EMEA IT SPOTLIGHT

DRAFT

Keynote Presentation

by

Michael Kagan, Chief Technology Officer, Mellanox Technologies

Think outside the computer

Michael Kagan

Okay, so good morning, everyone. It's actually great to be the first speaker and not the one before the lunch. So we'll start about some things about networking and we are in the midst of the revolution that actually will see the perspective in maybe many years from now but we are in the midst of a huge transformation over the world, and we'll talk about this.

So for me, from my perspective, all the things, all this starts with this move from phone being static to being mobile. I got this machine in my hands about 30 years ago; maybe a little bit more than 30 years ago, and I thought it's really cool things because I can talk to my kids, to my wife; I worked at Intel those days. So on the way - and I live in Israel, so it's 10 hours difference, so on the way to and from work, I could talk to my colleagues in the United States, which surely opened up new ways of operation and communication.

But I didn't realise those days that 30 years later, this machine will become something that you never leave home without it; it's actually what Steve Jobs told his engineers; I want you guys to design a machine that nobody would like to leave home without it. Now we use the phone for news, emails, gaming, social media. Sometimes we use it as a phone and you can do it all even without leaving the toilet.

So everyone now becomes a huge source of data. This thing generates an amount of data that nobody ever thought before. Data grows exponentially and every year, or maybe even every half year, the amount of data generated doubles. So every year we generate as much data as was generated so far before. Data becomes a resource. What used to be oil in the 20th Century is now data. Once we have access with data and once we have access to data in an efficient and productive way, you can do a lot of things, good things for yourself, for your business and for others.

I have here a very short example of a paradigm shift of getting the data. In 2007, Nokia was a huge company, a very large company, about $150 billion worth. Those days iGO and the like, navigation gadgets, took off, so people started to use the navigation in their cars. Nokia decided to get into this business as well but in order to get all this business with added value, they decided to build the navigation system that actually considers the traffic.

To make it work, they need traffic sensors, so they've gone off and both Navitech, which is the European company, they had about five million traffic cameras across Europe. So with traffic cameras you can know where the traffic is, builds up, and this way, by connecting this to your navigation software and navigation gadgets, you can actually navigate people around the traffic.

So [accidentally], actually the same year, in 2007, there was Israeli company started Waze. They figured out that they can get same data, which is traffic sensors, for free by writing this application software and every phone that runs this software actually generates the data to Waze and pretty much the rest is history. So Waze figured out that they can get the data in a much more efficient way out there and the rest is history. Nokia gone down in five years, or six years to less than what they bought Navitech for and Waze sold to Google for more than one billion. So once you figure out that the data is out there, and you find the way, efficient way to get this data to your hands, you can do a lot of good things for yourself and for community.

So the storage, the disks, and you know, they [unclear] they evolved and that's a whole other topic that I'll not touch too much on the media itself. But it turns out that there is enough disks to save and to store the huge amount of data that's being generated; more than 70 per cent of the data that is generated by the phones and the chats and so on and so forth are actually stored somewhere.

Only small fraction of this data is actually ready to be processed. It's not the database - it's not the databases; it's unstructured data. Very little is being analysed and a huge amount of data is wasted in terms of it is not being used to generate the real information and real value. So there is a gap between the resources that are out there and, in many cases, as we have seen with the Waze example, for free, they're free or almost free to the point that we actually can use them.

Okay, so to use the data is to process the data and we used to write programs on our machines to take the data and convert it to the information which is running the application. But the thing that is challenging here; the growth of the data and growth of the CPU power are not aligned. Actually the data growth is faster than even the [Moore's Law] and the CPU performance, if you look at the performance of the [one human] machine is not keeping up with the data and barely keeping up with the [Moore's Law].

The good thing is that the network is there and the network bandwidth and network performance are pretty much aligned with the data growth; maybe it's actually related but that's a different story. So we have these other technologies and now we need to think about how to do it differently. With the gap growing, it means we cannot do the evolutions anymore. Like electric light didn't come from the improvements of candles, we need to think outside of this box of what can and what cannot be done in order to improve our - take advantage of this data that is being generated. The CPU model, one human machine model is not good enough for us anymore.

