Bob Metcalfe Part 1

Derick: Connectivity, more than ever, drives how we live, work and play. Join us. As we engage the people who create, shape, and use the technology that connects us. I'm Derek and I've spent most of my life building and operating networks of every kind.

Brandon: I'm Brandon, and I've spent most of my life building software to empower network engineers like Derick. This is Seeking Truth in Networking.

Intro

Brandon: All right. So Derek, would you believe Bob Metcalf is up next?

Derick: What a way to launch a podcast. When you have someone like Bob on, do you just let them do the talking and enjoy the moment.

Brandon: Yeah. How often do you get to talk to someone who was really there? Someone who's the most qualified person in the world to talk about any technology, let alone Ethernet.

There's so many computing and networking moments that Bob was a part of, and it's such a trip, hearing him relive those moments.

Derick: And some people have an aura that just pulls you in. This whole episode, felt like we're at a campfire making smores and the fires crackling. And the stories are just flowing from this village. Elder who was there when all this legendary stuff happened

Brandon: There was so much good stuff here that we split the full conversation into two parts.

Derick: So grab a seat by the fire and let's get started.

Brandon: Bob, welcome to our podcast

Bob: thanks for having me. I look forward to it.

Moving to Silicon Valley: the early days

Brandon: Few people can say that they were there in the early days . And I was thinking we get started with that.

Bob: Well, I had been in graduate school, working on a project called ARPANET, which you should think of as internet 1.0 and the ARPANET was sponsored by ARPA, the Advanced Research Projects Agency of the US Department of Defense, and ARPA was the principal supporter of computer science research in the world.

I got nine job offers because I was the ARPANET guy and ARPA was giving out tons of money to do ARPANET. And so anybody who wanted to get some of that money, looked at people who are ARPANET people favorably. So I had nine job offers, but I didn't get the job offer I

wanted. I wanted to be an assistant professor at MIT, but I didn't get that offer. And so I stormed off to California and got a job that paid twice as much. Didn't require me to raise any money. Didn't require me to teach any classes and I could do whatever the hell I wanted. So it was really a... sad trade off.

I ended up moving to the Xerox Palo Alto Research Center . And I got there in '72 in the Computer Science Laboratory not far from Hewlett Packard in Palo Alto.

Brandon: Not too many people get their MIT professorship dreams realized that early. So that's not too bad of a second choice.

Bob: Oh yeah, I actually moved to Palo Alto because of the Beach Boys. You've never heard of them, but they were a rock group back in the sixties. And they sung about California in particular, California girls and surfing and things like this. So when it came time for me to settle where I was going to move, from Boston, ... you know I don't tell a lot of people this, but it was the Beach Boys that got me to move to Palo Alto! And boy was I disappointed! Because Palo Alto is an hour from the ocean! And when you get to the ocean, it's 45 degrees. You wouldn't even think of going in there, but I didn't realize that 'til I had already moved by Palo Alto .

Brandon: Maybe you can tell us a little bit about things that you experienced first being part of the Xerox PARC community and the epicenter of a lot of what we consider modern computing.

Bob: Well, Silicon Valley, wasn't always in California. The moral equivalent of Silicon Valley was Route 128 in Boston. That's where all the entrepreneurs were. That's where DEC was. That's where Data General was. That's where Wang was. It was hot stuff. But then in 1972, I moved to Palo Alto and of course, Silicon Valley followed me there.

In June of '72, when we arrived in Palo Alto - I think the term Silicon Valley had been coined, but no one said it with the same sense of awe - and the really cool thing was that the most famous people there, and this is way before Steve Jobs - were, you know, Bill Hewlett and David Packard, people like that, and they were wandering around. They went to events and you could meet them. And it was like a small town almost. Intel was beginning to boom. And then along came '76 when Apple was formed. I never heard of Apple till '79. When Steve called me. When I say, Steve, you know who I'm talking about, right? Well, I, where I'm from, when you say Steve, there is only one, but he didn't show up in my life till '79, and his company was by then three years old. Anyway, we all wanted to be Bill and Dave and the company that I eventually started in '79 was, designed to be like HP. The HP way was the 3com way.

They call it management by wandering around. And so Bill Krause and I, who were the leaders of 3com - we wandered around a lot because we thought that's, that's how you get to be like Bill and Dave. But that touches on an important point having to do with role models.

