

NETEVENTS

GLOBAL PRESS AND ANALYST SUMMIT

DRAFT

Opening Keynote Presentation

How to Pick the Game Changers

David Cheriton

Stanford Professor, Investor, Entrepreneur

Manek Dubash

Well I guess it's my cue to say good morning NetEvents. Good morning NetEvents. No, no applause please. I haven't ... bad jokes. Okay ... microphone. I hope you can all hear me at the back. Good. That's a good start. My name's Manek Dubash. I'm going to be your MC for the next couple of days. Now my job here is just to basically set the scene, introduce the sessions ... basically to give you the format of the event for those of you who've not been here before, this morning we're going to be having a [unclear] session [unclear] discussing the issues of the day. Issues ... coming out of this event are Artificial Intelligence, Security - this is a standard topic I guess ... that's an overlay on everything we do - and of course ... half of Silicon Valley for those of you who were at lunch yesterday, you'll remember money. It's all about the money. It's about where the money goes, where it's invested, where it's [safe] and what comes after that ... for the rest of the morning. You'll know from the topics that - oh, we're going to have two microphones.

How privileged is that? Yay. Who needs a microphone anyway? This morning's session will finish with a Shark Tank which as you will note there's no actual water up here. But the sharks will be venture capitalists who will be grilling potential start-ups, for start-ups who will be pitching for money. I think that could be quite entertaining. They'll

have five minutes to make a pitch to the venture capitalists. If they continue after five minutes we will cut them off vigorously. Oh yes.

So without further ado since we are talking about the money the first keynote speaker sets the scene. He is a professor at Stanford University who knows where the body is buried. He's directed lots of money in various directions for start-ups. He's going to tell us all about that, how it all works, where the mechanisms are and definitely where the bodies are buried I hope. Professor David Cheriton. Come on down.

Professor David Cheriton

I hope they're not expecting me to dance or something. Regarding where - if it's all about money I was going to - I had proposed to call this, how to invest like a university professor. But then I was told that I wasn't the typical university professor so I changed the title.

Let me dive - let's see - there we go - my actual title is Smart, Lucky or Prepared. I just want to talk a little bit about my investment history and thinking. Hopefully we can tease some interesting lessons out of this. I think there's a saying on Wall Street, I'd rather be lucky than smart. I think I have been pretty lucky along the way.

Right now one of the roles I play is I'm involved in a company called Apstra. I'm the chief scientist there. I'm a co-founder of the company. I am the investor. There's no outside investment at this point. So you might ask why am I doing another start-up? Why am I not a venture capitalist? I claim I'm not a venture capitalist. Then how did I get in the position to fund companies myself like this? Hey, university professors don't normally do this. So let me try and tackle some of these questions here.

Why Apstra? Well, I think Apstra is a good example of how I approach investing and start-ups. Some people say a start-up is a lot of fun. Well it's kind of - to me it's like one of the scariest rollercoaster rides you can go on where it's a lot of fun but it's pretty scary at times. So I look for more than one reason to do things, especially now when I still have my academic inclinations.

So with Apstra what I saw was major trends in the industry that were converging. That is, Apstra deals with the automation, the data centre networks. One of the things people are realising - there's a huge operating cost to running networks. There's a need for greater reliability in the networks and greater agility to adapt to new services and at a technical that there's now the ability to have APIs - two switches in the network - which means you're going to have programs talk to the switches.

Also people are building networks out of switches from different vendors which is breaking away from the single vendor model. That makes it more complicated to manage these networks. Of course your whole company runs on the network. If the network fails your whole company fails. If your network is slightly sick your whole company can be quite sick. It's a business impact.

Then finally there's proof points about the value of automating network operations out of companies such as Google that have invested a lot in automating their network

infrastructure. So all of these aspects converged to say there was an opportunity to do something in the start-up space there.

Then you can't do things without good people that have the right background. I was approached by Mansour Karam who I worked with previously at Arista and knew from Stanford as well as Sasha Ratkovic who has a great background in network management, was a fellow at Juniper Networks for a number of years.

Then finally - I think that you always like to have people bring something besides just a chequebook to a start-up. This is - the problems from a technical standpoint are really up my alley. He really needed to build what I call a distributed operating system, something that was distributed because a network itself is highly distributed. This has been my research focus for my entire career at Stanford.

I looked at this and said, I couldn't live with myself in the nursing home if I didn't try to make a company in this space. So it really was an opportunity too good to pass up on. I don't consider myself a venture capitalist even though I fund ventures with capital. At least - the people who are really venture capitalists are looking for a variety of investment opportunities. They're trying to maximise return with a given fund they're working with.

