

## NETEVENTS GLOBAL IT SUMMIT

SAN JOSE, USA

OCTOBER 3 & 4, 2019

**DRAFT**

### *How long is the road to AI?*

**Keynote Panel led by: Jeremiah Caron, Global Head of Research & Analysis – Technology Group, GlobalData**

Panellists:

Ravi Chandrasekaran	SVP, Enterprise Networking Business, Cisco
Nick McMenemy	CEO, Renewtrak
Prof. David Cheriton	Stanford Professor, Investor, Entrepreneur

#### **Jeremiah Caron, GlobalData**

Hello everybody. Thank you for that excellent presentation Ravi. Very, very interesting. Particularly like how Cisco and other companies as well are now finally truly integrating security as well as the network [unclear]. That's always been the issue.

I think a lot of you know, and for those who don't, I'm Jerry Caron. I'm the head of research analysis for a company called GlobalData. We look at - we provide advisory services and market data solutions across the whole tech and telecom spectrum. That's who I am, that's my role. What I'm doing today, what I'm doing now, is I have the great pleasure to lead a discussion with the three panel members you see here, generally on the subject of AI.

It's a bit scary when I saw that. I was like whoa, it's really going to be impossible to talk about AI in 30 minutes. What we really need to be doing is to drill down. We'll talk a little bit about definition around that term. The objective here is to talk about the premise of AI. What is actually being delivered? What are the tactical, practical uses today, and what we can see in the short term going forward? Also, what problems AI

is solving today. Obviously, Ravi touched a bit on that just now in the networking context.

Just a quick thought from me before I turn it over to the panel members. We do lots of work on AI related topics. You talk about a data science, you talk about machine learning. Computer vision is a big one. Really a very recent study just came out last week. Large global study of enterprise decision makers. Not just IT people, but also technology decision makers generally. There's huge interest in these areas. Huge expectation on spending money in these areas. Over 70 per cent in data science, for example. Nearly 70 per cent in machine learning, and over 60 per cent in computer vision type applications.

There's a recognition that this is going to be really important. What I found fascinating was the same study came back with less than 50 per cent. Surprisingly low number - 47 per cent to be specific - only that amount said that AI would be very important to their businesses going forward. You would think it'd be a lot higher than that, wouldn't you? I know why that is. The reason is because companies don't necessarily think about AI. They want to think about what AI will do, so apply the AI. That's actually what they care about. That's just a couple of thoughts there.

Let me quickly introduce the panel. I think the best way to do it is to have the panel introduce themselves. Starting directly to my right, let's start with Nick. Could you please let people know?

**Nick McMenemy, CEO, Renewtrak**

Hi, I'm Nick McMenemy. I founded a company called Renewtrak. We are a fully automated white label business which pursues support maintenance and licensing renewals. Sounds really turgid and very dumb, but at the back end here - it's interesting, I spoke before about this. We're using a lot of machine learning. We don't use AI per se. We can talk about that when we get into discussion. We're using a lot of machine learning to drive better renewal rate. Ultimately drive more revenue by using smart automation for the companies we work for.

**Ravi Chandrasekaran, SVP, Enterprise Networking Business, Cisco**

I'm Ravi Chandrasekaran. I am responsible for enterprise networking engineering at Cisco.

**Jeremiah Caron**

Okay, David.

**Professor David Cheriton- Stanford Professor, Investor, Entrepreneur**

David Cheriton, I'm a professor at Stanford University, and actually not responsible for anything at this point in time.

[Laughter]

I've been at Stanford for about 38 years and worked with AI people, but I'm more of an AI sceptic. I've been involved in networking. I started a company that's got bought by Cisco. Spent some fun years at Cisco, and then I was co-founder of a company called Arista Networks which I guess has gone on to try and compete with Cisco. I've been a co-founder of a company called Apstra which I think is represented here that's going on to tackle next level of management on top of Cisco and Arista and everybody else who switches to try and improve the network management situation.

I think my role here is to be the AI sceptic.

[Laughter]

### **Jeremiah Caron**

All right. We'll talk about it in a second. Just for those of you who don't know, David's been slightly humble. In addition to all those great things that he's mentioned, he also was one of the early funders of Google, and also of VMware and, of course, as he mentioned, founded RISC with some others as well. Great honour to have you here with us David.