That is, a lot of people, would like to be like Silicon Valley. And one of the things you need to become Silicon Valley-ish is to have role models, and we had role models - we had Bill and Dave wandering up and down Page Mill Road there in Palo Alto.

Arriving at Xerox PARC

Brandon: So tell us, you enter PARC; it's already been around for a little bit at this point . What did you see there?

Bob: Well, what had happened was that Xerox made some good personnel decisions and they recruited a guy named Bob Taylor, who by the way, is an alumnus of the University of Texas. I hasten to add where I work. And Mr. Taylor, he was a psychologist and he worked for DARPA and he ran the Information Processing Techniques Office of DARPA. Xerox recruited him to start the Computer Science Laboratory at the Xerox Palo Alto research center right there in Palo Alto . So you should think of Xerox PARC as a lab, a Stanford lab, although it was owned by Xerox, and Taylor having come from ARPA, knew all the best computer scientists in the world cause he was sponsoring them. So the computer science laboratory started recruiting them. So by the time I got there, they had guys, Alan Kay, was there Butler Lampson, was there, a long list of then-famous computer scientists like Alan Kay, literally the inventor of the personal computer, I would say - not counting Steve - and Butler Lampson was sort of the genius from Berkeley who had invented time-sharing computing in his spare time. So I became the networking guy there. I lucked out. I had put MIT on the ARPANET. And then I put Xerox on the ARPANET. But after those two projects were done, my next project was really important: the computer science laboratory decided to invent the personal computer and put one on every desk.

Ethernet: it begins.

And of course a network would be needed to connect them together. And who around here does networking? The question was asked. Me. I'm the networking guy here! And so they gave me the job of building a network to connect, believe it or not, a computer on every desk, not one per building ... on every desk.

Brandon: So was this as new a concept for you as it was for everyone else in the world at that time? I mean, did you know it was going to happen eventually?

Bob: Well, we all knew the computers were getting smaller and smaller and smaller, and my recollection is no one had a clear image of what was about to happen. I mean, you recognize it. It had a bitmap display, a bitmap display was unique! In those days, everything was character oriented, but the Alto - it had a bitmap display, so you could put not only characters on the screen, you could put pictures on the screen! Pretty awesome.

And then there was a thing called a mouse, by the way, I have a mouse right here instead of using keys with arrows on them. By the way, my Macintosh does have keys with arrows on them. was the method of moving the cursor around the screen. Instead, we, Xerox, stole the idea of the mouse from SRI. A guy named Doug Engelbart there had invented the

mouse, but it was Xerox that made the mouse very famous, and effective. And so the arrow keys, you haven't used the arrow keys recently.

Brandon: Occasionally I navigate a text document, but yeah, mostly I use

Bob: And then I got the job.

I had, prior to all those, been an operating systems person. I was down in the plumbing at the borderline of hardware and software. At the same time we were building this display that had bits on it. We built a laser printer. So you could print a picture. You could print the arbitrary images on these screens, including characters with different fonts. So Xerox became font central! We had lots of fonts that we would display on our screen, but we also had to print them. So we build the laser printer, the first laser printer, and they gave me the job of building the operating system for this printer. And then of course the network, and the only way you could print on this printer was through the Ethernet . So that was a bit of a requirement. In other words, when it came time to build the network, it had to be able to serve this printer. And this printer printed a page per second, 500 dots per inch. And if you do the math that's well over 20 megabits per second. And the current methods of communicating operated in the hundreds of bits per second, not the millions of bits per second. So we made the network - Dave Boggs and I - made the network as fast as we possibly could, and it turned out that the constraint was how many chips you could fit on the card. That is, the network was to be a card that you plugged into this Alto PC. So we had to design a network that fit on the card and that determined the speed.

Brandon: What was the speed?

Bob: 2.94 megabits per second. So on one day I had the most advanced computer terminal most people had called the Texas Instruments Silent 700 that ran at 300 bits per second. The next day I had 2.94 megabits per second at my desk. And if you do the arithmetic, that's a speedup of approximately 10,000, not 10%, not double. Not 10 times, not a hundred times, but 10,000 times faster. And this would later lead to the ability to upload cat picture .

