Celebrating 40 years of Ethernet! – from 2.94 Mbps to 800 Gbps …

Interview transcript: Bob Metcalfe, Inventor of Ethernet & UT Austin Professor of Innovation & Entrepreneurship with Jensen Huang, Founder, President and CEO, NVIDIA

Featured Speakers:

Bob Metcalfe, Co-Inventor of Ethernet & UT Austin Professor of Innovation & Entrepreneurship
Jensen Huang, Founder, President & CEO, NVIDIA

Jensen Huang
Oh, this is fantastic. Well, Bob, it is such an incredible pleasure to be able to spend time with you today, on the 40th anniversary of the first Ethernet specification, you're the father of the Internet of the Ethernet, you're the inventor of Ethernet. Of course, as always, we have lots and lots of collaborators. But your contribution to the foundation of the internet today and its impact can't possibly be overstated. You've had an amazing career. You were born in 1946, your career really started then you were born in 1946, which is also the birth date or the birth year of the transistor, the two of you enabled the internet that we know today. And now you went on to an incredibly celebrated career, you're a recipient of awards that that are countless: the National Medal of Technology, the IEEE Medal of Honour, the Internet Hall of Fame, the Alexander Graham Bell medal, the Grace Murray Hopper award, the Computer Science Award and these can't possibly recognize fully the contribution you've made to the world and to society today. You know very well, that the invention of the Ethernet is what enabled the internet today, where billions of people are connected to it, trillions of dollars of the economy is enabled by it. And because of the internet, it is possible for us to be in the same, although not in the same physical space, to be able to enjoy each other's presence in the same digital space. And because of the internet, it's possible for companies to become global companies to be virtual giant companies, with sites in Silicon Valley, in China and India and your European sites, all over the world to form today's great companies. When you go back and think about the beginning of it what inspired you? Of course, everybody knows what Ethernet is, everybody knows the essential technology event, the amazing breakthroughs enabled since. But the things that I want to talk to you today are really the inspirations behind it, what was the what was the inspiration behind the invention of Ethernet? What was the

problem you were trying to solve? And what were the surrounding conditions, the surrounding environments, at PARC, that inspired you to invent Ethernet?

Bob Metcalfe
Well, thank you for those many kind words, Jensen, I'm honoured to be with you here today. I got into the networking business in 1969, because I showed up at Harvard as a grad student. And the algorithm that grad students follow is they find out where the money is, and then they decide to major in that. So the Advanced Research Projects Agency of the Department of Defense was giving out money to computer science departments. I detected that and so I became an ARPANET grad student, and at MIT as a student at Harvard. But as an employee of MIT, I connected MIT to the ARPANET, the then internet. And I learned in the course of that how to send bits, one at a time down a very long wire. then moved to Xerox Corporation, Palo Alto, Xerox PARC, where I put PARC on the ARPANET. I just did the same thing. They built the same hardware.

Jensen Huang
Now when you put PARC on the ARPANET, it must have been one computer. When you say PARC I mean was it one computer at PARC that you put on the ARPANET or was it was a part that you put on the ARPANET?

Bob Metcalfe
It was one computer name was MAXC. It was a PDP 10 clone that we built in that lab. And I was the networking guy in the computer science lab there so I got the job of connecting it to the ARPANET. It was basically building a printed circuit card about this thing that would connect the time sharing system, the PDP 10 running the 10 x time sharing system connecting it to the packet switch the ARPANET packets which were in the same room. So I once again am sending bits one at a time down a long wire.

Jensen Huang
So you were a system designer, as well as the network administrator, so to speak?

Bob Metcalfe
I operated at the boundary of hardware and software. So I would build hardware. I would write software at the vendors, but just the plumbing level and then my whole career has been at the plumbing level.

Jensen Huang
Well, you're the world's most famous plumber, there is no doubt you and Mario both! Yet now this this printed circuit board that you made is made out of what kind of technology at the time is TTL and OSI type technology?

Bob Metcalfe

Well, actually, I just lied. I didn't make a printed circuit card, I made a wire wrapped of a big wire wrapped card and using the Texas Instruments 7400 series MSI chips. Remember that was a chipwhat was it, the 7404 had two flip flops on it. And the 7400 had six inverters. And anyway, getting the chips on the card was part of the limiting factor. So let me move on to the Ethernet because after we put Xerox on the ARPANET, we then decided to put a personal computer on every desk. And it and I lucked out, they gave me the job of connecting them together.

Jensen Huang
Bob, at this time the PC hasn't been invented yet. You were at PARC. And so with the personal computer, was it called a personal computer, it was probably alto that you're talking about, was the alto we did call it a personal computer.

Bob Metcalfe
No, by the way you can get into a big argument about which was the first personal computer. But the alto at Xerox research is my version of the first personal computer and I lucked out. Yeah, I got to build a network. And one of the requirements was that we were also building the world's first laser printer, page per second, 500 dots per inch. If you do the math 500 x 500 x 8½ x 11 per second. It's about 20 megabits per second. So the network had to be fast. We couldn't use RS 232, which was the then standard. So in my office, the day before Ethernet, I had a 300 baud, dumb terminal. And then when we installed Ethernet in my office, I had a 2.94 megabit per second packet switch land in my office megabit per second, whopping numbers. But just to save you the arithmetic the increase was a factor of 10,000. Wow. Well, it took us another 40 years to do much better than that, but not much better than that.