Let's dive into the conversation. Let's talk about this AI scepticism. Nick, I know you want to touch on that too. I think it comes down a bit to the definitions of what we actually have going here. First Nick, David started the conversation on scepticism. Why don't you give your thoughts on why you're sceptic, and I know Nick wants to jump in on it as well.

### **Professor David Cheriton**

I think one of the advantages of being older is you've been through a bunch of cycles, business cycles and market cycles. Well, I've been through a bunch of AI hype cycles. Back in the '60s, the view was we're just on the edge of producing machine intelligence. It kind of died out, and then didn't go anyplace. In the late '80s, the view was well expert systems were going to take over the universe. There were some very impressive projects done at Stanford there. There was great concern that the Japanese with their fifth generation computer projects - some of you might be old enough to remember that - were going to just dramatically leapfrog the US in technology.

What followed from all of these, there's a term referred to as the AI winter. When people have asked me over the last 35 years in Stanford what I thought of AIs, I say well it's a very promising technology. It's been promising ever since I encountered it and continues to promise. I think it suffers from being over promising.

For me, this is, as Yogi Bear would say déjà vu all over again. You look at the excitement over AI, you look at all these articles and say well people are going to be thrown out of jobs and AI machines will take over the universe and so on. I've been there, seen that and then what follows is an AI winter. There's a certain basis for scepticism just on experience.

**Jeremiah Caron**

Thanks. Nick, I know you have some thoughts on this system as well from your perspective.

**Nick McMenemy**

I agree with David, and I also think we've got a - the hype cycle is an inevitability of the tech industry. We live through that every day. I think we've also come to use the term AI as a catch all for everything which is a machine is going to automate something. It's interesting. I turn up and I'm quite specific. I talk about we do renewal to automation, and there's some smart stuff there. A load of people go oh my God that's AI. I'm loathe to agree because I just don't see that.

There is a place - to David's point - there's a place where some of the tech associated with automated stuff in a smart way may have an impact in freeing up humans to do stuff they really want to do. Ultimately all this automation, all this machine learning stuff they're trying to do is to make what we do every day easier, faster, better, more accurate, et cetera, et cetera.

I don't see AI as the great [unclear] which is coming this week, next week. I laugh to a degree, and we had this conversation before. You see some companies using the term willy-nilly, and there's probably a round of fund raising due, or they've had the analyst's report, or they've got to go to The Street next quarter and announce bad results. All of a sudden AI is flavour of the day.

To a degree I have scepticism, but I've got reality. I don't have a hedge position. I'm not with CISCO, I'm with Arista. I can't go and buy another company just in case AI's the next thing. To a degree I'm sceptical, but I'm also realistic that there are elements of what people are calling AI which are available in the industry. Certainly, what we do is - people are using it. It works and it's machine led.

**Jeremiah Caron**

It's an umbrella term used to describe lots of different technologies, and lots of different applications that in turn use those technologies. That's probably confusing in a lot of ways. Ravi do you have a comment on this?

**Ravi Chandrasekaran**

Yeah. I think it all has to do with what we mean by AI. It's not going to take over everything we're trying to do. In CISCO, we have a very simple view. We view these as tools that assist in certain things which we will not otherwise be able to do. For example, when I talked about the closed loop system, there is a perception or perceived notion that automatically somehow AI's going to make the closed loop happen. We don't believe in that. That's why I had in that picture I was showing, an IT system in the middle the policies and procedures that people do and you need to take that into account.

We look at AI. I even use the term internally it's more like an applied AI. It's not my - all of us are changing something. We're just using that as a tool. For example, if you

look on the encrypted traffic, you need some assist. You can't easily find out what is inside. We look at that which I think you're - the moderation start setting in. The reality will set in, where the reality is somewhere in between. It is not nothing, but it is not the glory that people talking about. It's somewhere in between. We are starting to see some good results of that. We believe these to be tools that can be used in [VMware] and the hype cycle will settle down. I think that's what should happen.

### **Jeremiah Caron**

Despite the scepticism, David, would you agree that there are things that we need to do now like cyber security, and perhaps some large Tokyo style networks that absolutely demand what we would generally call AI solutions?