Their ideal business is one where the business has proven that as it's working you just need to scale it up. To scale it up you pour in money so you can hire lots of marketing, sales, more engineers and more infrastructure and so on. They invest across a whole bunch of companies and see which ones succeed.

I have good friends who are venture capitalists who have said, well we really aren't that good at picking the winners so we invest in a whole bunch. My approach is different because I seek out opportunities where I think they add real value to the world. It sounds a little hokey. But also I go for investment opportunities where I think I can add real technical value. In some sense if it's a big technological challenge, I think I'm up for it and I know other people who join the team are up for it, it becomes a competitive advantage because there's relatively few other teams on the planet that are able to do the same thing.

Then rather than sort of spray and pray in effect across many companies I invest the time and money in a company to make it successful rather than just see if it's successful. I treat it as kind of my baby, not just one of the horses that I'm betting on.

Then the final element is that I think given I don't have anybody waiting in the wings looking at my return on investment and when do I get the fund - the lifetime of the fund and so on - I can afford to be patient and produce a really good product that ideally sells itself. So I'm not ruled by time to market or fund ROI and so on. A number of the things I've been involved - when people forget how long the gestation period was from when it was founded to when it actually was recognised as a success.

So my investment view is that I don't put all my eggs in one basket but I put my eggs in a small number of baskets and try and take very good care of those baskets. I think that is also a philosophy I apply in technical areas as well.

So how can I afford to fund these companies? Well I think I've been very fortunate. I came to Stanford in 1981 as an assistant professor. It was a crazy place at that time. I encountered this student there, Andy Bechtolsheim who was doing this cool stuff building what was then the Sun workstation. He was doing neat stuff so I started working with him and developed some software for his hardware, found a hardware bug and so on.

Then I turned around and bought a bunch of the early Sun workstations for my research project. Then Andy went off and founded Sun Microsystems that was sort of a leading - the leading workstation vendor in the 1980s. In 1994 Andy called me up out of the blue while I was trying to do a sabbatical as academics do, and talked me into doing this start-up.

We both hated this technology called ATM at the time and thought Ethernet was the right way to go. He knew hardware. I knew networking and software. We put ourselves together and developed this technology. It was bought by Cisco. Actually Cisco did an amazingly good job. They basically just brought us in-house, gave us the freedom to develop this.

As soon as it was ready to ship we were shipping enormous volumes. I'm told that the Catalyst 4000 line that came out of this was the second most successful acquisition in Cisco's history. It ended up being at one point a \$4 billion product line for Cisco. It's second only to - I think it was Crescendo - which was Cisco's first step into the switch market.

So that was one, an exciting adventure. I actually spent seven years at Cisco as a technical advisor. Then overlapping with that in fact, I got a lot of publicity out of - here's a professor involved in a start-up that gets bought by Cisco for a lot of money and so on.

Back at Stanford I met this guy, a PhD student, Sergey Brin, because we were both rollerblading. He was rollerblading. I was rollerblading inside the computer science building. So we knew each other only as fellow rollerbladers. But at one point they came to me for advice about their search technology. What they wanted to do was licence it. So I helped them and pointed them at an IP lawyer I knew in the Valley and so on.

But I also advised them, this is your baby. You're going to need to raise it. Fast forward a few months later they found that - they went around to various companies trying to get them to licence their search technology including Yahoo. It didn't go anyplace. So they came back to me and said, okay you're right. We need to start a company to actually take this technology further.

They were running Google inside of the Stanford computer science building with a bunch of pasted together hardware and so on. They wanted me to help raise money. So I helped. I got Andy involved and some other people. We raised the initial seed round. I recognised the value in this because what we forget from that time period was that the web was growing at an unbelievable rate.

As I [list] in the title this wasn't the first search engine but all the other search engines had been compromised by becoming internet portals and also the results were terrible. Google was just clearly head and shoulders above everything else. When you look at search - some people look at it narrowly - but in some sense it's the most human need and desire that we have beyond the animal level that our curiosity, our need for information and so on is all about searching for it.

People say, well what was the business model? Well David Sarnoff and the radio pointed the way. Look at how radio developed. You either have to pay for it or you pay for it with advertising. So it seemed like it was obvious that you could go one way or the other. Even then the subscriber model wasn't really kicking on. It seemed like a great way to build it from an advertising standpoint. So again - not the first mover but the most successful one.

VMware is another thing where - I was an early investor in VMware. Again I sort of label this as virtual machines rise again because this wasn't the first instance of somebody developing virtual machines. I was involved in hiring Mendel Rosenblum, one of the founders of VMware at Stanford. Sometimes I pick good people. Mendel was a good pick for a faculty member.