Brandon: And at that time, do you have any idea where any of this would go? Did you have a desire to download cat pictures, to do video conferencing, to use interactive apps?

Bob: Ethernet was not built to carry a video. it couldn't possibly carry a video because as you know, video is terribly synchronous and analog and Ethernet was asynchronous and digital. So no one imagined that we'd carry a video. Of course, the only thing that Ethernet carries anymore is video. So like this conversation right here is is video and most internet traffic today is video, but the original Ethernet was not designed to carry video or telephone - eventually Ethernet became a method for hearing voice call or voice telephone telephany, but it was never designed for that. We never imagined that. In 1972, there was no notion of the World Wide Web.

Addressing - how many bits are enough?

Derick: I actually downloaded and read the Xerox paper you had written. And I drew some important conclusions from this paper. There was eight bit addresses, right? And you talked about scaling in a sort of positive way, rather than dismissive, you said you could get so many hosts on a thick net cable. And if you needed more, then you could interconnect multiple of these networks with a gateway and you referenced Vint Cerf's 1974 proto -IP paper. And, in that paper Vint he did something that, people always love to point out this luminary of networking said that eight bits for a network address should be more than enough and it should be sufficient.

Bob: So in 1973, when Ethernet was invented, TCP/IP , the standard protocols of the internet, were also invented that same summer. And Vint had a seminar at Stanford, which I attended to take the ARPANET and update it. We had been working on this, on this Ethernet thing for awhile and Vint and I both came to the meeting, seeking a network number. We wanted to have in the protocols, a network number - which did not exist in the ARPANET. But we came at it from opposite directions... Dave Boggs and I and others were building Ethernets and putting one on every floor and every building of Xerox and we needed a network number so we could interconnect them.

So you could uniquely specify... "we're sending a packet to that network over there". Vint came to the same seminar, wanting a network number. But he wanted one network per country. He was interested in interconnecting an ARPANET in each country.

And as you may know, the number of countries is less than 255. I guess it's 190 or something. So it turns out one byte of eight bits holds 256 different values. So Vint was kind of pushing for an eight-bit network number, and I was pushing for a network number, but I had many more networks in mind than he did.

So we ended up with 32 bit network number, which was way bigger than Vint wanted. He only wanted eight bit, but it seemed like infinity to me too , but it was a little odd. We were both wanting there to be an internet. And the term internet hadn't been coined then, it wasn't 'til later. We both wanted a network of networks, but he wanted it to be the international one, and I want it to be the local area network, one network of networks.

Standardizing on Ethernet vs the competitors

Brandon: All right. So maybe let's continue towards Ethernet becoming more of a commercial thing. So you've got this almost three-megabit Ethernet ; it's deployed at PARC. Where does it go from there into 3Com? How does this all get started?

Bob: Xerox PARC was in Silicon Valley. We were surrounded by entrepreneurs and, everybody was starting companies, even college dropouts, like Steve Jobs, were starting companies. it was all almost automatic, that you would after a while, leave the big company and start a company.

So I did that. I left in January of 1979. I left Xerox to pursue entrepreneurial ambitions and that was my resignation letter. I'm leaving to pursue entrepreneurial ambitions, but I didn't

have a company in mind. And then, in February of '79, I had a consulting gig at DEC. DEC was then the second largest computer company in the world, right behind IBM.

And the head of DEC was an admirer of Ethernet. So he asked me to design an Ethernet for DEC, but I'd said to him, I can't do that. For one thing, I feel a sense of loyalty to Xerox. They supported me for eight years - and two, is I've just finished designing the best network I know how. So your network would not be as good as Ethernet. So he said, or one of us said, we don't actually remember who, who led it. We said, well, why don't we go to back to Xerox and work together?

So I drafted a letter and Gordon Bell, the head vice president of research at DEC signed the letter and we sent it off to Xerox and it said, Hey, why don't we work together and make an Ethernet that we can use to connect DEC computers and Xerox copiers and printers?

And Xerox said, yes. And then I found a guy from Intel. Who is looking for a new process, but he needed a standard to make a new chip around. The Intel strategy was to adopt standards and then make chips to implement the standards. So , we then created a three way partnership, DEC, Intel and Xerox called DIX. And I was behind it creating that thing and then we were notified that we could not meet to design this new network, because this would be in violation of antitrust law but it turns out that my fraternity brother, Howard Charney, had just finished suing IBM on antitrust.