### **Professor David Cheriton**

I think again there's a lot of value in figuring out where the pony is in this pile of hay so to speak. Having lived with AI, people speak for forty years, their interest is replicating human intelligence. My view us humans are relatively stupid, relatively slow, and relatively unreliable at almost any given task. If you look at the history - we used to have people computed trajectory tables. They did this math in their head. We thought that that was intelligent. Now we have computers that do it a billion times faster than human beings. You look at chess. We used to think that was an intelligent activity. Well, you know of the seven billion people on the planet, computers can beat every one of us. This objective of human intelligence is a little misguided.

I think you need to shift it as Ravi was saying, and say is there a pony in there, and can you apply it. That is when you apply good engineering principles, does this technique apply. Is this the best approach? When you look at security, playing my role as a sceptic, there's an article in a technical journal recently pointing out with machine learning how easy it was to fool it.

In one of the examples, it showed two pictures. One of a stop sign, and one of a stop sign. Just the one on this hand had been modified by an adversary, slightly. For a human being they both looked like stop signs. For this system, a vision system had been carefully trained with deep learning to recognise stop signs and traffic signs. This was a stop sign, this was a yield sign. I think wait a minute. How did it get fooled? Because it's using statistical inference. If you statistically perturb the image in a way which is not detectable really by a human being, you can fool this training that this thing is going under.

In many, many cases, we want predictable performance. In fact, I'd say that's what engineering's all about. You build things with predictable performance. I used to tell my students the reason people want predictable performance is because unpredictable is so horrible. If you have a system where it thinks - sometimes it thinks that a stop sign was a yield sign, and you try and put this in a self-driving car, you have a big problem. In some ways it's a security problem.

I think one of the things we have to tread very carefully on is that security is not a one-time event. Security is [unclear]. You do this, I'm going to counter with that. You do that, I can counter with that. With machine learning, we don't understand how to

counter adversaries. The striking thing about this article was it was written by a Google expert in machine learning, and the end of the article had no answers to this problem. I think we can't ignore these techniques, but we have to tread carefully.

### **Jeremiah Caron**

Just if I could segue real quick. You quickly get into the human discussion in any conversation about automation or AI. I think that now that we've all agreed, with David's encouragement, that humans are pretty bad at a lot of things. Ravi, particularly perhaps secure network management or automation in today's world perhaps Cisco would argue. What's that mean for the army of Cisco experts out there and their future and what needs to happen. What's the human side from your perspective and then Nick I'd like to hear what you think about that as well.

### **Ravi Chandrasekaran**

When we look at any IT organisation, they are growing. They are riding hundreds of people to run the network, or the digitisation that's happening. What we see is that the tooling they're providing the automation, and all the security related in the providing, is basically to free them up from dealing with the complexing that comes along so that they actually can get their job done. We look at all this automation, and we talk about all the security infrastructure building, the whole help [unclear] we are doing with [unclear] is around to free people so that they don't completely buried with the burden that's coming along with the number of devices coming on board, or the [unclear] that is being [unclear] world.

To us we feel this – when we talk to any of the CIOs or IT organisation, they're so happy providing the tooling because they thought that they cannot keep up. They cannot run the infrastructure, it's hard. They're getting more people. Some company goes and say I'm going to digitise, I'm going into a new world. You cannot say all of a sudden I'm going to double my IT budget. IT budget still continues to stay where it stays. You need these help. We look at it and then we talk to our customers. They understand it. They understand these are the tools that you need to have to run the infrastructure. It's not trying to replace people. It's not trying to reduce the people. It is freeing up so they can be a business [unclear]. That's how we see it.

### **Jeremiah Caron**

And presumably the skilling, both at the early stage and educational environments and also the later more - within your reseller partners et cetera. Presumably that's well under way.

### **Ravi Chandrasekaran**

Absolutely. You need Cisco programme whether it is [unclear] or whether it is CC, CCNA. All those programmes are being modified to reflect the new world. It is about more how do you do automation at scale? How can you programme that company use the API? Things of that nature. We are putting a lot of emphasis so that we can build a new toolset along with the education and the help they need. That's our way of looking at it. It's more augmentation. It is not replacing people.

**Jeremiah Caron**

Nick, from your perspective you've a very practical application using machine learning techniques. Your customer base, is it automatic that they get what the value is? Do they see the value or is that a process that you have to go through as you recruit new customers?