So in the 20th Century, we took one human machine, put people to program the machine, created the application or service so it gets the data, it generates the output. All good. Now we can't do this anymore. CPU doesn't keep up. So what do we need to do? We need to use the data to generate the services itself. This is the artificial intelligence model, neural network model; how do you take the data and actually create the service that using the data. The machine learns from the data to do the job and the application or services are created by data itself.

So when we create the services we need to basically think it's outside of the computer. It's not a single individual processing unit, a human machine anymore. It's the connecting different pieces of the network of machines and - by the network, and make them work together to create the services. It means take - optimised across the data centre, optimise across the cluster. With IOT, we had this just debate, discussion at the breakfast, it's actually optimising even outside, looking at the whole picture, even outside the data centre.

It's all about efficiency, but it's not just efficiency. If you do the right things we can do new things that we couldn't do before. It's not just doing the same thing faster, it's doing new things that we haven't done before. So one of the analogies is ants. A single animal like this - insect like this, it can't do much. But many of them together, working together in a way that people don't really understand this, I don't understand how it does work, they can build huge things. They build cities with ventilation and irrigation system and everything. This, by connecting to each other, by working together efficiency, by linking together efficiently these resources, you can do things which is not linear extrapolation of one guy what it can do.

So how do we optimise? First of all, we need to change - think about this paradigm, that it's not a single programmer, program a single machine. It's data programs the computer. You need to make machines that can learn from data. You need to make the machines that program themselves. You need to enable the data processing or computing - the data processing everywhere; on the compute, not itself, on the general-purpose computer, in the network, and in the storage. Computing and storage [unclear] is going to end. Last but not least, security is a really interesting topic that we will touch a little bit because, once we have more data, we can understand more and we can know more, and that comes the security from crime to privacy to everywhere, okay.

So Cloud. Cloud is the different approach; it is moving from the server approach to service approach. When we ask people, at least until not very long time ago, you asked people whether they have electricity in the home, the answer would be usually, yes, but what they mean is that they have a plug in the wall and the electricity is actually generated by the electric plant far away out there.

Until recently, when you asked people whether they have a computer at home, or have some storage disk at home, and the answer was also usually, yes, but what they mean is that you have a box under your table that you need to buy, you need to manage and so on and so forth, and that's your computer, that's your disk. But essentially there is not much difference between electricity as a service, and computing as a service; and that's what Cloud computing is all about.

Now, when you want to create such a service, it means that you need to take multiple pieces of hardware, let's say, or multiple resources and assign them, or chain them - this is microservices, together to create the application, or create the response to the request of the client. Now it becomes fairly tricky things.

If you think about the analogy of autonomous cars, today the car spends about 95 per cent in parking. But it's a very clear process, okay. I leave my car, I lock it, it stays in the parking lot for whatever it is, and then I come back and I pick the car up. So most of the time, actually car does nothing. When we go to the autonomous cars, or on demand cars, then basically I'm asking for a car and it shows up in a minute in the place when I am. So now you need to connect between these things and manage these resources.

The same thing in the data centre. When I am looking for some sort of service, there is multiple resources that need to work together in the [chain] to provide the services. That's what data centre is about. It's taking multiple resources and being able to configure them in a fraction of second to respond to the request from the consumer.

Okay, so we'll talk about cloud. Do you know when there was the first cloud download? First download from the cloud? It's not a new concept at all, okay. It happened about 3500 years ago when Moses brought the scripts - those days, cloud was pretty, you know, not as advanced as now; I'm not talking about cloud, I'm talking about the interface. It was single use and it took a day to download 10 sentences, not that huge.

But now we are talking about cloud as a service provider. So in order to make it efficient, again we talked about thinking outside of the box, or sort of [non-linear] thinking. Long time ago the communication was like pigeon mails, so if you want to move more data with pigeon mails and you think of this pigeon mail as the concept, you can move more data, you can move faster data, you can find, you know, a faster pigeon, okay. This guy can probably pick two scripts at a time and probably flies faster, but it doesn't inherently change the way you communicate and you work.