I got to know his wife [Dianne] because both of us wind-surfed. She's better than I am. It was kind of curious. I'm glad that I stayed friends with Mendel because he and I - I'd take off with his wife off in the San Francisco Bay and he was kind of left behind. At any rate we got to know each other.

IBM developed virtual machines in the 1970s. This was to allow two riders on the same dinosaur, that is, two operating systems running on the same big dinosaur or 370 machines. So that was that and didn't go very far.

When VMware was founded one of the things that I realised in [tying] to people in the industry was the practice had involved that you ran Windows as a single application operating system because what people found is if you put two applications on Windows they usually get in each other's way and is far less reliable.

So basically what Virtual Machines became is the ability to run a whole bunch of Window instances on the same server without relying on the operating system to keep the applications out of each other's hair or destabilising Windows.

Now in the research context there was this notion of developing what was called a microkernel - a much smaller, very effective core piece of the operating system to improve reliability. This was what I considered the Holy Grail. I had worked on this for years.

Mendel comes along and reinvigorates the whole virtual machine thing. I suddenly realised well I was missing the right approach to this microkernel approach which is to make it a true virtual machine compatible at the level you could run Windows. Of course the rest is history. So this is putting together my research, my friends and realising the technology was proven but just hadn't been directed at this particular need.

Another thing I got involved with was Arista which was yet another switching company. You know this is yet another - I'm trying to emphasise that it's not necessarily the first

company that does something that wins here. I had a PhD student, [Depang Zhou], had worked on some research with me which he thought was going to be commercialisable and so on. He wanted to start a company. I debated whether I wanted to do yet another start-up here. But I agreed to help him and provide some funding because it was kind of an interesting software challenge.

I went back to Andy like I'd done with Google. I guess he thought that when I asked him about Google that worked out pretty well. So he agreed to put in some funds with this start-up. Then a couple of my former students who are incredible software people, Ken Duda and Hugh Holbrook had agreed to join. Mansour Karam joined us as well as somebody I knew from Stanford and so on. We had this great team charging along.

The trouble was we gradually realised we had a great team but the wrong product. There was no market for this. I often think that this points the way that there's a lot of opportunities come along, sometimes it's harder to put together the team. It's important to not get too religious about what you thought you were going to build.

We did a pivot on this. I think a lot of the pivot was really orchestrated by Andy Bechtolsheim to recognise that 10 gigabit Ethernet was really happening, was really happening in the big data centres that as the cloud - cloud computing was taking off and there was a need for a next generation kind of switch.

So the David - our smallest networking company in the world - took on Cisco, the Goliath, in the cloud data centre. While they were off buying things like flip cameras and so on we were aimed at the area that was growing the most rapidly with the highest margins and also benefited from the fact that Cisco had put a very high premium on 10 gigabit Ethernet to protect the lower speed products in their line whereas we didn't have any lower speed products. So they kind of created what Andy likes to call a price umbrella over us so we could come in here and they had difficulty competing.

When I got out of Arista I was planning to go back to academia. This Apstra opportunity came along. It was essentially yet another network automation company, because people had tried to do this sort of network automation before. One of the things I find fascinating about networking is it's down in the plumbing. It's simple enough in its basic principle to explain to your parents. But it's surprisingly complicated when you actually look at all the bells, whistles, features and all the things that people do plus the complex dynamics that take place there.

So I was trying to go back to academia and then I started talking to Mansour and Sasha about this in tackling the biggest problem I see in networking. [Everyone's] been aware of this problem for many, many years which is how do you run a network as a system as opposed to a collection of boxes?

The typical networking company is really a switch company that Cisco or Arista or any - Foundry - on and on - will sell you a box, another box and another box and sell you a bunch of cables. Then you have to put it all together. So it's like buying a Lego kit except there's no instructions on how to put it together. You have your own preference as to what you want to build.

Or it's like you want to have a new home and what shows up is a bunch of bricks and a bunch of timber and you have to put it all together yourself. I was not planning to do another start-up but I realised that the time was right. The team was right. The technology was right and ripe for that matter, to do this. It's actually a pretty high technological bar as well. So there's a lot of value - well and there is a lot of value in doing this.

So the biggest concern that people raised when I mentioned I was planning to do this was this view, well everything is moving to the public cloud so there won't be any market. Nobody's going to manage their own network. They're going to hand it all over to our friend Jeff Bezos.

