THOMSON REUTERS STREETEVENTS

# EDITED TRANSCRIPT

## INTC - Intel Corp to Discuss Reports of New Security Research Findings

EVENT DATE/TIME: JANUARY 03, 2018 / 10:00PM GMT

### OVERVIEW:

Co. provided some update on security research findings and its approach to mitigation.

## CORPORATE PARTICIPANTS

**Donald Parker**

**Mark Henninger**

**Ronak Singhal**

**Stephen Smith**

## CONFERENCE CALL PARTICIPANTS

**Ambrish Srivastava** *BMO Capital Markets Equity Research - MD of Semiconductor Research & Senior Research Analyst*

**John William Pitzer** *Crédit Suisse AG, Research Division - MD, Global Technology Strategist and Global Technology Sector Head*

**Ross Clark Seymore** *Deutsche Bank AG, Research Division - MD*

**Vivek Arya** *BofA Merrill Lynch, Research Division - Director*

## PRESENTATION

**Operator**

Good day, ladies and gentlemen, and welcome to the Intel Investor Call Regarding Security Research Findings. (Operator Instructions) As a reminder, this conference call is being recorded.

I would now like to turn the conference over to Mark Henninger. Please go ahead.

---

**Mark Henninger**

Thank you, operator, and thank you all for joining us on this conference call to discuss the side-channel analysis security issue. The presentation that will accompany our remarks today is available on our investor website, intc.com.

I'm joined today by Steve Smith, Corporate VP and General Manager of Data Center Engineering; Ronak Singhal, Intel Fellow and Director of CPU Compute Architecture; and Donald Parker, Vice President of Data Center Platform Application Engineering. In a moment, Steve and his team will share prepared remarks followed by Q&A.

Before we begin, let me remind everyone that today's discussion contains forward-looking statements based on the environment as we currently see it and as such, does include risks and uncertainties. Please refer to our press release for more information on the specific risk factors that could cause actual results to differ materially.

With that, let me hand it over to Steve.

---

**Stephen Smith**

Mark, thank you very much, and welcome to our participants on the call. I'm Steve Smith and lead our Data Center Engineering. What I'd like to do is take you through our objectives for this briefing, a little bit of material, and then we will allow plenty of time for Q&A.

So you should have the slides available to you via the website. As we note on our foil #2, our objectives are: first, to provide some information on a new class of security issue that has been recently identified and also talk about our approach to mitigation; and then second, provide some context on this issue and how we brought the industry together to come up with a common approach to mitigating the problem.

If we go on briefly into foil 3 on the background. There was a security research team that notified Intel as well as other industry participants, including AMD and ARM Holdings, of a new side-channel analysis exploit of the computing system. And what this is, is a method for an attacker who uses this exploit to observe the contents of privileged memory in a way that circumvents the normally expected privilege levels in the processor. So the person running the exploit can execute some code and using this exploit can actually read the contents of memory that are not normally available to that user's process. And a little bit later, we'll tell you more in detail how they do this.

But in a high level, what they're doing for the first time is utilizing the speculative execution techniques that we have in most modern processors. And these techniques are used across the board in different families of processors, different architectures and so they're not unique to one architecture or to a particular processor implementation. So what this means is, whether it's an Intel processor, a different x86 processor, an ARM processor or conceptually, other architectures of processors that use speculative execution, this technique can possibly be employed on those processors. And speculative execution simply means that the processor can take some action in anticipation of computation that will be needed and bring an operation or bring some data into memory, and then this approach takes a look at that data that's brought into memory even though it's not typically available to that process.

The exploit is really not the result of a product errata. The processors are really operating as they should operate, as they were designed to operate and validated to operate. So in this whole process, nobody has reported to us or to any of the other vendors a processor flaw or bug that this technique exploits. It's simply what we call the side-channel attack on the cache.

Our approach to mitigating this is to update system software or update firmware of the processor. And some of the exploit can be mitigated by the OS vendors. Some of the exploit we actually have to change some firmware that resides with the processor. And those together will mitigate this issue for our current products. And we can -- now that we know about this issue, we can improve the processor hardware for the future, such that we have a more elegant implementation in the future hardware.

And in order to tackle this, to understand this, the issue, and to come up with a common approach, we've launched an industry-wide collaboration to facilitate this. And it's under the auspices of responsible disclosure. Excuse me, for my voice here. So responsible disclosure practice is that as a researcher points out a security issue, the party with the issue or in this case, the whole computing community, develops the mitigation and has the mitigation available at the time that the exploit is documented in a research publication. So we've been working with industry participants. And Donald Parker, who's with us today that Mark introduced, has been leading that collaboration across the industry. And that has involved multiple microprocessor vendors, operating system vendors and OEMs around the world to understand the issue, to develop the system software updates, to develop the firmware and to integrate and test those things. And we do expect those mitigations to be available beginning over the next few days, but from the Intel perspective, taking the next several weeks to get the mitigation setup for all the Intel products.