Jensen Huang
You know, in so I know we're doing a huge leap. There is a group doing 800 gigabits per second Ethernet right now. Yeah, right. And so that's 100,000 times about right.

Bob Metcalfe
47,000 times, but who's counting?

Jensen Huang
Yeah, well, good. And so my first experience, though, with Ethernet was it was a local area network that was it was called a LAN. And it was really about connecting computers that were close by. And my first use of it was really connecting peripherals. Now that the, the from that point in time from that point in time, to the time we are today, where Ethernet is clearly not just about peripherals, it's not just about local area networks, it's not even about coaxial twisted pairs in and it's so much more than that. The first step of that the first step of that, of course, was the standardization of the technology that you invented and that standardization process, what was what was happening at the time what inspired you to go to IEEE and made it eventually the 802.3.

Bob Metcalfe

So I left Xerox in January of 1979 with the intent of starting a company, although I had no clue what it would be. And in a meeting in February of that year with Gordon Bell, who was then the Vice President of R&D of Digital Equipment Corporation (DEC), the number two computer company in the world.

Jensen Huang
Yeah, I know Gordon, I know Gordon very well.

Bob Metcalfe
Gordon asked me to design an Ethernet for DEC. I said I couldn't do that, Gordon. First of all, I feel loyal to Xerox and second of all, I've already designed the best network I know how and yours would be second best. So together we had the idea of writing a letter to Xerox proposing that DEC and Xerox would work together to get their products to be compatible over Ethernet. And then I ran into Intel Corporation down at the National Bureau of Standards looking for a new chip standard implement with some new NMOS thing they had. So pretty soon we had DEC, Intel and Xerox co-operating on connecting with Ethernet, and we had it and suddenly they had a problem they couldn't meet, because in order to meet, they would have to violate anti-trust law. So I got my lawyer, fraternity brother I might add, Howard Charney, and he gave me advice as to what I should tell the lawyers that those three companies that would make it okay for them to meet. I gave them that list of five - what is no marketing people allowed in the meeting note pricing discussions, and the goal had to be an industry standard. So that was the moment in which the idea of creating an industry standard Ethernet came up. And you have to remember those days, standards were not made, like they are today. IBM would just announce its next product and that would be hit standard right away. So the idea being Ethernet to the IEEE was an important breakthrough idea.

Jensen Huang
Now Bob at this time, at this time in your in your life. How old were you? When you started 3Com, how old were you?

Bob Metcalfe
I was 33.

Jensen Huang
33, incredible. Well, at that time, there weren't that many technology startups, quite frankly. And surely at 33 that's quite young. And who did you work with at Intel? You know, for gossip sake, who was CEO at the time?

Bob Metcalfe
Andy Grove was the principal mover but a guy named Kauffman, forgotten his name... He was Groves. right hand man. And he came to a conference pitch I gave on Ethernet and then he invited me back to Intel to give a presentation at Intel and I met Mr. Grove and they were instrumental in getting Ethernet standardized.

Jensen Huang
That's really quite amazing. Now, at the time, there were a couple of other competing standards that were being proposed: Token Ring and Token Bus was being floated around, How did how did you guys decide, how did the community decide, that the standard would become Ethernet, you know, were there technical merits, were there economic merits ,were industrial support merits?

Bob Metcalfe
So on September 30 1980, the event we're celebrating today, DEC Intel and Xerox produced a spec, the Ethernet Blue Book, it was called, and gave it to the newly forming committee IEEE 802 with no dots just 802 and we thought that project 802 would then standardize officially Ethernet. But IBM and General Motors had two different ideas. They didn't like the idea of somebody else making standards for them. So they each submitted a competing proposal IBM's Token Ring and General Motors Token Bus, and the IEEE then began to consider the three of them. And eventually I didn't like this outcome, but eventually standardized all three. And that's where the dots come from 802.3 was Ethernet. 802.4 was the Token Bus and 802.5 was the IBM Token Ring. And the way we proceeded from there was fierce competition for the next 20 years.

Jensen Huang
Now, if I were a startup at that time, and it went through the standards bodies, and IBM's was 802.5. I might be inspired to go build that frankly. What were the conditions that cause caused Ethernet to have gained so much industry support? I know that one of the things you always believed in was interoperability and compatibility. And of course, for any platform for any standard. That is the essential core value that allows industry to be formed around it. But at that point, IBM I put out Token Ring to be standardized. And why wouldn't anybody want to build peripherals or computers that would be connected to the world's leading computing company at that time? You know, what caused..

Bob Metcalfe
IBM wasn't used to fighting standards efforts like Ethernet - their heart really wasn't in it. Their dark black heart was intent on pretending to have a standard but shipping a product that was not standard. The way I know this is at 3Com, my company, we shipped IBM Token Ring ahead of IBM and we had a hell of a time selling it. We never sold very much of it because IBM, whose heart was not in interoperability, had software dust they sprinkled over their Token Rings. And so whenever we tried to sell a Token Ring into a customer, we couldn't get it to work. We had this software, Dust, they called it SNA system network architecture. Just to annoy them. I used to call it SNA.

Jensen Huang
Well, if you added technology to it, it would be called SNAT.