My view here is that we've ended up in a bit of a - what I'll call medieval computing landscape - because you have Jeff Bezos who's built this walled garden if you like called AWS which is a castle. He basically says, come under my protection and you will be safe. Ha ha ha.

Exactly what happened in medieval days, you think the king was - felt that the highwaymen were a problem. They weren't a problem. They were part of his overall plan. You build this castle, the peasants are beholden to you because you protect them from the highwaymen. So we're ending up in the same thing.

Now AWS if I can [be mean] to this, is sold as the public cloud. Well what does public mean? It means that if you pay for it you can use it. Well, by that definition I claim Disneyland as a public park because you can pay and you can go in and use it. So public cloud is a complete misnomer.

Then AWS says, well that's not the only piece. We don't just provide computing service. We provide thousands of services. One of the AWS invents - they mention some ridiculous number of new services they're adding every day. Well what people are learning is those services are there for lock-in. They're not all that useful.

There's this disease you get under junior software engineers that say, gee I never use them like that. Why don't I try using it with this application that somebody is paying me to develop? You find these start-ups and big companies end up where they develop some application and it's using 100 of Jeff Bezos' little services that are unique to AWS.

So somebody gets a big bill from Jeff Bezos. Some of the bills I've seen are enormous, unbelievable, \$1 million a month easily and upwards, It's very hard to move that application because it's not highly dependent on that service. It's just that highly - just slightly dependent means it doesn't work without it. So you can't get off of AWS without re-writing the application.

Then of course you're sold this notion that you're not in the IT infrastructure business. I view that as a myth because Bezos basically says you don't need your own brain. You can do your thinking with my brain. Ha ha ha. Well you look at Andreessen's comment about software is eating whole industries well Jeff Bezos is eating whole industries.

Remember people used to say well it's just an online bookstore? Well what's happened to the grocery business? What's happened to a whole - the video distribution, music distribution and so on? This guy has enormous clout. If he decided to start a bank

tomorrow it would instantaneously be one of the largest banks in the world. There's 100 million prime customers at this point with Amazon. So this is paying customers.

So I think there's a real risk here that this guy is not the richest man in the known universe for no reason. I'm not faulting him but I think that anybody who wants to have an independent business should be aware that he does want to own the entire universe here. He's got incredible capabilities to move into almost any industry.

In fact I was told that Boeing is moving off of AWS because they regard Amazon as a competitor. You think, what? Boeing - Amazon? Well Boeing has a drone business. That's one of their fastest growing areas. What is Amazon doing? What is Amazon good at? What has Amazon invested in, in delivery? So it's now competing with Boeing. So name a business and I think they should be concerned about this.

I view that there is an alternative which is the private cloud. You can get the same efficiencies as Jeff Bezos more or less, buying your own hardware except for how you operate it. So what separates the efficiency of your operations from AWS is purely automation. Part of my enthusiasm is providing this automation that people have developed - a lot of smart PhDs have developed - in places like Amazon and Google - I know a lot of these people, Facebook and so on - and providing it to other companies.

I think that we're in this very exciting time of what I'll call pervasive automation which you might consider the third wave of computing where the first wave was human productivity. You say well remember the productivity line of Office Suite and so on? You take people doing stuff. You're going to write a letter. How do you make it more efficient than having to redraft it when there's a spelling mistake? Well use Word and so on.

Then the second wave was human connectivity. Remember when Ma Bell was the only way to communicate in real time? Now the internet has created an enormous number of alternatives. I was going to call the third wave human automation. But somehow it didn't sound quite right. But I think the third wave is this automation of things we normally associate with human beings having to do.

So when you look at Uber and Waymo and so on, they're automating transportation. It's not just those companies. It goes beyond that. You have Caterpillar tractor developing self-driving bulldozers. Airbnb is automating hospitality. Katerra is another company I'm involved in which is automating big parts of construction and supply chain. Caspar is another one that's automating smart apartments.

Of course Apstra which I'm heavily involved in, is automating the data centre management. So I think we're seeing this whole trend to automating a whole bunch of things. Automation - none of this is easy but it's incredibly valuable to do this.

I tried to answer the question, am I smart, lucky or prepared. Well, I think lucky - I've been very fortunate to end up at Stanford which was a series of curious incidents given I was born in Canada and lived there for some time until I was actually an assistant professor for a while in Canada. So it was great to end up at Stanford at the time I did and then of course be able to meet and work with a lot of amazing students and other people that came through Stanford or came through Silicon Valley.

In terms of prepared, there's this great line of Louis Pasteur which is, fortune favours the prepared mind. I guess my research has been a great preparation for the things I've been involved with because I always like to focus on real world problems. My dad grew up on a farm in the dirty thirties. My grandfather was a homesteader.