Moving on to Page 4. Just a little bit more of a clarification of what this issue is and what it's not. So first, it's a method for this attacker using this exploit to observe the contents of privileged memory. And that's memory that's not normally available to the user process and thereby, you're circumventing the expected privilege levels. And when you think about this malware that's using this method and running on the machine locally can expose sensitive data that attackers might be interested in finding on the system. But I have to point out at this point that this is a very complex type of attack, and one can read the contents of the memory at a given address and that address may or may not have information that's actually relevant or useful to the attacker. So there's some steps that are required beyond just reading memory to figure out the value of any data that's read.

What the attack is not. It's not a denial of service attack. It doesn't stop a computer from running. It doesn't stop a user process from running. It's not a network attack. It's not something that one can launch over the network or compromise the network. And I think perhaps most importantly, it's not a way to inject malicious code. It's not a way to corrupt memory. Even though one can read memory that's outside of the normal privilege zone, one cannot actually write to that area, change the contents of the memory or corrupt the data that's in-memory in that area that's read. So it's simply the ability to kind of look outside the normal memory protection zone and observe some of the contents of memory.

The other thing that's important to point out is, since the researcher shared with us in a responsible disclosure mode the results of their work, we've had time to work on that with the researchers and put the mitigations in place. And to this date, we've observed code that is, I'll call it, proof-of-concept to show what the type of attack is, but we have not observed any active deployment of the exploit in the real world of computing.

Okay. And with that said, we can move on to Page 5. Our approach to mitigation is that we look at this and we want to provide for Intel platforms the most comprehensive approach that we can, such that we will offer the highest level of security for our users on Intel compute platforms. And we've been working to put together a combination of operating system updates on the broadly used operating systems and some firmware updates that we developed that are specific to the configuration and operation of our processor, sometimes called microcode. And that has all been developed with industry partners, tested with industry partners, working with OS vendors and with OEMs. And we've been working at this for some time, such that we'll be ready the beginning of the next few days to start the deployment of the mitigations. And again, it will take probably a few weeks before the mitigations we have in mind will all be available to customers.

And with that, I want to hand it over to Ronak Singhal to describe a little bit more about the security issue because there are several different modes of the possible attack and they require slightly different mitigations.

So Ronak?

**Ronak Singhal**

Thanks, Steve. So what the researcher had disclosed to us are 3 different variants that are all variations of doing the side-channel timing attack in order to gain access or information about data that you normally would not have access to.

So I will briefly walk through each of these 3 and what our mitigation strategy is from the Intel side. The first one is what we call a Bounds Check Bypass, essentially take advantage of existing code that has access to privileged information and use it or abuse it to speculatively access memory that an attacker typically would not have access to. This one is a fairly fundamental exploit and the way for us to mitigate is via software. So we've been working with software partners on both the operating system side and the browser side for mitigations for the first exploit.

The second one, which we term the Branch Target Injection, is that the malicious code finds a way to essentially redirect the internal structures inside the processor to speculatively execute code that they want to see executed. Again, this is all done speculatively, which means it doesn't affect the basic functionality of the processor, but it allows a side-channel cache timing attack to occur. For this variation, we are issuing a microcode update that Steve alluded to that provides a new interface between the OS and the processor that the OS can then take advantage of to mitigate this variant. So the strategy here requires both updates on the software side as well as on the hardware side.

The final variant is what we call the Rogue Data Load, which is the ability for a memory access from an application to speculatively access memory that it should not have access to. And this is the one where you may have already seen some of the mitigation strategy as we've already pushed patches to Linux for isolating the page tables between the kernel and the user space. And by doing that, that serves as an effective mitigation for this third variant. Now in all of these cases, as Steve also mentioned, we will be pursuing hardware initiatives or hardware improvements for both performance and security going forward.

**Stephen Smith**

Okay. Thank you, Ronak, and appreciate that. Let me just make a few more comments before we turn it back to the Q&A.

So again, just to summarize, after receiving the report of the security research, we have led the industry initiative to respond to these in a responsible manner. We have brought together teams from Intel, from AMD, from ARM, from various operating system vendors and from OEMs around the world who are working with us to put the mitigations in place.