With the new technologies like video conference sync, it is not just faster than pigeon, it enables to do things that you couldn't do before. Today we can have an interactive meeting across the ocean, across multiple countries, talking with multiple people. Soon, even in different languages with the automatic translation on the way. No matter how fast your pigeon flies, you couldn't do this - you can't do this with this technology. So you need to have a new technology, a new way of thinking, a new way of working, and then you can do things that cannot be done before.

Now we are in the networking business and we deliver the networks that enables you to do more than you could have done before. It enables you to do new things, more things. We worked with Oracle since 2002 and they actually internalised that using our network, the communication between different machines is not [botonic] anymore. They re-wrote their Oracle rack system and by moving from 10 Gig networking to 40 Gig networking, they got 50X performance boost. Not 4X, 50X. The reason is because they used their resources more efficiently.

The other thing which is more recent is by using our networks, [Bidoo] using Nvidia GPUs and Mellanox Interconnect, the first time they build a machine that the first time was classifying images, or recognise images better than human. Today, computers can see better than humans see. The image recognition on the computer is better than the human and the first time it happened in 2015 when Bidoo build this machine using our network and they couldn't achieve such a results with any other network.

So network is not just about efficiency on the interface to the application. It's not just about how fast you move the data between the - on the wires. It also processing the data while it moves. In our most advanced network products we have a floating-point units, actually computational units within every switch. So we can do the data aggregation operations on the fly and actually when you ask for multiple results on different - on thousands of machines, you can get the end result from the network and not getting all the results and then figuring out on yourself.

This technology actually is being used now today in the machine learning. When we do the neural network training on sending different data sets, the multiple instances of the neural network model and after individual training, there is a consolidation of the training results [which are called] parameter server - usually takes as much time as training itself; and by speeding up this parameter server by factor of 10, we can actually get down the training process from days to hours, or from weeks to days. So this is our [tower sharp], technology sharp stays for the scale [of the hierarchical] aggregation and the reduction protocol, it's available in our most advanced network products today.

So one thing that I will touch a little bit on the storage; and on the storage it has to deal with the efficiency of the storage and performance of the storage, and how we can deliver storage services to applications and to systems that are actually [can not] necessarily advance to take full advantage. So SNAP is the software defined networking and accelerator storage processing. If you look at the internet cloud, internet fabric, you have multiple machines connected to each - to the network and each one of these machines have its own storage, which is - and your application running on this machine uses the storage; and it's obviously not highly utilised because it's real [hardware] that is on this machine.

You can use [NVMe] fabrics, which is the great thing to access the storage over the remote. The problem with this is that not all the computers and not all the platforms are actually ready and can take advantage of this. There is an issue of the updating of the software and application and the APIs and the platforms. So it works but then you need advanced technologists on each machine to take advantage of this. It's not always the case.

So what if you could have storage services? You could be able to represent anything on your cloud as the really true local device. Not network attached device that has some driver that presents APIs that you can use, which implies some changes on the host. But you see on your machine on the legacy operating system, you see the local device, but this local device somehow magically can touch everything.

So what we deliver, we deliver this with our SmartNIC technologies, BlueField can represent - or present, sorry, to the host, to the operating system, the local device and it can [unclear the training] on this machine can actually allocate resources from different machines on the network and make them available as the local storage devices, or local storage services resource on the local machine. So this way, the utilisation of the storage goes up. You can actually utilise, and you can even [theme] provision, you can over-subscribe your physical storage resources and it still works. We have a pilot project with the major cloud providers right now and they hopefully will move to production with this, either this year or next year.

So by using the advanced interconnection, that's again about the efficiency of the interfaces, and efficiency of how can you access your resources from one machine, from one physical machine to another. You can actually have a much more efficient - much better utilisation of your resources and, again, it is about service chaining. I have a multiple stops, multiple microservices that needs to be chained together to deliver service my application is looking for, and this is one of the technologies that enables to do it.

Okay, so these things are usually called virtualisation. We virtualise the resources, this way or another, by running - by abstracting out the physical infrastructure from the application; and as with this analogy of the power plants, then before the plants [unclear] there was no workers at the power plants. Now once we build these things, we need to invest, and we need to invest computing resources, computing resources to manage the physical resources and to run, actually, this virtualisation layer.