So I called them up and I said, Howard, they won't let us meet! And Howard gave me a list of five things we had to do to meet. No marketing people allowed in the meeting. There must be an observer from the US government. The goal of the activity must be to create a standard that's open to the industry and a few others.

You weren't allowed to fix prices, for example. So then the meeting started and DEC and Intel and Xerox agreed to produce a standard Ethernet and the way standards were created in those days - there was an organization called the IEEE , the Institute of Electrical and Electronics Engineers, that had a standards process.

In fact, the one we liked was called the HP instrumentation bus, which had IEEE 488 had just been standardized. So we found the guy who standardized the, uh, HPIB and we got him to help us create a project at IEEE to standardize Ethernet. Project 802. And you notice there's no dot in the name of project 802, not 802.3. Not 802.11. Note, the dots came later, but this committee's purpose was to produce a standard Ethernet. And so the meeting began at the Jack Tar hotel in February, 1980 in San Francisco: IEEE Project 802 to standardize Ethernet.

And in September of that year, DEC, Intel, and Xerox produced a spec called the Ethernet Blue Book and submitted it to the IEEE to make a standard. And everything would have been fine except IBM and General Motors decided they didn't like Ethernet. So they submitted their specifications. So it then became a three-way race, Ethernet versus IBM Token Ring versus General Motors Token Bus.

Brandon: So how would you describe Ethernet vs. Token Ring?

Bob: Token Ring was designed by mainframe people and they were not cognizant of a key architectural feature of the internet. The Internet had seven layers. The Open Systems Interconnection model was an attempt to organize all the functions of networking and Ethernet was designed in the context of those seven layers. In fact, layer one layer, two Ethernet was a standard up at level one and two.

But the IBM Corporation, they had a lot of problems, but one of them is they were used to making their own standards because they were 95% of the computer industry. So they were used to just issuing standards like SNA - System Network Architecture - became the way you connected computers together, thanks to IBM. And by the way, I used to give talks and then I would annoy IBM by calling it " snah " and they liked to call it SNA, but I called it "snah" just to annoy them. And so the token ring had too many features. For example, it had acknowledgements, you know, you send a packet and you want to know if it got there.

So back comes another packet. That's an acknowledgement that says, got it. Send another one. Ethernet doesn't have acknowledgements. It just has packets. Now some of those packets are acknowledgements, but Ethernet doesn't know that they're acknowledgements, it's just carrying packets - whereas the token ring felt obligated to know that the packet had been received. So in the very lowest-level bowels of the token ring was this acknowledgement mechanism, which added to the complexity of the chip , so they would take turns sending based on the circulation of this t oken, but Ethernet didn't have that feature. And so it was much simpler, faster and cheaper than the token ring.

Brandon: But it sounds like Ethernet's a Wild West. How do you know your packets came back? How do you ensure access to bandwidth? It had to push people and make them go, this is crazy. This can never work.

Bob: God, that's what they said. Well, they forgot that on top of level one and two was level three, four, five, six, and seven. The function of acknowledgements and retransmissions was handled at the next layer up. So there were acknowledgements, they just weren't built into the low level hardware; they were a higher-level protocol. Later the World Wide Web came up here at layers three, four, five, six, seven - above levels, one and two without changing one and two, that was the miracle.

The worldwide web worked on top of 1973- standardized technologies.

Forming a Company: 3Com

Brandon: So you're doing this standard - this is right around 1980. How do we get to 3Com?

Bob: June of '79, it was clear to me that DEC, Intel, and Xerox were going to make a standard and we had lined up the I EEE to do it. So I figured that there had to be a startup in here somewhere. So, started 3Com on June 4th, just me. And then I, I thought of three people with whom it would be useful to start this company. John Shock, Ron Crane, Greg Shaw. And I invited them to dinner at a restaurant, which has only recently closed called the Bella Vista, overlooking the twinkling of the Silicon Valley. And, one of the guys, John shock immediately said, no. The other guy, Ron Crane, said yes, but not for another year because I

have work to finish up at Xerox. And Greg Shaw, the third one said, yes, let's get started tomorrow.