The second is a reminder that the processor is, in fact, operating as designed. And in any -- in every case, it's been this side-channel approach that the researchers used to gain information even while the processor is executing normally its intended functions. And with this in mind, we still want to make sure that we're improving security for the user, so we've been providing these firmware updates and we're working with the industry to provide the software updates to this.

And Ronak also mentioned that we've seen some data posted on the web and contrary to some of the reports, the performance impacts are really very workload dependent. And for the average user, we really don't expect those to be significant and then those will be mitigated over time as Ronak talked about.

So we're committed to product and customer security. We've taken the approach to lead and resolve this for the industry. We're having this disclosure today because there were some information in the media that was posted, I'll call it, ahead of time and potentially misleading, so we wanted to just clarify the situation. And then again, since our products are performing to specification, we really don't anticipate any material impact to our business or to our products since they continue to operate properly and since we're adding these firmware and software for the security mitigations.

With that, I'll turn it back over to Mark.

**Mark Henninger**

Great. Thank you, Steve, and thank you, Ronak. (Operator Instructions) Operator, please go ahead and introduce our first question.

## QUESTIONS AND ANSWERS

**Operator**

Our first question comes from the line of John Pitzer with Crédit Suisse.

**John William Pitzer** - *Crédit Suisse AG, Research Division - MD, Global Technology Strategist and Global Technology Sector Head*

Steve, just a couple of questions relative to the mitigation issues that you're taking. On the software/firmware update, you mentioned that the performance degradation will be somewhat negligible for the average user. Is that the average user inside of the data center group, inside of your compute group? And I guess I'm more curious about the high-end user within the cloud or within your data center business and what kind of degradation they're likely to see with the software mitigations.

**Stephen Smith**

Yes. Thanks, John, and hi. Hey, it depends on the workload specifically in use and a little bit less where the workload is. I'll ask Ronak, if you can, just describe the attributes of most workloads and then the attributes of the workload where maybe there'd be a little bit more performance impact.

**Ronak Singhal**

Yes. Thanks, Steve. So as you were saying that this isn't really -- we don't differentiate on the performance impact between PC Client and the data center part. It's really about the attribute of a workload. A workload that is largely running in app -- in user space will see limited to negligible impact is what we've seen. So some industry-standard benchmarks, we've seen 0% to 2% type of impact. There are other workloads that spend a lot of time going back and forth between the application and the operating system. And those are the types of workloads where you can see higher levels of impact based on the mitigations that have been put in place. And so there are, for instance, synthetic workloads that have shown 30% or more performance impact as a result. But the statements that Steve made on performance are around, if you look at the broad spectrum of workloads as well as taking into account real-world scenarios, you get to that level of impact.

**John William Pitzer** - *Crédit Suisse AG, Research Division - MD, Global Technology Strategist and Global Technology Sector Head*

That's helpful. Then Steve, maybe as my follow-up, when you think about hardware mitigation, how long will that take? It sounds like this is probably beyond the Skylake microarchitecture and Purley. And if you think back to both floating point and Cougar Point issues, there were also some financial impacts there, charges. Do you envision anything from a financial charge perspective around solving some of these issues?

---

**Stephen Smith**

Yes. So first of all, the mitigations that we've talked about are OS and firmware that are deployable in current systems. So with that, I'm looking at my own PC or I'm thinking about a data center, those things will get updates via the normal software/firmware update mechanisms. And as such, there's no meaningful cost to Intel. So no, we do not expect any financial implication around Intel's products as a result of this one. So that's kind of answering backwards order.

Your second, your earlier part of the question was, what's the time frame for implementing changes? And I'll just say that as part of the mitigation, as we learned the root causes of these issues, Ronak and his team are accountable for doing the microarchitecture of our cores and they started making the changes in the products for our future. And you can look at those as, I'll call it refinements, so that the OS and firmware have to do less heavy lifting. We just kind of build the updates into our hardware in a more -- in a way that's more transparent to software. And you'll start seeing the first of those products within this calendar year. We have a general direction that all of our products going forward that are not yet, if you will, in silicon and committed to launch in a very short period of time, all those future products will incorporate those enhancements as Ronak and team learn more. And that will just make the mitigations more efficient.

---

**Operator**

(Operator Instructions) Our next question will come from the line of Vivek Arya with Bank of America Merrill Lynch.

---

**Vivek Arya** - *BofA Merrill Lynch, Research Division - Director*

Steve, I know you can only talk for Intel. But for some reason, AMD seems to be quite vocal that these issues do not apply to their product because they are perhaps not doing some of this speculative execution-type techniques. Can you clarify that this is a general industry issue and not specific to Intel, that it applies to everyone? Or is there a difference between the way that your products are designed versus AMD products that this might not apply to their products?