**Nick McMenemy**

It seems to me, what I've found in the business, it was 100 per cent sceptics. They didn't understand the value, they didn't believe the value. It's David's point before. They were I'm not really sure this is for me. Inevitably, they understand that the automating something to support a human. In my case, in some of what we do, we do replace the human element to make it more efficient. [Unclear] not left behind.

We see, and now our clients are increasingly seeing it. The majority of our clients at the moment are tech companies. They see that they've got to do a better job. The predictability is a mandatory, not just for the CFO and the street, but then clients want that predictability that something's going to happen. The staff want predictability that they can understand something has taken place. They can apply maybe what they've learned if this, then that next logical purchase, they now start to understand.

We talk about the Nirvana of cross-sell and upsell. I think pretty much every tech company I go to after they tell me about cloud computing, they tell me about cross-sell and upsell. It's that contingent if this, then that. It's not difficult. It's a simple if and, or if then statement. It's not complex to articulate. For many of them it's really been difficult to make it work. They've got humans driving that process. It was scepticism, and now I think people - they get it. When I talk about this stuff to people I say you know what, let me just show you. I show it and you have that lightbulb moment because that the hype - it's not hype. It's real. It's a practical application. There it is.

**Jeremiah Caron**

I just want to - a conversation we had at breakfast - David I'll catch you up real quick here, and throw it to you actually first. We're moving along, there's scepticism about what AI really is, but generally it is seen as making its way into applications today. What do we need to be doing about compliance and governance? Is there enough attention being paid to that? In your view, is that something we need to be worrying about now or is it all cool?

**Professor David Cheriton**

I don't think that - I'm not a big fan of regulation. I think we have enough already. To cite a recent unfortunate example, if there'd been machine learning involved in the 737 Max 8, and Boeing couldn't explain why planes were driven into the ground at 500 miles an hour, it would be viewed as irresponsible engineering, in the same way as if the wings had come off. I don't think any of these things introduce new requirements for regulation. I think they - we have enough regulation. In fact, I think we have more than enough regulation. I think that we just all need to be reminded that when we're

building real products that we can't just have a magic black box and defer everything to that.

I think responsibility still rests with the company, with engineers, to make good decisions. I think, again, I come back to predictability. If you deploy something that's in a safety critical environment, and you can't predict how it's going to behave, I think that has to be viewed as irresponsible. Whether it's putting a baboon in the pilot seat, or whether it's putting machine learning in control of some vehicle. I think we need as a technical community just to be cautious that the world hasn't changed. In effect we're still doing engineering.

### **Jeremiah Caron**

Very quickly, I'd like you guys to come in on this too. Before I do that, I just want to know - we're going to have a few minutes. If anyone has some questions, think about them, and we'll come out to you in a couple of minutes if you have a question of these guys as well. Any thoughts on...

[Over speaking]

### **Male**

...because if you look at it, you don't want to keep closed book systems which are autonomous in nature because we are far away from that. You got a lot of interesting side effects. There are poisoning things that could happen, that are unpredictable. I do believe that you need to have tools which tells what's going on. You need a high level system which takes care of the policy and other considerations which activates it in a way where you can do it in a responsible way so that you don't have a system which is off. We believe in that and that's part of reason we always talk about every part of the IT process.

There already things in place. There's some infrastructure where you can have policies ... let it decide based upon the policies and rules of different organisation then you have the right way of doing the ... but it's not really the ... we have human intervention that says that's what we believe in.

I'm not a bit fan of auto remediation. In networking that's really crazy if you make a mistake you can bring the whole thing down. [Unclear] talked about. I would rather give hints and enough insight into what is happening so a higher level intelligence can come in and play a role. That more often includes policy because different organisation need to do things differently. That's what we believe in, and that's the path we are going in.

### **Jeremiah Caron**

Nick, anything there?

### **Nick McMenemy**

Everything we speak about when we talk about this has got data is the DNA. Ultimately data is the DNA that makes this work. I don't believe we need more regulation or



anything that. I just think we need to probably adhere to what is currently out there. You got those ISO standards, you've got those things. Core things like data, and the security, and the policies around the management of that, I think we're probably good enough. I think we're so [unclear], I just don't think we know enough to know what we should regulate, and how we should regulate it.