Bob Metcalfe
Like, I hadn't thought of that. But going back to your question, the reason that people went to Ethernet eventually, is that number one, Ethernet was designed in the context of the seven layer, a reference model of the internet. And so we knew our place. And so the Ethernet was designed to be level one

and level two and not the other levels. And IBM didn't understand this and nor did General Motors. So for example, Ethernet does not have acknowledgments. You send packets, the packets may or may not be acknowledgments, but Ethernet doesn't know they're just packets. So the whole idea of acknowledging things was we assumed would happen at a higher level with TCP basically. Whereas IBM felt odd because they were a bit old fashioned, who were felt obligated to have acknowledgments. Namely, that token in their ring would go around acknowledging that the packet that you sent has gotten there. And that was built in and level one and two. And that was a mistake. Yeah, so IBM Token Ring was always slower and more expensive than Ethernet. So that was another feature.

Jensen Huang
Yeah, in the wisdom of specifying Ethernet was to go as far up as possible to specify as much as necessary, but not more than that, so that you can enable innovation and evolution on top. Right. Yeah, that was really that was really great wisdom. And so now, 3Com founded. Ethernet was standardized. You're competing against IBM, you know who was 3Coms first customers? Who are the people that that got you off the ground?

Bob Metcalfe
Well, there were a lot of people over a long time got us off the ground, I think the Ethernet bandwagon was joined by principally by Hewlett Packard and DEC, Intel, Xerox and 18, 19 other companies. So it was all of us versus IBM and a few of its close friends. And the hard part is that there weren't any PCs. And the whole purpose of Ethernet was to network PCs. And there weren't any. I mean, there were 100,000, maybe Apple two's floating around. But the apple two was an eight bit 6502 microprocessor. And it wasn't worthy of a 10 megabit per second Ethernet. So we had to wait. And we had to wait for the IBM PC to come out, announced in August of ‘81. And then we built a card that we shipped in Aug, September of ‘82, that you could plug into an IBM Personal Computer and put it on the Ethernet. And then it got difficult. Because people would say, ‘well, what good is Ethernet?’ And I would say it's good for everything. It's great. It's a horizontal platform, it's, you know, it's not a vertical or anything. And it turns out, that is not a good value proposition. So if anybody's planning to say they're selling a platform that's good for anything, customers don't respond. They want to know what it's gonna do for them. Right? Exactly, would you like to know how we sold the first Ethernet?

Jensen Huang
Well, this, I think this is such a great story. And Bob, this is really your second dimension. You know, the first invention that was the technology and your second invention, and arguably even a more insightful invention? Because this leads to my next question. And I'm interested in hearing your thoughts on this. What is Ethernet? Because Ethernet, Ethernet could be many things. It's a technology as a standard, of course it's networking, it’s a market, it’s the fabric of the market, the nature of the market, the people in it, the competition of it. And then is it a platform, and I think in a lot of ways, it's all of that, and your insights. And it's so funny the first time I heard it, your insights, and I'm gonna let you tell it but your insights about you when you were at the time in sales and marketing. You were probably the most educated sales and marketing person in Silicon Valley at the time and because you weren't trying to evangelize a

brand new, a new way of thinking about computing, in fact that intersection between your background and your ability to think about customers and creating markets came together at the time. But it was it was your selling to a customer that I thought was fantastic. Tell this, tell the story of how it came about.

Bob Metcalfe
Ah, well, I have 50 stories that meet that description, but I was just this, the value of a network is proportional to or equal to right? All right, so we're selling Ethernet, to people with personal computers, of whom there aren't any. So we had to make it easy to try our product. So we made a trial kit, a three node kit, three cards, option cards, and a diskette with software on it. And this cost $3,000. And what you could do with it is you could share a printer, among the three PCs, share a disk among the three PCs, send emails among the three PCs, etc. And we sold a bunch of these and it was only $3,000. So people were curious, and they so they bought it. And then they all came back and they said works just like you said it would. We can share the printers, we can share this, blah, blah, blah. And by the way, sending email among three people. Isn't that interesting? Your product isn't useful to us. And I was then, as you pointed out the head of sales and marketing, so that was my problem. So I one night I went over to Stanford and we had, Xerox had given Stanford, some altos. So I snuck in and use the alto which I knew how to use, and I made a slide, a 35 millimeter not a PowerPoint, a 35 millimeter slide that showed that the cost of a network about $1,000 a card was a linear with the number of nodes that you had in your network. But the number of possible connections, which I claimed was the value of the network grew as the square because every node you added could talk to the preceding and minus one node. So and times n minus one is n square, roughly n squared connections. So I argued that the value of a network grew as the square and therefore would eventually overtake the linear at some critical point I made by the way, I made six copies of this slide, because that's how big the salesforce was sick. And the slide basically said, your network is not useful, because it's not big enough. And the remedy for that is to buy more of our products. And they believed us.

Jensen Huang
Now, who is the first customer that believed you. Or, what was your first large installation of number one card?

Bob Metcalfe
Why I remember a young man in Maryland, who worked for a bank, and he was going to build a small workgroup of PCs connected by Ethernet with a multi user accounting package. And he, so he may have been an early customer. By the way, I learned a lesson from him, I learned why people buy products from startups. Why would you take the risk of buying a product from a startup? Why would you do that? And he explained it to me, he says, I want to be the president of this bank. And this workgroup that we're building here, this workgroup, multi user PC network, is going to save us a lot of money and make me look good. And I'm gonna get promoted and move up. So the reason the reason people buy, yeah, the cutting edge technology is because they have competition. And they want to win.

Jensen Huang

Yeah, we save them saving money and saving time, saving money, saving time.