I kind of inherited some real practical gene in spite of becoming an academic. So I've been interested in how you build complex software systems, how you build these in a distributed way and of course networking which is underlying the whole distributed direction and what many people call the cloud now.

Smart, well it's hard to feel smart in Silicon Valley because there's so many other smart people. But I think I've been smart enough to treat students well, to recognise the ones that are smarter than me and start companies with these folks.

I don't know - takeaways - well if you want to invest like me, you get a position at Stanford 35 years ago and run with it I guess, a little hard. But I think we live in incredibly transformative times which is great because there are enormous opportunities. But like I say, so was the American Revolution. You want to be on the right side of history because the wrong side is fairly painful.

I think - this is perhaps a little bit of personal philosophy but I think - I never invest in something that doesn't add real value that I can identify. I've missed out on a lot of great investments because of this thinking. So take it with a grain of salt.

Never invest in a start-up that only exists for one reason because I've seen so many cases where the one reason goes away. So you want multiple reasons. To me, three is the first large number. So there should be three good reasons.

Never invest in anything you can't explain to your mother which may be duplicating the add real value thing. People that come to me with ideas for a companies and I can't figure out how you'd ever explain it to a non-technical person, of course that's who you're going to sell to in many cases.

Also there's this whole notion of first mover advantage. You've got to come up with something totally new and so on. In fact what I've seen is a lot of the first movers are often the first losers. You can go through the list. Arista was not the first switch company. Google is not the first search engine company. In fact I had a student involved in a previous search company much earlier and things.

So there are a lot of lessons you can learn from the people that go out there and fight on the front lines the first time around here. They often prepare the market and so on. So I think a lot of success really ties back to - I often say the success of the company is similar to the French Revolution. It's all about execution. So the important thing is how well you execute. Some of that goes both ways. What you kill off and what you keep and how you go about doing it.

Finally it's important to be lucky and be prepared and obviously choose your friends well since I've benefited enormously from the friends that I've developed over the years in Silicon Valley.

So let me stop there.

Manek Dubash

Thank you very much. A fascinating presentation. Thank you so much for that. If Jerry Caron would like to come down - Jerry from GlobalData and - oh no - maybe throw you a ball or two and see what happens there. We'll have questions afterwards.

*Keynote Interview & Audience Q&A
with Professor David Cheriton and Jeremiah Caron,
Global Head of Research & Analysis - Technology Group,
GlobalData*

Jeremiah Caron

Hello everybody. I'm Jerry Caron, I'm the Head of Research and Analysis for Global Data. Very happy to be here to ask David a few questions. I had the pleasure of a few minutes this morning over breakfast to talk about some of these things and what I want to do is just ask a couple of questions that sort of follow on what you just spoke about and then open it up for questions from the folks out here as well, so go through some things.

The first thing I want to start with is that you articulated at a couple of points during your presentation sort of a clear vision about what you invest in, why you invest in it and everything like that. So, if we could dig into that a little bit more about what it is - what are the key attributes you look for in an investment. Maybe to frame that, use an example from the early days, say Granite. Because Granite and that Ethernet switching era, there were a lot of people doing Ethernet switching and they were not the first. What about that set up made a difference for you and drove you to kind of go in with that at that time?

David Cheriton

Well, let's see. I think - let me start off with, Andy Bechtolsheim was I think one of the brightest people I've ever met in my life in terms of not just technically but astute in terms of developing something that the market needs. I feel like the opportunity to work with him and he's a delightful person too, so the opportunity to work with him was incredibly seductive. I think the other thing that, this is going to sound a little arrogant, but I think one of the things that was under recognised, which I think Andy tuned into, is that networking had become, even at the switch level and switch chip level which is what Granite was focused on, become very complicated.

Most of the people developing hardware for Gigabit Ethernet at that time were all hardware engineers so they were dealing with - they were taking their traditional approaches to developing hardware and just trying to scale it up, whereas I was coming from the software world where we dealt with much higher degrees of complexity than what hardware had dealt with. I was simply scaling it down so I felt that there was an opportunity to do a much better structured approach to hardware design applied to complex chip like an Ethernet switch and putting Andy's brilliance together with my networking and software background I felt was an opportunity to do it great.

I think it paid off because I just remember there was an interop where we were part of CISCO at the time, but we showed up at CISCO with a single chip Ethernet switch operating a Gigabit and there was all these companies that were the also-rans that had implemented this across five or six different ASICs. Five or six means you've got a cost that's five or six times higher than CISCO. I think they basically just packed up their bags and went home when they saw one chip doing what their five or six took to do.