We're talking about AI, we talk about this machine learning this stuff. Really what it is, is distilling the essence of what a human has asked the machine to do, and making sure that it does it in a way which is predictable, very viable, and all that type of stuff. That doesn't require any regulation, it just requires adherence to the rules the company have been living by.

### **Professor David Cheriton**

I have one comment here. Besides figuring out where to apply this, I think we're still trying to understand where we are with some of this technology. One of the examples I like is there was a lot of excitement that a colleague of mine [unclear] at Stanford released this collection of images. Something like two million images. There's been these contests, if you like, of having machine learning systems recognise these different images. There was a lot of excitement of the accuracy of that going up substantially. For instances, it could tell - it could recognise a wolf versus a dog.

What I understand recently is people discovered that the reason it was telling the difference between a wolf and a dog was that wolves were in a background that was snow and bushes, and dogs were in a background that was grass. The point of this is that we thought we were further along in solving this problem than we were.

Another example that came out was that it - some work at Google demonstrated that in recognising a zebra, it was very significant in the machine learning that there was grass involved. You can have a system which would recognise a zebra, but if you put a zebra in a parking lot it wouldn't recognise it.

I think we're still discovering these sort of effects that are hidden in these data sets. I'm not faulting anybody, I'm just trying to illustrate that we can think we're further along that we actually are.

### **Jeremiah Caron**

That gets to the point, Ravi, you were making earlier about the key thing the data being data, and the data fidelity and integrity, and then it builds from there. Any questions from the audience? Scott?

## *Audience Q&A*

### **Scott Raynovich, Principal Analyst, Futurium**

Scott Raynovich, Futurium talking. Couple of interesting points. Jerry you made an interesting point about applied AI. It seems to me like AI works with very simple applications. Sometimes it doesn't work. Siri to me still doesn't ever work. Shazam is amazing to me because I'm a music fan. It seems to be flawless, and recognises a song within a number of seconds. As we get up in complexity, it seems to be where AI falls apart. David was referring to the Boeing mess, which actually seems more like an engineering and management problem, rather than an AI problem. An aeroplane is more complex than a search engine, or a music search engine.

Then we get into cloud and we're seeing this presentation about AI and networks. The complexity increases by order of magnitudes. You have machines talking to machines on different levels, different domains. Cloud domain, virtualisation domain, hardware domain. That's where things fall apart, especially with inter-operability rate. When you talk to IT people, their problem isn't oh I have this magical AI. It should fix everything. The problem is getting things to talk to each other, and having the data sources interoperate so that you can apply the AI.

My question to the panel is are we just going down the wrong rabbit hole? It's really about getting the systems to interoperate and have clean data so that we can build AI machines to actually do what we're talking about. That's not happening today.

### **Male**

I can comment just really around the data cleanliness problem. The challenge we have got is just data is not clean. We look at certain counters and things that's fine, you cannot compare if it has come from different versions of the same operating system or different products, and they're not compatible so there's a lot of challenge that's there. We are facing that because we are just – as you start using data [unclear]. You start realising those play a big role. The small deviations lead to a very different conclusion unless you take those into account.

So far as that's a big challenge in how do you bring data which is clean, consistent, they are compatible, they are what they are across different versions of software, different products, and even different vendors, and different parts of the infrastructure. That's a big challenge. Data fidelity, and ability to use it consistently and how they bring this data together is a far bigger challenge than applying some simple ... to that. That trips us often. What we think is the right answer is not necessarily right as you go through different data. That's a big challenge. That's why we have to be very careful. We thoughtfully it. I keep using the term it's an assist. It's not a replacement exactly for that reason.

**Male**

I would dispute your characterisation that AI runs out of gas, or at least the techniques we're talking about with complexity. I think if you consider, say, the problem that Google tackles of how do you do a search for images of cats for instance. It's a complex task to figure out what is it. Starting from pixel level, what's a cat? They successfully applied that. I think the - and it's sort of pretty predictable in the sense that most of the images you get back are cats. I think that as machine learning is a statistical technique it can handle extremely complex things. When you have enough data to train it on, and also when being mostly right is good enough.