Bob Metcalfe
Yep.

Jensen Huang
Yeah, that's really fantastic. But it's such a great story. And of course, you started the company before the PC revolution really started. And so you must have written a business plan that that excluded the PC revolution. Did you did you see it coming? Or how did you scope out the market, because the market was everything that was left over from IBM, which wasn't much. I mean, DEC was very successful. And, I don't think Sun, Sun wasn't there yet. And, so what was the market you were you were addressing? And how did you…

Bob Metcalfe
Well, I had a business plan. It was spiral bound, it was on brown paper. And it was a writ that was available on its first draft on September 30 1980, the same day as the Ethernet Blue Book, because my plans were predicated on the issuance of the Blue Book, and that the specs would be adequate to build a company and they were. So I then began showing this business plan to people and their first complaint was, there's no market research in your business plan?

And my answer was to this company that what I'm proposing is market research. If there's a market out there, we'll find it. And if there isn't a market, we'll go under.

Jensen Huang
Gosh, I wish I had that line.

Bob Metcalfe
I had that line.

Jensen Huang
It sounds a little bit, that sounds a little bit like if you build it, they will come.

Bob Metcalfe
And I'm glad you came up with that phrase, because in answer to the question, what is Ethernet? One of the principles of Ethernet is build it and they will come. That is when we built that first Ethernet to run a 2.94 megabits per second, we did not have a requirements document from anybody saying they needed 2.94 megabits per second. We just built it as fast as the semiconductors would allow. And that's been a principle of Ethernet. They're building an 800 gigabit per second Ethernet right now. I'm certain there no one requiring it.

Jensen Huang
That's right,

Bob Metcalfe
Build it and they will come has worked for the last 40 or 50 years. So I think, yes, build it, and they will come.

Jensen Huang
Yeah, and I still remember when we started the company, it somebody said that to me, too, is you build it, so what you're saying Jensen, is that if you build it, they will come. And my response to it was if we don't build it they can't come. And so and so all innovators, all integrators that are out there taking a chance, and they believe in something that they have to build it based on dreams and hopes and confidence that there's a there's a world of innovation out there, that so long as the platform could reach millions of people, or millions of developers and smart people, somebody is going to build something amazing at some point. Now, we go back and most people think of Ethernet, if I said Ethernet, most people think of it as the networking technology. Of course, that networking technology is interoperable because it's a standard. The technology as you know well has evolved tremendously since the beginning, since your 7400s, and the MSI's and specification, and Manchester encoding and those the foundational layers of the physical layer, the data layer, the technology itself has evolved tremendously and one of the things that you said earlier, which as far as semiconductor technology could enable it in the case of Ethernet and switching technology. Its engineers are pushing the boundaries beyond semiconductor physics. And in fact, in fact, Shannon's Law was thrown at you in the early days when you were developing and seeing the evolution of Ethernet that had gone beyond Shannon's Law. You know, tell me about that moment. How did you react to it? What was you? Did you did you think that there was a physical limit that that was very close, did that concern you?

Bob Metcalfe
Well, if you thought having IBM and General Motors attacking you wait till AT&T attacks you.
And AT&T attacked us

Jensen Huang
You have law when it's physical information law,

Bob Metcalfe
Yet with Shannon

Jensen Huang
Yeah.

Bob Metcalfe
The Bell Labs, geniuses and many of them were geniuses, said that the copper that had been put in by the telephone company for over 100 years, connecting all those telephones had a Shannon limit of 14.4 kilobits per second.

And so all this talk about megabits per second wasn't going to work over these, this pitiful copper because of Shannon's law. And then one day someone started selling a 50 kilobit per second modem to run over that AT&T copper, and it worked. It worked just fine. I remember going to conferences rubbing AT&T's nose and that I said what about this Shannon limit thing? You guys?

Jensen Huang
Yeah, well, you remember FDDI versus CDDI. That we had reached the limits of copper and we hadn't gotten fibre and fibre was really expensive at the time and so the idea of hundred megabits per second was going to be too difficult to reach. It was too expensive. And that company, the startup company that Cisco ultimately by came up with 100, 100 megabits over copper and unleashed another, you know, several orders of magnitude innovation after that.

Bob Metcalfe
That was Grand Junction networks.

Jensen Huang
That's right, Grand Junction.

Bob Metcalfe
I helped found that company with my partner Howard Charney But you know, I have a theory about your company, which I'd love to test on you right now. My theory is, you know..

Jensen Huang
This is, this is not about me, but this is about you. This is my time, hundreds of thousands of people, millions of engineers are going to watch this, they came to watch you. And so let me let me get something out first. So, so you know, you and I, we've built, we've built companies and we've done a lot of things, we failed in a lot of things and failures are great teachers, and we learned a lot from failures. Failures are great teachers, partly because of the deep emotional anger, the surprise, the disappointment, that forces you to re-evaluate, to be to be introspective into deep analysis. To understand what happened and as engineers will always go back to first principles and recreate the model that predicted the first outcome and refine that model so that we can predict future outcomes. And so failure is a great teacher. And I wouldn't mind hearing lesson two failures. However, the thing that's really extraordinary also, is when something is a spectacular success. It's also an excellent teacher. I found also, that when you when you look at something that is an extraordinary and spectacular success, to be mindful of analyzing. What was what were the dynamics of the events or their circumstances that caused the extraordinary success? You know, some amount of serendipity help some amount of accidents help, but not really. And what were some of the lessons that you took, that that enabled Ethernet to be the magnificent success that has become the fabric that has connected the planet? That all of computing is built on today? What did you take away from that? What were some of the learnings that you would repeat, when you start another company, you weren't, that you would advise the young entrepreneurs that you're advising these things?