So Greg Shaw and I founded 3Com, June 4th, 1979. And our goal was to change the world in terms of connecting computers. Until then, every computer manufacturer had their own network architecture. IBM had "snuh" - what I called "Snah" - System Network Architecture. DEC had DECnet. Wang had WangNet.

So we were going to do internet using standards, connect computers together with standards. So we named the company 3Com: Computer Communication Compatibility . That was our business: making computers compatible so they could talk to each other... and so our strategy was to adopt standards. We adopted the Unix operating system. We adopted the TCP/IP protocols that were about to be standardized, and we adopted Ethernet. And so we started making products to implement those for people who wanted to connect their computers together. For awhile, we were consultants.

The Ethernet Handbook

Our first product was a book - a local computer network vendor list. And we sold this book for 250 bucks a copy.

Brandon: Bob, let me jump in - what kind of product was that exactly though? I don't know if I caught it.

Bob: Yeah. The idea was, to create an industry and the industry would be the local, the local. computer networking industry. Later, they would be called LANs. But back then I was calling them local computer networks, which is in the paper. The title of the paper was Distributed Packet Switching for Local Computer Networks.

Cause I felt it was important to make the distinction between terminals - dumb terminals, which was the dominant way of networking computers. So it was called the 3Com Local Computer Network Vendor List. And it was basically Xerox copies of product literature of companies that I thought were in this industry.

And they didn't have to think they were in the industry - this is my book! So the book came out describing these products. I tell all these products and sort of described how they were part of the local computer networking industry, and I wrote a preface describing it, and then a young lady named Sarah Shevock (sp?) was our receptionist, and we gave her an embossing machine.

And her job was to not let anyone out of our offices without buying a copy of this book. And eventually my wife, Robin, and I became the publishers of this book. And then eventually the venture capitalists made us stop doing that, but we made a hundred grand a year selling this book to people and what it did, I mean, it's real impact was getting people to think that there was an industry. And we were about to be the lead player in this industry. We, 3Com, by selling our Ethernets.

Brandon: So you were selling the book that would give people information to help create the network that would give them the ability to get that information in the future.

Bob: Yes. And it worked. The book later became known as the Ethernet handbook. I got a bunch of copies here, and Robin and I, every Saturday morning would ship 10 or 20 of them, but with our own hands would put them in the envelopes with the invoices and then we would carry them down to the post office there in Portola Valley and mail them, and then this money would come raining in a 250 bucks a copy. But then the venture capitalists who had invested in 3Com said did not invest in a publishing company. And we did not invest in a consulting company. We expect you to have products. So we develop some products.

Brandon: How many books did you end up selling?

Bob: What's 100,000 divided by 250 - that's 400 copies for five years. So 2000 copies roughly. And every VC - oh, by the way, at 3Com's office, I met my wife, I founded my company, all on a place called Sand Hill Road- which is where all the VCs were. So every VC on Sand Hill road had a copy of our book. Because they would come visit us and then Sarah would corner them with her embossing machine and they would you'd would be forced to buy a copy. Some paid cash - that was welcome.

The IBM PC Emerges

Then we started working on our products and there was a problem. Our goal was to network personal computers in a local area network.

But there weren't any personal computers, as you can see the problem there. There were Apples - Apple was in '76, so there were maybe a hundred-thousand Apple computers in the world, but there really weren't any PCs. Fortunately, in August of 1981 IBM introduced the IBM Personal Computer. So then we decided that it was our destiny to network them; even though IBM felt they had their own network, they made their machine open so that we could sell Ethernet for IBM personal computers and the clones. There were a bunch of like Dell and Compaq.

Brandon: Wow. So they, they made the machine open cause they knew they wanted to network these computers. Were there any other uses they had in mind, things like graphics cards or other typical extensions.

Bob: Oh, you know, the option slot idea was general. It wasn't specific for networks. So they had their own network cards. They had their own graphics cards. They, you know, everyone was making cars, but it was open; the idea of the IBM PC in 1981 was an open machine. It was running Microsoft DOS and eventually 1 2 3 became the killer app.