With cats you say well gee this is great Google. I can search for images of cats and nine out of the 10 pictures you've shown me are cats. I think the area where it's difficult is when the statistics are working against you, where you say I'm looking for some anomalous situation. If it's only right nine out of 10 times, and I'm having events occurring at basically trillions of times a second in a large network, the statistics work against you. If that filters all the way up to alerts or any action by an operator, it leads to being overwhelmed by alerts.

I think even Google has a great technical paper where it describes the difficulty of composing machine learning systems because we compose an engineering system by having clear well-defined interfaces, and have semantics across them when the internal guts are something that does unpredictable things, we can't compose them in a reliable way. I think that there's a composability problem we haven't solved.

**Manek Dubash**

We have a question over here. I think we have one here too.

**Steve Cassidy, PC Pro**

Steve Cassidy, PC Pro. There's a very interesting gap in the topics you guys have been covering which evolves from observing the problems with things like bitcoin. That is that AI seems to be counted as a free resource. Nobody has any sense of whether a solution is expensive either in compute terms or in dollars and cents sitting on EC2. Or whether it's actually a really cheap easy fix that just happens to be quite intricate and quite a lump of logic. How can businesses predict the costs of AI in their development exercises?

**Jerry Caron**

Any thoughts on that?

**Male**

I think the cost of running it is highly predictable because it's basically just doing classification in real time, so to speak. You can determine that just by building a system and running it. I think the data collection there's at least - you can scope this reasonably based on experience. What I think is harder is to understand in certain domains how much it's going to cost you to maintain it or deal with unpredictable problems I think

that's where you end up being - that third piece I think is the part that I think is harder to capture just - if we build this thing, how many times are we going to have problems when it does something we didn't expect.

**Jerry Caron**

We have, I think, time for one more question. This gentleman here at the back. Yes?

**(Unknown)**

The state of art appears to have moved on beyond deep learning or statistical learning. Instead of trying to find [patterns] in noisy data, the shift is happening towards more casual markets reinforced when learning and so on. One learns from people who actually use AI. Is that helpful?

**Jerry Caron**

I'm sorry, I didn't understand what you moved on to. Could you repeat the question?

**(Unknown)**

The question is that instead of finding patterns in data which is what deep learning does, a new stage is emerging in AI which looks at causal relationships to narrow down the outcomes, [unclear] outcomes, so you don't have too many alerts and you understand why something is happening. I'm trying to understand if that is helpful.

**(Unknown)**

I'm seeing the evolution of that, but I fundamentally - I live in a different world. I live in world of practical reality and the application of this. Fundamentally, if you're given them a [data lake] which is essentially what a lot of this machine learning is applied to, you need to know what to look for.

The point about the cat before, or the zebra. If you don't know what you're, looking for, you don't know what outcome to expect. One of the things we were finding a lot of time is being spent, particularly in our business, is educating our client in what is it you are trying to achieve. What are you looking for? The cat, the zebra. I'm looking for the zebra. Are you actually looking for the zebra or are you looking for the black and white lines?

Having that conversation has moved on, to your point, from that deep learning, it's moved on. But we haven't made that quantum shift that - Shazam has probably done us a favour, and it's also been a bad thing. Because it showed us that you can apply some smart stuff, and it does work. But it's set expectation in an exploit sense, that oh, this is really easy, and it's commoditised, or it's cheap or it's available, or whatever it might be, and it's not that simple.

I live in a world where nine out of 10 is not good enough. If I get nine out of 10, I get 90 per cent of what I do right. For companies like Cisco, that leaves a few billion. Probably quite serious, so one's focus has to be very specific. It's a rules-based order.

Ultimately, everything we hear about - in machine learning, the machine learns what it's supposed to learn, if you can understand what it's telling you at the end of the day, against a set of rules you've defined. As a child, you grow up, what do you do? You learn as you go. This flame over here is hot. The rule is, don't put your finger in that flame. The cause is, I'm going to burn my finger. Or, the effect, sorry, is I'm going to burn my finger.

All that stuff, to your point, it's not a quantum shift from where we were maybe 18 months ago, even, 24 months ago.

**Jerry Caron**

All right, I think we're going to have to wrap up. Thank you very much. Please join me in thanking David, Ravi and Nick for their comments. Thank you very much.

[Applause]

**Manek Dubash**

Thanks to the panel, and thanks, Jerry for the erudite chairing, and insights.