Bob Metcalfe
You almost said this, but let me say it, I think you learn more from success than from failure. Because there are so many more ways to fail than there are ways to succeed. So if you find a way to succeed, you have found something of great value. And they were, remember now Ethernet was 1973 invention. So it's gone through a lot and the technologies have evolved, and almost every mistake you can make has been made. Let me try to recall a few. In the building of 3Com, one of the mistakes we made a few times was to focus on our competitor instead of our customer. So we went against Novell that made PC networking software, we spun our wheels for 234 years, fighting Novell, it was a complete waste of time. And then later, the year we went public, a startup started called Cisco. And we, and we didn't give them enough respect. We didn't pay attention to Cisco and it took them a long time, but they eventually blew past us out in the '90s or so. And the mistake we made there was that they introduced routers with 14 protocols and we had routers with three, two or three protocols. And being standards experts we figured that you only needed three but Cisco went out there and offered 14 and guess what the customers wanted 14 protocols because they were filled with uncertainty about the future trying to work

Jensen Huang
They wanted 14 because they're not sure exactly how many they want.

Bob Metcalfe
They wanted to future proof their acquisition so they wanted to be sure that no matter what protocol one they would have it so that's how Cisco slipped by us in the in the early fights over who would sell routers. So that's a lesson. Are you gonna let me propose what I think about your company?

Jensen Huang
(Laughs) Okay fire away.

Bob Metcalfe
We will get back to me.

Jensen Huang
There are so many entrepreneurs on this that would love to hear from you. The lessons that you've learned, you've dedicated the last decade of your career, investing in companies creating frameworks for innovation, for entrepreneurs that think about the starting of their company. I personally want to hear about it, and I know that there are millions of young entrepreneurs who would like to hear about it as well. And so long as we can come back to that, I'll answer any questions you have.

Bob Metcalfe
Well, my theory of NVIDIA is I go beyond the GPU. So clearly NVIDIA is the GPU company. But I don't think that's your secret. Your secret isn't the GPU per se, it's the interconnectivity among the GPUs. That's your secret weapon. So you're in the same business, I was in connecting things together. I was in the

business of connecting PCs and mini computers together, you're in the business of connecting GPUs together. What do you think of that theory of your company?

Jensen Huang
Yeah, the connectivity of small processors that are orchestrated, and scheduled, micro scheduling of a large network of processors, is the technology that is the fundamental (inaudible) playing large processors is quite heavyweight. And it's very difficult to take very, very complex programmes, to have it be distributed across all of these processors, have them be coordinated, to communicate their intermediate results, and then to get everybody going, again, that orchestration, that scheduling is a very complex problem. I mean, it's the supercomputing problem is to large scale, massively parallel problem. And so that the technology is complicated from that perspective. You know, people think that on chip networks must be easier to do, and it is easier to do than then the internet, of course, but the speed at which we communicate with thousands of processors, all able to talk to each other simultaneously, this fabric that we call it inside, inside the chip, our internet, if you will, is enormously complicated. It's a machine in itself. 10s of clock cycles are propagating across the fabric from thousands of processors, all at the same time. And we have to keep the latency down the throughput high. And it has to be it has to be resilient to two very large ranges of functional temperature and operating temperature and process skews and, and we're trying to keep everything down into the nanosecond. And so the bit rate is really high. I mean, it's a technical, it's a technology miracle. And so, so that part of it is the giant breakthrough of a GPU, if you will, the breakthrough but that's the technology. And there's something you said earlier, customers don't care about your technology, they care about their problem. And I would say that the wisdom of our company was to recognize that a GPU is a is a component, like a DRAM like a CPU, like a network processor. And so it's a technology component. How we how we formulate the entire solution to serve large domains of problems is really the breakthrough from a business perspective. And it was a very challenging business model. And the way we articulated It was, like there's all these microprocessors they're moving at the speed of light, or they're moving at the speed of Moore's law. It everything we build next to it gets sucked into it. Everything we get built next to it got sucked into it. And I still remember our company was teed up as a business case that was in the early days and I wish I have the presentation now, but it was presented to a visiting professor Andy Grove, and at the end of the business competition, we were described as a company that had no hope and would likely go out of business very soon. And the reason for that was a very good one. First of all, the articulated strategy of our company was to go solve problems, to solve challenging problems that are very, very large and sustainable. has great impact in the world if solved, but that are not solvable by ordinary computers, otherwise code for CPUs. And so if you're choosing a problem that's not solvable by CPUs, either that market doesn't exist because the solution doesn't exist, or that market is extremely small, and it's a super computer market. And neither of those conditions leads very well to a successful startup, or a safe, sustainable company. And so that was the articulated mission of our company, this is the problem, we're gonna go solve. The second problem is you sit next to the single, largest, the single most powerful black hole of technology in the history of humanity, that everything near it gets sucked in. And you could be travelling at the speed of light, and still only be stationary next to it. And so the ability for a startup company to achieve R&D scale and not be sucked into the CPU, because of Moore's Law, was probably close to impossible. So between those two