So if I had my machine that I could use all the compute course to applications, now part of this computing is actually need to be invested in virtualisation. So now my machine becomes less efficient from the application perspective because I have lower computing power. Fortunately enough, all these things can be optimised by off-loading fairly significant part of the virtualisation operation to the hardware, to network adaptors, or to the storage. But still, some part is out there running on the compute server and this would not be a big deal because this is a really small part.

The problem with this is security. Today, crime - you know, you don't need to really rob the bank and go with the guns to the bank. You can do it - you can get much more money if you're doing crime online and therefore the security becomes an even bigger topic.

So when we talk about the classical enterprise data centre, this is - the security model there was that, based on the fact that in the enterprise data centre, the owner or IT manager of the data centre has full control of all the software that's running on that data centre - all the applications running on the data centre. So you have a security policy on your machine and you have the application that are not going to do any harm because you know this application; and the only thing you put is you put the firewall on the edge of the data centre, so the bad guys would not come in.

This is the M&M model, which is, you know, hard on the outside, soft on the inside. You have a periphery protection on the periphery of the data centre, but inside the data centre, it is not protected. Your data centre is unprotected. If somebody gets inside, and there are some data centres that are protected sometimes for this firewall is actually implemented with the people with the machine guns, like in the power plant, there is no connection to the internet, so it's firewalled with real fire. But once you have this control, you are good.

In the cloud, the things are changing, okay. You have a security policy running on the compute server and you have multiple applications. Unlike the enterprise data centre where you have full control on the application that's running on your machine and you are deciding, as the IT manager or CIO, you decide whom to run on your machine and whom not to run.

With the cloud space, we are inviting other applications which we have no control on to run on my machine because I'm [actually sending NIPS], right. So I am bringing the new applications and to the customer application and this server, or running on the same machine that runs my security policy. Now what may happen is that one of these guys are a bad guy, and the problem with this bad guy, that once he runs on my compute server, it can actually take over the security policy.

Once he took over the security policy, he took over my data centre. So data centre is gone. Data centre is compromised. Entire data centre is taken over by single guy that I invited to run on my machine. He took over the security policy on the machine that he was hosted on and now all the data centre is under his control.

If you think it's not possible, it is possible, okay. There is 30 years of history of host-based security problems. Until '17, actually, they were all based like the Heartbleed is the most recent one; it was - all of them were taking advantage of some sort of software bugs in the operating system that the bad guy could sneak in. These three are actually based on the side-channel attack, which is not a softer bug, it is taking advantage of some [unclear] of how people build microprocessors and this type of thing happened on any operating system, on any processor.

So in order to protect our infrastructure, we need to make sure that attacker and the victim are not sharing the same computing domain, they're not running on the same computer. If you let attacker to run on your computer, no matter how good your operating system is, no matter what you are doing, but if you let other guys to run on your machine, you will be - you are under the risk of being hacked.

So in order to make it work, what we need to do, we need to change the security model from being soft inside, we need to make it hard inside. So every single machine on the computer, on the data centre must be protected. So we enable this, again, with our BlueField technology that actually implements SNAP as well but what it does, it enables to run security policy on the BlueField Card that runs its own operating system, and it's not the application server. No one will get there.

You can think about the some sort of - this network that connects these guys, this network is protected by the people with the machine guns, okay. It's not network that is connected to the application network, it's not network that's connected to the data network. There is no connectivity between this software and the software that - and the compute domain of the applications. So with the BlueField technology we can segregate, completely segregate the infrastructure computing tier and the application computing tier.

Besides security, by the way, it has additional advantages. At the table at the breakfast we talked about updates, software updates of various things. Some things it takes years, or in hundreds of thousands of dollars or Euros to validate the upgrade. Here we can upgrade your compute server and your infrastructure compute server, completely independent, they are not connected to each other; they're independent. This way you're both more secure and much more efficient in terms of operation.

So this are sort of, you know, those are the tips of the things that we can take to the 2020 year and on where we'll have billions of people connected, and even tens of billions various devices connected that will require, actually, [sensation] of the data centre architecture all the way to the base, or to the telco, to the base of the antenna, to the self-driving car, to everywhere, all the way, even maybe to the smart watch, okay, to make sure that bad software that you could put on your watch doesn't compromise your watch operation, or the data that's being collected on that watch.