Steve Jobs demos a killer app

Steve Jobs called me one day said "You got to come down here!" so I went down to Cupertino and he said, "Look at this!" - he opened up his Apple II and pressed a few keys. And on the screen was a grid of little squares. He said, "Watch this!" - and he put the number 29 in one of those squares. Oh, that's interesting, Steve. And then he put a formula in another little box that took the number 29 and multiplied it by another number, and he

was demonstrating the electronic spreadsheet to me 'cause I had never seen a real spreadsheet. So I didn't really know what this was, but you know, Steve thought it was insanely great. And he correctly identified this as the killer app for the IBM personal computer. And so spreadsheets became, and there was a big fight about whose spreadsheet, but eventually as you know, eventually, well, first it was VisiCalc and then it was one, two, three, and now it's Excel. And so Microsoft always wins the third time. And so we now had spreadsheets.

Making Ethernet Affordable - The Curve

Oh, so we had the problem - that there were no PCs. So we made Ethernets for a mini-computers - DEC minicomputers, PDP-11's, and faxes and so on. I sold my first Ethernet card for $5,000. You plug it into a big option slot and it would hook you into a cable and make you part of an Ethernet.

Of course when Bill Krause joined 3Com - I helped recruit him to 3Com - he was my adult supervision starting 1981. First thing Bill wanted, he said, "Bob, $5,000 is way too much." He said, "I want you to draw a graph showing an exponential declining curve that would project the cost of Ethernet in five or 10 years so that I can show this to customers... I want customers to know that this standard's going to get cheap someday. " Of course those of us who were there ahead of Bill says, "God, what does this guy want? He wants us to predict the future?" We don't know how to do that, but he insisted he a very planful guy.

So we had an offsite and we drew the curve. Of course. So if you draw the curve, you have to explain how you're going to reduce the costs. One chip solution, two chip solution, three chip solution, the movement of the transceiver inside the card. And each of these became came a milestone in this exponential down to a hundred bucks per Ethernet connection. Of course they're free now. So it's kind of hard to think about that. So we had a plan, thanks to Bill, to drive the cost of Ethernet from $5000 down to $100 over a 5 or 10 year period.

Inventor(s) of packet switching?

Derick: You mentioned in there that you started off with Unix workstations and I have in front of me a presentation from November it was like a seminar you did. And I see, I didn't know this, that Paul Baran was a director at 3Com - he was one of the two people that independently invented the concept of packet switching. Is that right?

Bob: That's right - that's exactly right. I Met Paul Baran in an odd way. I was, I was job hunting. This is during that time when I got nine job offers and I visited the head of ISI at the University of Southern California and I was explaining to him, I often make this mistake - I was explaining to him as if he didn't know that the Internet had been invented by a guy named Larry Roberts, who Bob Taylor hired to build it.

And this guy looked at me like I was a complete idiot and he turns around and there's a 12-volume set called On Distributed Communications - I believe it was 1964 - by Paul Baran.

And he said, Bob, Larry Roberts and Bob Taylor, are Johnny-come-lately. Is it Paul Baran? Just the guy who invented the internet and packet switching and all that stuff.

So then I made it my business to meet Paul Baran and recruited him to be on our board. And he's a fabulous guy. And may he rest in peace. There are two inventors of packets. The packet switching is the fundamental technology of the Internet that girds the rest of it. And for a long time, it was generally believed that there were two inventors who separately, independently, uh, one in England, Donald W Davies and one in the United States, Paul Baran both came up with packet switching. In fact, it was, I believe it was Davies who gave it the name packet switching. And now there's a fight about who really invented packet switching, and Paul and Donald are not around to defend themselves. Anyway. Yes. Paul Baran: the inventor of the internet. And actually I agree that Vince Cerf should be called the inventor of the internet. So I've agreed to that. But packet switching, I'll give to Paul Baran.

Applications of Ethernet

Brandon: So it's, the mid eighties , your startup is growing. You're getting Ethernet deployed in computers, personal computers, as they're now called, or desktop computers, on your way to a million ports before the end of the decade.

Does it go according to your expectations or are there things that are holding its adoption back that you didn't expect?

Bob: This all happened in the 1980s. That was the, uh, peak of advance. And the were a number of problems. One was, there was no software to run on top of. We couldn't figure out what everyone was going to use Ethernet for. And I would frequently say there's millions of uses. And then Bill Krause would say, well, name one. Well, I can't think of one off-hand, but there's millions of them, but eventually we settled on PFMT S - print, file, mail, terminal, and stubs. So that became the, uh, initial killer apps for Ethernet .