impossible conditions Andy assessed that this company, though interesting, is not likely to succeed. And I would say that the assessment was not wrong. There was one observation or one insight. And as you know, these market insights, or strategy insights, or innovation insights, make a tremendous difference in how technology companies are ultimately formed. And the insight that we had was it was that the PC was coming. And it's not likely that the most important killer app is recipes, it's not likely that the killer app is spreadsheets, it's important for office automation, but it's not likely to be the application that puts a PC and everybody's homes. And we you know, of course, we came from the from the workstation industry, the semiconductor industry, so we know very well what workstations can do and they were all connected by Ethernet and the killer app was a flight simulator, where we're chasing each other around the Bay Area and we're all sitting around just, you know, laughing, chasing each other, having a great time because we're suspended in this virtual reality world networked together. And so we theorized that someday, the single largest consumer application is going to be virtual reality, video games, 3D games, and that everybody is going to be a gamer. Who wouldn't want to be a gamer? And this platform that we're going to create, would enable this new industry, which at the time, I think Electronic Arts had like 12 employees, the 3D graphics video game market was completely non-existent, there were no standards, Direct X didn't exist, OpenGL didn't exist at the time. And the PC had no sound, PC had very poor graphics. The PC didn't even have a have joystick ports that allows you to interact with games. But we imagined that someday the PC would evolve into this thing, if we created this, if you will, accelerate a computing device that would enable these applications and the first application would be a video game. It could be the economic engine that gets the flywheel going. That because this industry, video games, who would want to play it, and the technology was so hard, the technology, complexity x the size of the market would drive our R&D. Well, inside that equation has proven to be true. And this device that sits next to the CPU, the microprocessor that accelerates domains of applications that we have to go figure out. And we have to go evangelize that we have to create the software stacks for . This approach to creating markets has proven to be quite a good one. And now as Moore's Law has ended, we found ourselves in a in a time where this new era is coming along called AI. Computer Science has found a new way of writing software. And this new way of writing software is a very specialised domain that requires tremendous acceleration that microprocessors can do. We found ourselves in the right place at the right time with the right set of skills, the right scale, the right perspective on how to invent and create new markets and real new computing stacks. And so I think it's part insight, part technology differentiation, part strategy, part serendipity.

Bob Metcalfe
I never went to one, but I've heard of things called LAN parties, Local Area Network parties. So apparently, they were playing video games.

Jensen Huang
That's where my work and your work came together.

Bob Metcalfe
Now I know what they were doing.

Jensen Huang
You know, it's really quite remarkable. The video game industry, not only has it become the largest entertainment industry in the world, it's also used as a form of art. The graphics are now so beautiful, and you can mish-mash them together to tell new stories. People use it, people use it. eSports it's one of the fastest growing sports. The people who play eSports have incredible reflexes and you know, snap of a finger thinking and strategizing in real time is like playing real time multi-dimensional chess. And they're incredibly good at it. And we're now we're seeing real sports coming into eSports. And so this is a F1 Racing NASCAR. The actual pro-racers are sitting in these networked, Ethernet worked. pods that are racing each other. And they say its as intense as driving the real thing. And so video games, sitting on the on the fabric that that you invented, makes it possible for us to network all of these gamers together. It absolutely follows Metcalfe's law. And Metcalfe's law has made all of these social games you know, massive online multiplayer games, which is which is built on a network. It is the reason why the video game industry is so large today. It is absolutely based on networks network laws.

Bob Metcalfe
So today we are celebrating the standard, the first spec, the Ethernet Blue Book, September 1980. And I argue that the standardization of Ethernet explains its success that is by making it standard so multiple people could invest in it and multiple people could interconnect with it. I'm wondering, is there now a standard that needs to be made? Or two standards or three in gaming to advance gaming? Are there standards issues in the..

Jensen Huang
No, I think you've got to write the standards makes it possible for horizontal, what is what is described, typically as competition, but it's really horizontal co-opetition, that between a lot of different players, you're able to create a total solution for the marketplace that no single company can. On the other hand, the standard of Ethernet was sufficiently high level that it could be a foundation by which but not so high level that innovation on top could happen. And so therefore the internet was built on top of it. And because the internet was built on top of it, now, that's a platform by which all kinds of applications are built on top of. And so the brilliance of Ethernet, the standard it's really to me a platform, the ultimate platform strategy, the ultimate platform innovation, and sufficiently specified, but sufficiently open ended. That fine balance is what makes it possible for this rich ecosystem to a developed and I think it's completely genius. You know, like poetry, it's completely genius to figure out exactly what the words are, and where to stop it. And that that genius I think has enabled all of these industries to happen. Now, the videogame industry, we have specifications like Direct X and OpenGL and those industry standards have absolutely enabled the video game industry to flourish. Now one of the things that Bob that I want to talk to you about now is what's next, you know, the thing that's really exciting? We started the conversation before we started today, about an acquisition that we announced two days ago. And about the announcement of ARM and my interest in it and the reason why I think it's going to be so important for the next era of internet. The basis of my theory is that the internet is today used largely by humans that are interacting with information and enjoying content and connecting with each other, it's largely its value, its value, the n squared value, the n is humans, largely humans, and