There's going to be much more applications and, again, the data; more and more data is being generated, more and more data will be available, and more and more things we can do; $4 trillion of revenue opportunity. It's all out there, we just need to take advantage of this.

Thank you very much.

[Applause]

Manek Dubash, NetEvents

That's great, thank you very much, Michael, and Jean-Baptiste Su will now interview you in his own inimitable style.

Audience Q&A

Jean Baptiste Su

Hello, everyone. Hello. Fantastic. You get it? Thank you, Michael, for this wonderful presentation. Hello everyone, so I'm Jean Baptiste Su, I'm a Principal Analyst at Atherton Research in Silicon Valley. I also write a column on Forbes about disruptive technologies in the enterprise IT space and cloud. So I think we're right on topic. But before we start, I had to wish you a happy 20 year anniversary.

Michael Kagan

Thank you very much.

Jean Baptiste Su

So 20 years ago, you left Intel, you were building, you know, chips and the Pentium; you were one of the architects of the Pentium chips and I know we are a very, kind of a - I see a lot white hair. So we all know what a Pentium chip is. So 20 years ago, you left Intel to start Mellanox and become an entrepreneur. What were you thinking?

Michael Kagan

What was I thinking? I thought that it's - this idea of computing being a service actually trickled through our minds pretty much those days. People that - we believed that network is - becomes the central part of the world, of the communication, and few people didn't quite agree with us on this but it was great days, you know. It was before the bubble burst, so it was not as hard to start company as it is today. We were a bunch of guys that worked at Intel for ever, for a long time. We are a very [unique] company in the sense that we had the nine founders and all the founders have stayed in the company for more than 15 years. We decided to start the company to take advantage of this thing; of course, we didn't realise and [generalise] all the details but it's so...

Jean Baptiste Su

So you started by doing processors and then in Mellanox you switched to doing networking?

Michael Kagan

That's right. That's right, so yeah, at some point, as somebody told me that I decided to make something useful out of the processors we were building, so...

Jean Baptiste Su

I know, you see, somebody else thought that that was very useful and just a couple of months ago you sold Mellanox for, I think it was $7 billion to Nvidia...

Michael Kagan

[Unclear]. That's right.

Jean Baptiste Su

Yeah, so why do you think Nvidia was interested in what you were doing at Mellanox?

Michael Kagan

Well, what Nvidia is interested in is - I think the best would be to ask Nvidia. They probably know what they're interested in. We worked with - but we knew Nvidia for quite a while; we worked with them for, I think, maybe a decade or so. Actually, they've the fastest artificial intelligence machine in the world, which is the Oak Ridge Lab supercomputer built based on the Mellanox technology interconnect and Nvidia GPUs and the IBM build this [series]; they basically built the machines but if you look at most of the computing in the AI space it's done on the Nvidia.

So they know us for a decade. Our technology enables much more efficient scaling out of the GPU computing and it's proven time and again and the Oak Ridge machine is one of this proofs. I think there is a lot of - on the technology side there is a tremendous amount of [synergy]. But, you know, about the details exactly why did they buy us now and not before, and so on and so forth, you can ask Jensen.

Jean Baptiste Su

Yeah, I did actually but he was busy with his fast cars.

Michael Kagan

Oh okay.

Jean Baptiste Su

So you were mentioning about the Spectre bug in...

Michael Kagan

Spectre bug, yeah.

Jean Baptiste Su

...in the processor and Intel [suites]...

Michael Kagan

It's not the bug, it's...

[Over speaking]

Jean Baptiste Su

...so you're an expert in building processors, so did you know at that time that this was a possible failure?

Michael Kagan

No, of course not. I mean, we didn't [take our] processors knowing that there is a security hole in there but, you know, these machines are fairly complex and at some point people figure out how to integrate, if you will, various characteristics, [capabilities] of the machine in the way that you - that will make something that you didn't think about; and that's the Spectre bug, okay. On the nutshell, it's the understanding of cache behaviour of the CPU, which is the core technologist in the CPU design and taking advantage of this by figuring out who does what and getting the data leaked out. So, I mean, I'm sure Intel will find a way, or Intel - it's actually not on Intel, it's all...