Today we call stubs APIs. We would offer API access to the underlying packet mechanisms. So you could write apps that sent packets back and forth. And then we implemented PFMTS, and began to sell Ethernet.

Getting customers to see the product as useful

I remember we came up with a three node trial. So if you had three IBM PCs, we would sell you three of our cards for $3,000 with a diskette with the software on it, the PFMTS software - and this software would share the printer. Imagine three PCs sitting on a table with a cable, connecting them together with the card in each one, connected to the cable. And then you'd load this diskette software onto each of the three of them, and then if there were a printer connected to one of them, it could be shared by the other two PCs. And if there were a hard disk on one of the PCs, you could share it with the other two PCs - you wouldn't have to buy our disk for all three of them, and the terminal service, and so on. So we started selling these trial kits. They sold pretty well. Of course, three grand was a

small-enough number that most people could do it without getting management approval. And so we started selling these like hotcakes. This is in 1983 and our customers all came back. I was running sales and marketing. We had six salespeople - and they came back and they said, Bob, "The kit works exactly like we said it would. It meets specs, but our customers do not find it useful." Whoa. So being head of sales and marketing, I had to solve this problem. Our customers did not think our products were useful. So I went to Stanford one night and this, slide is on the Internet. And I drew on an Alto. We had given it, we Xerox had given Altos to Stanford.

I was a professor at Stanford, so I snuck over there and I used an Alto and I made a 35 millimeter slide by the way, PowerPoint. I was on the board of PowerPoint, the company that made PowerPoint, we sold it to Microsoft in 1987 for $14 million. It was 1987. This is 1982 or three. So there was no PowerPoint.

So I made a 35 millimeter slide that explained why our networks were not useful. And basically it said that the cost of the network is linear in the number of our cards that you buy, but the number of possible connections you can make with those cards goes up as the square of the number of cards you buy.

So there's this point where that quadratic overcomes the linear and your network becomes more valuable than it costs. And the reason, the reason that all these node networks were not useful was because they weren't big enough. And that's what the slide said . They had not achieved critical mass and the remedy of course, was to buy more of our products and build your network to 30, 40, 50 nodes.

They believed it and they bought it and it turned out to be true! Which raises the question "Was I lying?" You know, marketing sales and marketing people are sometimes accused of lying. And I came up with this slide, predicting that if you grew the networks, they would be more valuable.

And they were. Was I lying? And the answer is no, I was not lying because I had a time machine. I had been to the future through the Xerox Palo Alto Research Center. I had been in a world where everybody had a PC and all of them were connected with LANs and we had laser printers and routers and we had internets and it was good. I knew it was good. So I was not lying when I made that slide. And the customers believe this slide, I guess customers are pretty smart. And we went public in 1984, based on the sales generated by the proliferation of PCs and then the PFMTS uses of those PCs.

Brandon: So there's a value of connecting things going up as you add more of them, some have called it Metcalfe's law. What would you call it?

Bob: You can Google the slide up. It's out there. The systemic value of a network , V, is approximately equal to N squared, where N is the number of nodes. In 1995, in Forbes Magazine - a man named George Gilder had a lot to do with Moore's law - calling it Moore's law. So he wrote in a book, he wrote called Microcosm, but then he wrote another book called Telecosm. And in that book, the law, it wasn't Moore's law - it was Metcalfe's law. That the value of a network grows as the square. So George made Metcalfe's law

famous. I've been defending it ever since 1995. One of my hobbies is defending Metcalfe's law.

Brandon: Nice. I'm looking on my screen at a picture of that, and as an engineer, I love when there's a single graph that explains a deep concept and you can draw it on a napkin.

The Critical Mass Point

Bob: Yeah, the critical, you know, a key part of that slide was the critical mass point. The internet passed critical mass a very long time ago. And so you need to get above critical mass before you see value. That point has sort of been made a long time ago. That is, unless you recurse. So if you look at Facebook as one network, well, then it passed critical mass a long time ago, but if you look at each group of friends, as a network on Facebook, they are going through this n-squared growth rate and they have a critical mass phenomenon at another level up.

So there's the base level of Facebook. And then there's each of the friends groups, which are another network and they also obey Metcalfe's law.