it's created trillions of dollars of economic value and of course, accumulated over time, trillions of trillions of dollars of economic value. And the thing that I'm excited about is the confluence of a few technologies that we've been working on one particularly, that's going to activate a new internet. Now, of course, it's an overlay it's of connection to this internet, but it's really a new internet, if you will. A new internet of autonomous machines and these autonomous machines are going to be interacting with the internet, they're going to be interacting with each other. And because of AI, you could write software for the very first time and, you know, computers, computers at the box, unless they're software. And for the very first time, we can write meaningful software that activates these computers at the edge to do interesting things and otherwise, seemingly intelligent things like drive cars, do agriculture, you know, keep cities and streets safe, and, you know, move trucks around, and robotic arms and things like that. And then that internet unlike the current one we're in, is not sporadic in the sense that we use the internet today, sporadically, and in the context of nanoseconds, you know, we interact with a computer once every infinity practically. But in the future, these machines are going to be interacting with the internet and interacting with each other continuously. The size of the next internet we're going to build is thousands of times bigger than the current internet, we enjoy. And so that's that the reason why we bought we bought ARM. Number one, we would like we love their business model, their business model makes it possible for companies far and wide, whether their industrial equipment or manufacturing equipment, or automotive or whatever it is consumer electronics or whatever it is, for them to be able to build computers. And what we'd like to do is to add to this, this network that ARM has with AI computing technology that we're very good at, to offer it to the ARM ecosystem, so that they could build interesting computers using the same business model. And by doing so, will be able to automate, bring AI and make possible these autonomous machines to be intelligent and connect to the internet in a meaningful way. And so we bought ARM for that reason. So that we could accelerate, if you will, we could bring AI if you will, AI computing to the to the ARM ecosystem. And so that's the journey really laid on the foundation that you laid. We would like to to enable the next generation the next chapter of the internet. And so what do you think?

Bob Metcalfe
Would you call that the Internet of Things?

Jensen Huang
Yeah, absolutely. It's the Internet of Things. And I would say, Bob, that the Internet of Things has been floating in tech for practically a decade now. But we've just not been able to activate it, we've not been able to turn it on per se. And the reason for that is because the things need software. And the things oftentimes can't be connected to the cloud. And the reason for that is because they're too chatty. They're just talking continuously and you can't put the autonomy, the autonomous software, the software in the cloud, it's just too expensive to do it. They're supposed to be things and they're chatty, they need control, they need to make decisions continuously. And so we need to put, if you will, the data centre that's otherwise in a cloud and the software artificial intelligence software in the car itself. And so for example, autonomous driving cars, autonomous driving tractors, autonomous driving, you know, box movers and manufacturing arms and so on so forth. And you know, agriculture, farming equipment and boys on the on the ocean, you know, air quality measurement systems all over the all

over the world. Whatever it is, these things are so chatty. You really want to find a way to put the computer closer to the edge, or sometimes. And oftentimes, the data has sovereignty issues or privacy issues, or it's impossible to move it to the cloud for whatever reason. And so, so I think that that AI is going to activate that those things.

Bob Metcalfe
I was reading recently, there are 7.7 billion humans on the internet. But there are 8 billion things on the internet already, even before, even before what you've been talking about bringing AI in for edge computing.

Jensen Huang
That's right. That's right, I think the I think the internet is going to be thousands of times bigger than it is today, the amount of traffic that's going to go through it is going to be just phenomenal. Now some of the things that I'm particularly interested in and beyond that, in the in the field of computer vision and computer graphics, is to move beyond transmission of video, as the way of being together. You know, of course, of course, today we transmit video, which is a two-dimensional image, but in the future will probably transmit perception. And from that information, I can reconstruct, I can reconstruct you in 3D. And so you and I could be in the same, basically, have a greater sensation of being together. And so virtual presence is going to be possible. And so because of your invention, today, we have the ability to share to be in different places, but sure, one digital space, in the future will teleport ourselves to different spaces, you know, So the teleport, teleport is going to happen. It's just going to be moving photons, not atoms.

Bob Metcalfe
So to build all that and to trigger Metcalfe's law, as I'm fond of calling it, we're going to have to have standards carefully balanced, as you were saying earlier balance. They can't be too low level, they can't be too high level, and you don't need zero of them. You don't need one, maybe two or three will do. But 10 standards is too many, generally speaking. So there's this whole art of standard and today we're celebrating a very successful trade off in that decision in that standards making the Ethernet standard of 1980. Which, by the way, there's no relationship to the Ethernet standard of 2020. There's been a lot of evolution over the 50 years.

Jensen Huang
Well, amazing, amazing work is built on top of your incredible invention. And so I, on behalf of all of the engineers and computer scientists in the world that participate in this industry, Metcalfe's law, this power law, this quadratic has made it possible for us to enjoy an industry far, far greater than any of us could have possibly imagined in the early 90s. And I thank you for that. And for all of the people who are watching, they must, they have surely learned a great deal from all of your body of work and your teachings and your previous online videos. And I enjoy watching them frankly, your humour is incredible. And I want to thank you for the time you spent with me today. And I hope I'll have the pleasure of interviewing you some more in the future, and to tease out the body of the body of knowledge that's trapped in there for everybody else to enjoy.

Bob Metcalfe
Thank you very much.

Jensen Huang
The contribution you made to computer science. I as I mentioned, the technology is one thing, the wisdom about the standards. Another thing, the specification of that standard, you know, the thing that that there's a there's a phrase in our company
as much as needed, as little as possible. That is the specification that I put on almost everything - as much as needed as little as possible. And the reason for that in almost everything. I mean I could find almost no application where that's not true. It could be seasoning of steak, you know, as much as necessary as little as possible. People tend to do too much. Because they're so clever. They want it they want to contribute more. They want to control more. They want to do whatever it is more, more more. It is the genius that keeps it just enough it is Jobs genius to not add that last stroke and that genius requires discipline and vision and self-control and wisdom and just all of that bottled up together and you know Ethernet did that. And you should be proud of that.