[Over speaking]

Jean Baptiste Su

Yeah, you have ARM, you have AMD...

Michael Kagan

Yeah, yeah, all the - it's not something has to deal with the - even with the [x86] instruction set. But, you know, CPU makers will find a way to mitigate this bug but there will be another one.

Jean Baptiste Su

Right, well, you heard here first. So yesterday we heard from VMware about this evolution of the compute infrastructure, about a data centre. So we were, as you were saying, you know, we had a host, [a base] systems, we have virtualisation. Now we are talking about containers, server-less. So how do you - how is your vision - how does your vision fit into this next wave of computing infrastructure?

Michael Kagan

Well again, it has to deal with the ability to concatenate, or to chain the services, or microservices to deliver the application because, as the end of the day, all these things run on some physical machines that are built in a - not in a particular way to serve a specific application, but they're built at the way, like, you know - if you want to have a many chip use serving a single application, not many people, if at all, build servers that contain hundreds of GPUs. Now how - but if you...

Jean Baptiste Su

Nvidia does.

Michael Kagan

Well, yeah, Nvidia [show you their biggest box] and it's not 100 GPUs. But the point is that if you can take advantage of hundreds of GPUs, now you need to reconfigure this network, or these services to take advantage of this and this is where network is. Actually the network is the execution engine of the data centre operating system. The data centre operating system is about assigning the resources to the applications, or to the clients and in order to make them work together, you need to configure the network at a fraction of the seconds so people will not get bored because - before they got the response; and that's what network is all about.

So [let's say about] how do you configure the network? How do you deliver the network connectivity? Or are your resources to the application containers that you mention is just lightweight processors that many of them require it to deliver a service that - a distant service to another and they can be physically located on different machines just is a matter of machine utilisation and the viability of the resources; and now you need to make them work together as they would be a single - on the single machine, and that's all about network.

Jean Baptiste Su

Right, so are there any questions in the room just raise your hand. So we're coming into this new - it's funny, you were talking earlier at the beginning of your talk about smart phones and we are now into - beginning into the 5G revolution and one of the things that 5G is going to bring, aside from higher speed, is almost real time responses. So how do you see 5G changing the cloud paradigm? Are we seeing now - are we going to have - basically the idea was maybe having the cloud - you were talking about having it maybe on your smart watch. So how is that going to change?

Michael Kagan

Well, yes, I think that these type of opportunities actually drive cloud providers outside of the data centre. It is for various reasons, you know, starting from the simple physics, okay if we want to have things like decision-making at the fractions of the seconds, like for self-driving car or other things, then you need the decision-making point to be physically closer to the place that effects the decision. If you have multiple sensors that are physically located, one in the traffic lights, one in the car, one in the paving and one on top of the roof that all these sensors, so data is coming from the sensors that are aggregated to some point, which is a few metres away. You can make or you want to make a decision based on this data, you actually need to build the [incidence] of the data centre that makes the decision out there. It's in the base of the antenna, it can be in the car itself, it can be anywhere. It drives the data centre services and data centre paradigm outside.

Now the point is that in order to build the things, you don't need to reinvent the wheel all over again. This efficiency of accessing the network, and the security topics that we discussed, they are pretty much the same no matter where you are physically located, if it's the antenna based, if it's the autonomous car or if it's the huge data centre like Google, Facebook or Amazon or whatever Azure.

It means that, if you take the same architecture and you take the same or subset of the same software that used to run the data centre, developed by people like Microsoft, Google, Amazon all these guys are actually providing their data centre management tools. As you can buy them as a service, and now you can run your data centre inside the car or on the base of the antenna or any other aggregation point using the same tools. Now it's a matter of the implementation of different pieces of hardware to feed the dimensions and the power and the cost of that particular implementation.

But it turns out that software is much harder than hardware, and much more complicated than hardware. Being able to take advantage of this development in the data centre all the way down to [H] that's what IoT and 5G is doing. Actually, for companies like ours it presents a great opportunity, because we have the best engine for the data centre operating system out there, and now it creates a huge volume for these types of products.