Jeremiah Caron

Yeah. I think that one of the key lessons there and it seems to be the case throughout your key investments, throughout your career, is that even though there are other organisations perhaps doing something that at a high level looks to be the same thing, really there is - if you can innovate and you can do it in a better way the market will be there for you. I think bring it to present day times, the SD-WAN market has been talked about and it's been driving a lot of interest for now quite a few years. While a lot of bets have already been made in terms of in the managed WAN services market on what suppliers will be using, the game isn't over right. It's still early days in terms of what the market might need and ways to potentially innovate and help drive that market.

David Cheriton

Yeah, I think it's still early days but I [must admit] I think, you know, my perception is service providers in the wide area, there's more innovation that's been driven from the datacentre outwards than for evolving the WAN itself. To me, if you think about it from a sport standpoint, I think the datacentre is like playing high speed soccer and the WAN area seems like it's more like lawn bowling, so it's difficult for start ups to stay alive while it's evolving in that sense. I think it's evolving but it's a long, long play in my mind.

Jeremiah Caron

Right, okay. Moving to a more sort of generalised question. Oh, just back to your philosophy about investment. You made a point of saying you're not a VC. Why do you think more VCs don't sort of take the approach that you take? Is there - you would think the logic is there that hey, rather than taking the shotgun approach, being more targeted, building it around a belief in the value and then investing and making it work as opposed to just casually [absorbing] seeing which ones work and might be the way to go. Do you have any thoughts on that?

David Cheriton

Well I think that my approach is complementary if I can put that into the way the VCs operate. I think that the VCs have to operate in the way they do because my understanding is you're part of a VC firm, a whole bunch of people want to invest in your fund because of your success, so you end up with a billion dollar fund and your job is to invest it. You only have 10 partners or something or rather so you end up viewing, well, the 10 partners can only be on so many boards or so many companies and so I have to put in a lot of money.

So I think that you end up having to, to get the return on investment, you have to figure out a way to invest it across a large number of companies and so I think you're forced to operate in this mode. Plus I think that most venture capitalists have a lot more business background than I do and a lot less technical background so they're looking at these companies from the standpoint of the business dynamics which are not clear very early on. I think there's things that are ready to launch in a significant way where the VCs can be very useful and help and then there's things that take a longer period and require a lot more technical input. I like to think those are better ones for me to tackle.

Jeremiah Caron

Yeah, okay. One of the things we talked about at breakfast, the sort of global nature of innovation and driving new technology development and thinking obviously a lot of it sort of ends up being funnelled back to, or through, this area where we're sitting now, in Silicon Valley. Are you, in your work and what you look at, is it global or do you look at things going on in other countries or do you tend to focus on what's happening in this particular area of the world?

David Cheriton

Well I mean I think the trends that are happening are global. One of the amazing aspects of Silicon Valley is that people come here from all over the world either to tell me what they're doing in other parts of the world or else to do their start ups here or move their start ups here. I'm a pretty laid back casual guy. I like to ride my bicycle to work and things and I don't like to travel particularly, so, so much comes to me that I feel like I'm paying attention to the rest of the world just by sitting here and taking meetings with people.

Jeremiah Caron

Yeah. I mean there's lots of areas of the world who have tried to replicate Silicon Valley and it just hasn't happened. It is actually global what happens here. People come with their ideas and end up here because of the leveraging of everybody else really. Another area I wanted to ask about, we talked about, was this, you know, when you talk about security, talk about automation, talk about AI, a number of other things, a point that often comes up is this notion of a skills shortage. That there isn't enough technical competence around the world to make these things happen as quickly as we might like

or in the best way as we might like. You work at an organisation and you're colleagues others, like California Institute of Technology or Massachusetts Institute of technology. You have the benefit of the best and the brightest coming through your organisations, but what is your view on this perceived skill shortage? Is a problem or not?

David Cheriton

Well I think there is a huge skills shortage. I think and certainly we see this in the salaries that software engineers are commanding these days. I think that this is also a bit of a moving target because I think there is this wave of automation taking place so I think talented young people are smart enough to figure out where they want to direct their energies. So they don't want to do things that are feasible to automate and that covers a lot of territory. I think that we're kind of working through this problem but I think that there's almost a bigger problem that we have legacy systems that need to be automated and there's simply - I think there's an automation shortage in some sense right now. I think we're moving in that direction but there's just a lot of work to do to get there.

Jeremiah Caron

Right, okay. I want to stop for a second and check and see what questions we have from everybody assembled here. There must be a couple, one or two. We have a hand shoot up right in the centre so why don't we go with that one first. But the microphone is over there, my apologies, so we'll start over here on the right.