Bob Metcalfe
You know, one of the slogans like your slogan just then that we had during the Ethernet standardization process, the slogan was, anything that is not prohibited is mandatory. And the idea there is that standards in the past have been destroyed, but options you some branch of the standard could use this option others, but in the Ethernet world, anything that was not prohibited was mandatory.

Jensen Huang
That's right, therefore, if it's essential.

Bob Metcalfe
Essential.

Jensen Huang
Yeah, as minimum as possible as little as possible. Everything on that list therefore is essential.

Bob Metcalfe
Yeah, it's really guys at Intel wanted the Ethernet to have a 16 bit hardware address, but we wanted a 48 bit address. So the IEEE, standardized both 13. And we had to build chips that could do both. But no one in the history of Ethernet has ever sent a packet with a 16 bit address it just..

Jensen Huang
Well, I guess. I guess if you wanted to hook up a few peripherals, if you want to hook up a few peripherals in an office, that might be okay.

Bob Metcalfe

Never happened. Yeah.

Jensen Huang
Incredible, incredible journey. And I am genuinely excited about the next step and in a strange sort of way Bob, Coronavirus has put a ton more pressure and a ton more focus onto the internet and onto networking. You know, one of the things gosh, I wanted to talk to you about that. And it was something that that I bet you and I wouldn't have conceived of or even realised anyone would crazy enough do in computer science, you know, I would have never thought this was going to happen. So we knew that a computer is something that runs applications and that computer was going to get faster and faster. And we're gonna run bigger and bigger applications and in sometimes the application would be so large that you would you would connect a whole bunch of computers to run that that application and that would that that that application of course at some point became a fairly complex application and it runs in a virtualized environment, what is known as the data centre, and this data centre was used by enterprise and then one day it moved down from the cloud. They were smart enough to build very large data centres that run consumer services using off the shelf, but a fully integrated servers they call them hyper converged infrastructures otherwise, became the word hyperscale. And so, one application would run in a in a server and millions of people could be taken care of in the cloud. And at some point, the applications became so different, some of the applications are very, very complex, some of the applications are really just files, file retrieval, some of the application started to integrate artificial intelligence and speech recognition and image recognition and things like that. And so, the emergence of this new way of doing application composing which is called micro services, and the micro services are very, very efficient. And just the bare minimum system resources to create a container, virtualized container and this container could be run anywhere in the data centre. Or this container will find the computers it could run in a data centre. And that second part that I said is really important because the data centre at this point is no longer homogeneous. It's got some of us got CPU some of as a storage server. Some of it is a GPU accelerated server. And you have an application with modules that are composed and these modules are micro services that are running in different parts of the data centre. So that led me to the question, I never thought that Ethernet is the connection between macros or modules within an application. That is pretty amazing. Today's data centre Bob, is composed, these composable servers compute what is compute, composable this aggregated infrastructure. And these micro services are orchestrated using this this framework called Kubernetes. And you could scale up and scale down, and the data centre is fully utilized. And the shock absorber, the system that's making it possible is the Ethernet switched Ethernet network. And it's putting enormous pressure on the networking. It's now called East West traffic, as you know very well.

Bob Metcalfe
Now, is this why you bought Mellanox?

Jensen Huang
It is it is exactly the reason why I bought Mellanox. It is exactly the reason why I bought Mellanox. It's always good to think about how to build your hardware. By understanding software. It's, you know, understanding software and understanding direction of software inspires you about what's the best way

to design hardware and evolve hardware. And if you look at an application today, it used to be largely microprocessors and interactions with memory, then it became microprocessors, and GPUs, and interaction with memory. And now it's one computing node and interactions with another computing node with the Ethernet network in between that now becomes the critical path of a data centre, if we can make the latency low, both by making the latency lower. And by making the bandwidth higher the throughput of data centres the energy consumed in a data centre, the efficiency of a data centre, the cost of a data centre, you and I were talking about earlier, that customers really in the final analysis, want more for less, the cost of a data centre would drop dramatically. And so we saw this coming. And it's absolutely true, the east west traffic in in data centres is off the charts is going to keep on going off the charts,

Bob Metcalfe
Would you consider or are you considering Ethernet on chip?

Jensen Huang
Ethernet on chip it will probably be not necessary. And the reason for not necessary because the protocols between our processors are very simple. Yeah, they're very, very simple. And we want to keep it as simple as possible. And everything that that one processor could send to another processor is very well known. And so we could make a lot of assumptions about the sender and receiver and so that the protocol doesn't have to be nearly as sophisticated. But inside racks it's basically networking. And in the future, Bob, we're gonna have an application that partly runs at a data centre partly runs in another data centre at the edge, for example, and then partly that applications running on autonomous machine roaming around the world. And so that this way of composing applications connected by networking is fairly essential going forward.

Bob Metcalfe
That's a grand vision.

Jensen Huang
Yeah, it's pretty exciting.

Bob Metcalfe
So NVIDIA, Mellanox, ARM.

Jensen Huang
We might be able to make something of this.

Bob Metcalfe
Yes.

Jensen Huang
Bob is such a pleasure. It's such a great pleasure. It's such a great honour.