Audience Q&A

Jean Baptiste Su

Great, so I know we have a question, if you can tell us your name and publication.

Joel Stradling

Sure, thank you. Joel Stradling from GlobalData. [Unclear], very interesting presentation. First, I have a little joke I want to share, and then I'll get to the question and share an observation and it's this. How many software engineers does it take to change a light bulb? The answer to that is they don't, because that's a hardware problem.

Yesterday we had Joe Baguley from VMware, and one interpretation of his presentation is that the application becomes a network. Looking at your presentation it seems as if you're expressing the view that network becomes a computer. In these three separate areas of app, network and computer, perhaps the app becomes a computer. That was the one question. Is that a possibility?

But then the inherent vulnerabilities in security are because of those three separate areas today, so I can understand how the network and computer part can become more secure. But with apps running on top, it seems to be we will always have these weaknesses because of that software/hardware divide.

Michael Kagan

Okay, so first of all network is not a computer, network is the essential part of delivering - being able to deliver services. Network can and does some computation, even computation beyond just moving the data. In a sense, in terms of security and in terms of all those things, there is many ropes here that you can hang yourself if you're not careful.

But the trick is there is a huge amount of many, many, many little things that you need to get them done right in order to have a solution. This is what it is all about, is to make all these little things right. I think we as a company we have figured quite a significant amount of things out of them. I am not claiming that we know everything; there is always new things to learn and to do. We are climbing up the stack since we actually started, and we always discover a lot of new things and opportunities that we can take advantage of.

But back to your comment, network is a part of the computing, an essential part of the computing of data centre or all this computing revolution, data processing revolution.

Joel Stradling

Thank you.

Jean Baptiste Su

One more question, yeah we have two more, if you can tell us your name and your publication.

Antonio Dini

Antonio Dini, Antonio Dini from [Corriere delle Comunicazioni] in Italy. You had a little moment when you were talking about Spectre the first time, you said it's not a bug, then you mention it as a bug, I would like to have a better perspective of what you consider those chip problems, bugs or what?

Michael Kagan

Okay, so it's a terminology topic, so let's bring it up. From my perspective, bug is something that I was expecting machine to work this way and it actually works a different way. If you take it high enough up, and saying I was expecting the machine to be completely safe and it's not safe, then it's a bug. But it is not a bug in a sense that it is not like a hard bug, that there was a softer bug that didn't check - the software didn't check the boundaries of the access right.

Here it is like if there is some other site channel attacks that I heard about based on the temperature. For example, the fact that when there is many people in the room the temperature is higher, and therefore I can figure this out by just looking at the temperature sensors. Therefore, I can - from that point on it depends what room it is and what type of people are getting there, although I get to the security issues. This is the problem with - this is the type of the problem with the Spectre, so it's taking the basics of CPU operation and finding the way by observing what's going on to learn about what's going inside.

Jean Baptiste Su

One more question. Antony.

Antony Savvas

Antony Savvas, ITEuropa and Data Economy amongst others. The data centre business is struggling. They want more automation. They want more artificial intelligence. Why didn't Intel buy you? They could have done so much with your technology. I know they were interested, why do you think they didn't get a bid over the line and buy you. They would have been a more natural fit for your technology.

Michael Kagan

Why Intel didn't buy us, you should ask Intel.

Antony Savvas

That's a [beautiful] answer.

Michael Kagan

Why didn't Intel buy us? First of all, I think the price that we got is much lower than it should have been, but that's outside of my jurisdiction. I think we have a huge value in terms of - our technology has a huge value in this new 21st Century data processing story. I can't answer your question why Intel didn't buy us.

Jean Baptiste Su

Well Antony, I can answer your question. They were on the list, so they were - but unlike what Michael said, they thought it was too expensive, and so they didn't bid up, but they were in the bidding process. I can tell you that.

Michael Kagan

Are you saying they were cheap?

Jean Baptiste Su

Yeah, well they're a chip maker.

Antony Savvas

[Unclear] $2.50 a share.

Jean Baptiste Su

Well thank you so much, I know we're having the - we have to go. But thank you so much Michael, thank you for your time. Thank you all.

[end]