demand for high performance persistent memory. And so in automotive applications, they've been using things very simply called data loggers that log information from various sensors or from the engine or different things that the driver is doing like their infotainment system.

What's happening with these new cars as they have more sensors, more data coming from the sensors more frequently, and more rapidly. And they're finding out that some of the memory technologies that they were using before are not really cutting it to provide this the adequate endurance or durability nor they perhaps don't have the temperature capability that we have. We have an automotive qualified process minus 40 to 125.

And if you go to now assisted driving and ADAS type of systems. You're going to see even more data now, not just sensor data but camera data whether they're taking snapshots or streaming video. They want to store it in some kind of black box device and this could be a centralized black box or it could just be at each relative system as a microcontroller makes a decision.

This data logger has to capture that data, the sensor data that went with it and perhaps some video data. So that if something goes wrong whether it's an incident or a failure the manufacturer can look at that data log and then process the information and determine what went wrong. Obviously insurance companies are interested in that as well because now they have something that actually logs what happened in the vehicle, when this incident occurred.

So we see MRAM being very interesting in that application because we have the capability from a temperature range. We have very high durability, you can read and write to it and virtually not wear it out. You can write very fast and once you write to it, it's nonvolatile, that means if power were to be severed the data is protected.

<Q – Kevin Cassidy>: And I think Everspin is very unique for the size of your company that with your heritage of Motorola and Freescale that you have all the systems in place and the qualifications to address the automotive market.

<A – Phillip LoPresti>: Yeah, it was actually something that we looked at our 200-millimeter fab that where we produce our Gen 1 products is in – was Freescale, now it's NXP perhaps soon to be Qualcomm fab is an automotive qualified fab. And because we operate our back-end of the line process, the very tail-end of the processing of the wafers inside that fab. We also have the operator and automotive qualified within that fab. So it made sense for us to take advantage of that. We qualified our parts and now we're certified in that area and we've got some significant design wins in infotainment and as well as in eCall systems.

<<Kevin Cassidy, Analyst, Stifel, Nicolaus & Co., Inc.>>: Thank you, Phill, Jeff. Thanks for coming.

<<Phillip LoPresti, President, Chief Executive Officer and Director>>: Thank you everybody.