---

**Stephen Smith**

Yes. So Vivek, hi. What I can say is that the researchers have demonstrated some of these exploits running across a variety of product implementations. And some of the mitigations that are called out are in software, OS and such, and some of them are in hardware. So Ronak has gone through what we're doing in our product when you look at our Page 6, describing where we're working with the OS vendors, where we're doing virtual machine manager updates and where we're doing firmware that's specific for an Intel processor. So we've given some detail on what we're doing, and we've also given you the general statement that the sum of mitigations includes the combination of hardware and -- software and hardware.

---

**Vivek Arya** - *BofA Merrill Lynch, Research Division - Director*

So again, I want to clarify, is this an Intel or an industry issue?

---

**Stephen Smith**

It's an industry issue. And you'll have to ask each participant what their particular product-specific mitigation actions are.

**Ronak Singhal**

Let me briefly add to that. The issue was originally found from Google Project Zero, which is a security team within Google. And you can now find information that they have published about this on their website at security.googleblog.com.

**Stephen Smith**

Yes. So let me, Ronak, let me ask you to walk through that again one more time so people know what that link is.

**Ronak Singhal**

The link is security.googleblog, all one word, .com.

**Stephen Smith**

Okay. And then that maybe we just need to describe a little bit, I'll ask Donald to comment here, about Google Project Zero and then our effort to work with them and the whole industry.

**Donald Parker**

Sure. So Google Project Zero is a set of researchers within Google that work with the industry when they find security issues in products, whether they be software, hardware, firmware, different types of products. Google works with those impacted by that. And in this case, they worked with CPU vendors across the industry and they mentioned that in their blog today. And they will work with us and -- so that we can help with mitigations and then they will publish their results. They published a portion of those results today, and we expect to see more details in the coming week.

**Operator**

Our next question comes from the line of Ambrish Srivastava with BMO Capital Markets.

**Ambrish Srivastava** *- BMO Capital Markets Equity Research - MD of Semiconductor Research & Senior Research Analyst*

I had 2 questions and I'll stick to the 2. First, on the performance side, based on the change that or the modification that you would be making, would that lead to any performance impact as -- and pardon me for my lack of knowledge of how the kernels work, but just going with all the reports that are out there in the media, would that impact performance in the future architecture? That was my first question.

**Stephen Smith**

Let me ask Ronak to speak to that. I think that Ronak has already covered this, but I think it's worthwhile to talk about again the expectations for performance versus workload.

**Ronak Singhal**

I want to make sure I understood the question. It sounded like the question was more around in the future as we make hardware changes, what will the performance impact be?

**Ambrish Srivastava** - *BMO Capital Markets Equity Research - MD of Semiconductor Research & Senior Research Analyst*

That's right, yes.

**Ronak Singhal**

Okay. As you cited, there have been some reports around what the performance impact is. As both Steve and I have mentioned, the impact will vary by workload that you see today. Going forward in time, as we make changes in our hardware, you should expect to see that the impact of those mitigations have a lower performance impact than the mitigations that we're using today for products in the field. Basically, as we have the ability to move this from the firmware and software domain into the hardware, we're able to lessen the impact while providing security.

**Ambrish Srivastava** - *BMO Capital Markets Equity Research - MD of Semiconductor Research & Senior Research Analyst*

Okay. That makes sense. And just the second is a clarification, when you talked about the -- and I think you said 30% performance impact, I'm assuming when you made the difference between the user versus nonuser, where the workload is interacting more with the software as opposed to the user, is that on the data center side that you see that kind of impact?

**Ronak Singhal**

So again, the 30% number is more of an outlier case. It's not what we see as the typical case. But like we said earlier, it's not a statement about data center versus a PC usage. It's more about what kind of application you have. You could see that application in either scenario.

**Operator**

Our final question will come from the line of Ross Seymore with Deutsche Bank.

**Ross Clark Seymore** - *Deutsche Bank AG, Research Division - MD*

Lots of the good questions have been asked and answered. So I just want to take a higher level question and get some perspective to this. Steve, how common are these sorts of attacks and this sort of industry consortium coming out to address it? And when you talk about having the -- what was the word you said, industry-wide collaboration for the responsible disclosure, just walk through the steps that we had here because a lot of questions were asked before you guys responded today and the people were speculating that, that was indicative of some bigger problem. You went through a lot of answering those questions, but just walk us through the commonality of this sort of event and how you are addressing it and how common that methodology is as well.