[Unidentified Male Participant]

Yeah, I got the mic and it's a short question. There's lots of press in here and one of the odd things about the press sector is ideas turn up all the time for new titles, new sites, new everything else. There's an analogy with the investment process you were talking about...

David Cheriton

Can you hold the mic a little closer? I'm having...

[Unidentified Male Participant]

Closer or further? So we spend our time doing triage on new titles for websites and even articles and so on and I'm interested in hearing about your kind of ratio of start-ups you pick, because your story was picked a start-up, picked a start-up. Was that out of thousands of proposals, hundreds of proposals? How do you triage your decision on what to be involved in?

David Cheriton

Well I think it's, you know, if you're asking how many sort of proposals I look at versus how many I've invested in, I think it's tens that I've looked at not hundreds. Maybe it adds up to a few hundred. It ranges. There's ones where I've invested in that I haven't spent as much time. One of my former students has founded a company called AISense

that's doing some interesting stuff. I think they're in the Shark Tank here and I invested with him before and he's very independent and was successful before and so when he came to me with [a lobbed] crazy idea I thought if anybody can pull this off he can and I just invested some money along with people in that and let him do his thing. There's other cases where there are students that come to me with something that's really half baked where they're just looking for input. If I can see some hope for it I suggest a hope and if I can't see any hope I'm fairly direct about that too.

One of the things that I experienced early on when I was a professor, assistant professor, I got involved in various consulting operations where I'd end up consulting with some company and you encounter these start-ups that are really on their last legs. You see engineers have put heart and soul into something and just hit a brick wall. It's really an upsetting experience. I mean the money has gone but I think people are really psychologically traumatised by buying into a vision and working so hard on it and it not going any place. So it makes me a little more conservative to say before you get everybody going and charged up and committing part of their career to something look at it very carefully.

Jeremiah Caron

We would be remiss if we didn't dive in to talk about AI a little bit because it's such a big concept that everyone writes about and everyone is sort of talking about. You're involved in a number of AI initiatives, correct, so that's something you're fairly passionate about and have belief in, at least at the applications of AI.

David Cheriton

Well yeah, I think - I'm glad you added applications of AI. I have the advantage, if you like, of working with AI researchers at Stanford, one of the leading AI places in the planet for the last 35 years, so I'm familiar with the technology and I'm familiar with how its evolved. Some of the early days of AI was defined, in fact that's where the term came from, was to try and replicate human intelligence. If you would like, there's seven billion of us on the planet, why do we need a machine doing the same thing? This is kind of like saying, if the Wright brothers said what we want to do is come up with something like a bird that can land in a tree and lay an egg and you think well that's not what's needed, what's needed is transportation.

I think one of the things that's changed about AI is that we've gone from people want to replicate human intelligence to people want to apply those techniques to real world problems. Again, I'm Canadian, I'm practical, I tend to say when is the technology ripe, when are all the pieces put together and where is it applicable? For instance, AISense is doing an automatic speech recognition and harvesting that. It wouldn't be practical 20 years ago, it wouldn't be practical 10 years ago. It's now practical, deployable and I think - I actually was partly funding the autonomous vehicle work at Standard when it started and that became taken over by Google which became part of Waymo.

There, again, was some smart people where I talked to them, the technology was getting right to do that and it was moving ahead. I think that when you get beyond the label of AI and you say, here are the techniques being used and are they right to deploy. You

look at machine learning, for instance, it's gotten a lot of play, it's very expensive to do and deep learning is incredibly expensive to do. But then when you have our friends at Google deploying things that have more computing horsepower than the entire universe had five years ago, which is insane, I mean the numbers are just unbelievable what they're deploying, you realise that there's something you can do which you can't even imagine having been done some years ago.

I think this is where it's not just saying, well AI is this magic thing, but looking at the specific technologies and then how you can apply them. I think one of the other struggles that I think I have and everybody should have is that we cannot, as human beings, understand what is going on here in sort of ordinary day senses. You have computers, like with machine learning, looking at billions of variables. Remember you took algebra and there was three or four. Well maybe you can imagine 100 and then 1000 might be a whole page of the, but we cannot conceive of the size of the computations we're doing now. We cannot begin to understand what the power of this is until we apply it. Now we have examples where people are applying things that were impractical five or 10 years ago to problems, applying them to those problems and the results are just astounding.

I think that the potential is enormous if you take the right techniques at the right time and apply it to the right problems. If you do any of those pieces wrong you're going to end up with something that is just not quite there.