Surprising use cases of Ethernet

Derick: No one would have ever predicted that Ethernet would be, probably by number, the biggest point-to-point protocol there is . In fact, it kind of weirds the mind, right? So many DSL lines run ethernet. But what is probably the most surprising use case or use of Ethernet that you see today?

Bob: So I remember at Xerox there was a guy named Dan... Dan Swinehart, but he decided as we began to spread Ethernet through Xerox, that it should carry telephone calls! So Dan started writing telephone protocol software that ran on Ethernet carrying voice. All the bell-heads in the world, said "You can't carry voice on a network that has stochastic retransmissions and lost packets and stuff. You can't do that!"

Of course, the network is running a million times faster than the networks they're used to. And so it could, even though it was stochastic, it could carry voice. So you had voice over IP. So now most of your telephones, if you have any, if you have a line to telephone now, they carry voice over IP and voice begins with V but so does video.

So very soon, VOIP became video over IP and people started carrying video streams over the internet. Of course they were laughable. They were little postage-stamp videos of, you know, a small fraction of the screen. And they were jerky as the characters went like this, and we all mocked the video on the internet and then a few years passed and now we don't mock it anymore - it completely dominates the internet carrying video, as we are doing right now. And it works pretty well, doesn't it?

Derick: Yes, it does.

Brandon: It's working pretty well for us. I mean, in 2020, the world changed and the importance of video...

Bob: The internet came just in time for COVID. So COVID-19 comes along. We can't go to work anymore. We have the internet. And the amazing thing is all those videos goes onto the internet and it works! Look, it's working right now! Look at this! So it's like the internet was built for COVID-19!

Brandon: I get to see your eyes light up, even though our listeners can't see that. Did you foresee any of that?

Bob: So yes, I foresaw the video internet. What was happening is the media was getting hotter and hotter and hotter, and the next step was video. It was just the next step. I could see that the video internet was inevitable and the speed of ethernet in the internet in general , was sufficient to carry video. The protocols needed to be tuned up, but after that it was automatic.

On to the browser

Brandon: So the other question I wanted to ask related to that is, how we got to a browser .

Bob: So this is quite a bit later. A guy named Andreessen at the University of Illinois wrote a browser based on the World Wide Web which had been developed at the physics lab in Europe called CERN. The grad student at Illinois writes a program based on some protocol - three protocols actually, URLs, HTTP, and HTML, three sort of lame protocols; each one of them was sort of ... adequate. And I was then publisher of InfoWorld magazine writing a column on the internet and I had to give awards each year to the best coolest technology .So I decided to give Mark Andreessen an award for Mosaic! He had written Mosaic, this Graphical User Interface for the World Wide Web, whatever the hell that was. So I flew off to Champaign- Urbana to present my award to Andreessen ...and he wasn't there! He had left in a huff ...he had stormed out the door because the University of Illinois mistreated him - you know, claimed ownership of it and so on. So he went off to California and he founded Netscape Corporation. And he wanted to call it Mosaic Communications, but the University of Illinois wouldn't let him so he had to call the company Netscape and they claimed ownership of the code .But it was only 9,000 lines of code. So he hastily rewrote his browser and founded Netscape and Netscape started proliferating and it became an increasing fraction of the traffic on the Internet. . So this this is what I would you call this web traffic what we're doing now this is video..

Brandon: Think so ...

Bob: It's on the web, isn't it? And the web is on the internet, which is on the Ethernet. So these layers happen. So, so the World Wide Webs sorta proliferated. The Xerox Research Center did not anticipate the World Wide Web , or at least the meetings I attended to - no one said, Hey, there's going to be this web someday.

Ep1 Outro

Brandon: Stay tuned for part two, where we take a step back and talk about what the contributions of Ethernet ...really are. Plus Bob walks us through the tech he's most excited by today - probably not what you'd expect. See you soon.

Ep2 Intro

Thanks for joining us. If you're coming in fresh, take a moment to check out part one, which covers the origins of ethernet and Bob' stories from early Silicon Valley. With that said - let's get started with Part Two.

Ep2 Outro

Derick: Thanks for listening. Just a quick reminder, check out the show notes, at seekingtruthinnetworking.com. We have links to some of the stuff we talked about, including some of those early papers, as well as some pictures of Bob, at work, inventing Ethernet.