**Stephen Smith**

Yes, Ross, hi and thank you. Let me just start with the high level. We get security issue reports quite often. And it could be anything from a particular issue with a particular LAN driver on -- that someone has observed that it's a quick fix in software. To this one, which is kind of a more complex, and I'll call it a new approach with this side channel. Side-channel attack concepts have existed for the decade and decade or so, and we've responded to a number of those, but this one using speculative execution kind of moved us to a new domain where it was quickly realized that this applies to, I'm going to call them most modern microprocessors that are high performance and utilize speculative techniques to gain the performance advantage. And so that has gotten us to drive in a more kind of industry collaborative way because it's not really one vendor's problem. It's not an issue with our product. It's not an issue with someone else's product. It's a general design approach. You can design for correctness or -- and then from a different view, you can design to improve security. So what we did was put together an industry coalition. And Donald Parker, who's here with us, has pulled that coalition together and has been working for months with the participants in the industry to align the different pieces. And you can see that we're talking about OS releases, we're talking about VMM releases, and we're talking about support and changes from multiple microprocessor companies. And then all of these need to be deployed through OEMs. So it becomes a project in its own to put in place a mitigation. And Donald, maybe you would just want to comment again briefly that you've been pulling together a pretty collaborative team of those various players, right? And it's -- we hope it's a new template for how things are done. Go ahead.

---

**Donald Parker**

Yes, I think that's the case, Steve. Thanks. As you pointed out, we worked with these other CPU vendors, operating system vendors, OEMs, cloud service providers, to bring the industry together to really ensure that we had a mitigation across all impacted products simultaneously and really a deep focus on ensuring mitigations were available as the publications came out this week and next. So yes, it's been a -- I think it's a sort of a new model as these threat landscapes evolve over time to seek out that collaboration across the industry.

---

**Stephen Smith**

And I think you'll see that if you watch carefully in the coming weeks as different vendors talk about their products and their mitigations you'll see the different releases, patch strategies, et cetera, from various vendors that are the outcome of this collaboration that Donald put together.

---

**Ross Clark Seymore** - *Deutsche Bank AG, Research Division - MD*

So maybe my quick follow-up. Yes, that answered my question perfectly, Steve. The quick follow-up is, just if we put this all together, do you expect much, if any, of a financial impact or even a market share impact because of this issue on Intel?

---

**Stephen Smith**

As we look at it, we don't see any financial impact. It's been a lot of work by a bunch of people so you can think of that as opportunity we could have spent elsewhere, but that's not anything you can put as cash cost on the bottom line. And I wouldn't expect any change in acceptance of our products, and I wouldn't expect any -- kind of create financial impact that we would see going forward.

---

**Mark Henninger**

All right. Thank you all for joining us today. And operator, if you could please go ahead and wrap up the call.

---

**Operator**

Thank you. This concludes today's Q&A session. Ladies and gentlemen, thank you for your participation in today's conference. This concludes today's conference. You may now disconnect. Everyone, have a great day.

---

## DISCLAIMER

Thomson Reuters reserves the right to make changes to documents, content, or other information on this web site without obligation to notify any person of such changes.

In the conference calls upon which Event Transcripts are based, companies may make projections or other forward-looking statements regarding a variety of items. Such forward-looking statements are based upon current expectations and involve risks and uncertainties. Actual results may differ materially from those stated in any forward-looking statement based on a number of important factors and risks, which are more specifically identified in the companies' most recent SEC filings. Although the companies may indicate and believe that the assumptions underlying the forward-looking statements are reasonable, any of the assumptions could prove inaccurate or incorrect and, therefore, there can be no assurance that the results contemplated in the forward-looking statements will be realized.

THE INFORMATION CONTAINED IN EVENT TRANSCRIPTS IS A TEXTUAL REPRESENTATION OF THE APPLICABLE COMPANY'S CONFERENCE CALL AND WHILE EFFORTS ARE MADE TO PROVIDE AN ACCURATE TRANSCRIPTION, THERE MAY BE MATERIAL ERRORS, OMISSIONS, OR INACCURACIES IN THE REPORTING OF THE SUBSTANCE OF THE CONFERENCE CALLS. IN NO WAY DOES THOMSON REUTERS OR THE APPLICABLE COMPANY ASSUME ANY RESPONSIBILITY FOR ANY INVESTMENT OR OTHER DECISIONS MADE BASED UPON THE INFORMATION PROVIDED ON THIS WEB SITE OR IN ANY EVENT TRANSCRIPT. USERS ARE ADVISED TO REVIEW THE APPLICABLE COMPANY'S CONFERENCE CALL ITSELF AND THE APPLICABLE COMPANY'S SEC FILINGS BEFORE MAKING ANY INVESTMENT OR OTHER DECISIONS.