Jeremiah Caron

Okay, great. A question here.

Dan Pitt

Hi. Dan Pitt, I am with MEF. I am based here and David it's always interesting to hear your comments. I want to focus on the comments you made about the use of public clouds versus a private cloud or retaining your own infrastructure. Coincidentally last night I attended a program hosted by AWS showcasing all of their new applications and Amazon applications that use some aspect of AI. We talked about the relative cost of maintaining your own infrastructure and as you point out with automation, it's a lot lower than it used to be or it could be a lot lower than it used to be. But there are other aspects of maintaining an infrastructure versus putting it into a large public cloud that I'd like to get your views on.

There's not just AWS. There's also Google Cloud and there's Microsoft Azure as viable competitors and in other places things like AliCloud. The question came up in the discussion at dinner last night and AWS paid for very nice food, about the role of security and can a private enterprise afford the latest, the best resources for security of their own network, of their own assets, as a very large cloud provider that has deeper pockets. What are your thoughts on that David?

David Cheriton

So the question is, can I have a secure private cloud that's comparable or better than Jeff Bezos' cloud. Is that right?

Dan Pitt

Yeah, or Google Cloud or [unclear].

David Cheriton

Sure. Yeah, I think you might - let me divert it slightly initially. You might make the same argument about reliability and say well Jeff Bezos has got all these guys who are experts in running the cloud, how are you going to maintain the same reliability? Well there's a date that sticks in my mind which is 14 March 2017. On that date 150,000 companies disappeared off the internet. The reason was some guy in a datacentre in Virginia made a little error in configuring one of Jeff Bezos' networks. The reason he made that error was they were changing the configuration so they could make the datacentre bigger. The reason they wanted to make the datacentre bigger was because Jeff Bezos was talking on more of the universe. You have 150,000 disappear, not because they wanted to get bigger, not because they wanted something better, but because Jeff Bezos wanted to grow his empire bigger.

I think that that's an example where you can get hosed by his ambitions where this has nothing to do with what you want to accomplish. When you come to security there's two forms of security. I think when I talk to people in datacentres, a lot of their concern is not infiltration it's exfiltration. It's we have the barriers around the datacentre fairly secure and there's a lot of good firewall technology and so on and monitoring mechanisms there that you can get right now. I think the problem is getting things that get inside, where everything does get inside and then you have things flowing out.

Now, in that world you'll say, what do you view as safer? Your insider cloud with thousands of other applications, some of them being run by your competitors, some of them being run by who knows what, whoever Jeff Bezos decides to sell time to and you have very low visibility into what's going on and what the separation is. Plus we've seen recently these nasty things like side channel attacks which are incredibly hard to prevent and I think we've only begun to see the beginnings of those. So, to me, this is like if you're concerned about getting a disease are you going to share the same bathroom with all sorts of people? Jeff Bezos is just saying, I clean it very carefully or do you want your own bathroom? I think those are - I shouldn't put it in those terms, sorry [Murray], but I think it comes down to that sort of comparison.

In theory you could say, well, they've got all these experts, but in practice I think looking at the reliability of AWS, they've had some significant outages that are nothing to do with what the companies are doing. I'm not sure that - I can't point at any particular security problem but I think it's unclear that their emphasis is on growth and not on looking after the customers. I think - whereas if it's your datacentre you can protect it and I think there's a lot of services that enable you to do this with your own private cloud?

[Unidentified Male Participant]

Would you say the same for Google Cloud though as you say for AWS?

David Cheriton

Well I think it's the same thing. Now I've used the analogy of your private home and a hotel. If you say I've got a family and we live comfortably in a three bedroom house I'd buy a three bedroom house. The relatives come to visit at Christmas time and I need five more bedrooms, I rent hotel space. To me, the public cloud is great from that standpoint to say if I'm not sure where I want to live in a new city I rent a hotel for a while and figure it out, or if I buy my three bedroom house and then the relatives come to visit I'll rent hotel space. It's basically great to have the public cloud there for research space, for start-ups who want to experiment initially and when you're trying to figure out what you're doing.

I think once you figure out what you're doing and get to a certain scale you want to have your own private cloud, just like you want to have your own private house or private condo. I think that's the thinking so there's a role whether it's Azure, whether it's Google Cloud, whether it's AWS. The reason I focus on AWS is they're by far the largest, the most dominant and they've got the greatest track record of getting people on board and then taking over their business.

Jeremiah Caron

Alright. I think we're going to wrap it up now. I really want to thank you for sharing your experience and insights with the group. Thank you very much.

David Cheriton

Thanks.

[Applause]