



Interview of Geoff Ralston

Interviewed by:
Marguerite Gong Hancock

Recorded: May 13, 2015
Sunnyvale, California

CHM Reference number: X7501.2015

© 2015 Computer History Museum

Geoff Ralston: I graduated in 1992, in December of 1992. And in my last semester at school I had taken a class in entrepreneurship and I got together with three friends, one from France, and two from Israel and we created our project. And our project was to create a global and international shopping network where we would provide products from different countries, especially for the diaspora of those countries but for anyone-- like you could move products all around the world and make them available via this international network. This is 1992.

Marguerite Gong Hancock: So early.

Ralston: And so, we went and we worked on it for about three months and then we all sort of said, how do you even think about putting together the infrastructure for this thing? It was totally an entrepreneurial failure on our part, because we didn't probably try hard enough, because there wasn't really anything there. I knew a lot about the Internet exists but just thinking about getting that Internet to consumers so they could buy things... It was starting to happen but there were these islands of CompuServe and Prodigy and AOL starting but it just seemed impossible. The timing wasn't quite right. But that being said, the year before Tim Berners-Lee had created HTTP.

And so we gave up, and I look back and Amazon was around the corner. It was about to happen. So it was really interesting for me to look back. And then, of course, I had been thinking about a worldwide network for a long time -- not because I was particularly farseeing. It's because I read science fiction. <laughs> And people like John Brunner and others had written about networks.

An interesting aside if you're interested, John Brunner wrote this book called, "Shockwave Rider." I read that in the late seventies when I was in college and it talked about this Net and it was-- you might guess, I don't know if you've ever read it but it was connected to future shock and the advance of technology. But there were two concepts that were really important and it connected to me in interesting ways. One is this concept of a Net, a world Net that anyone can connect to which stuck with me forever. And, of course, was very farseeing because 30 years later, 40 years later here we are. But interestingly, he also invented the concept of the computer worm, like the computer worm, right, is that virus that goes from machine to machine. And it just turns out that the founder of Y Combinator where I work and I was a partner, Robert Morris, invented the first real computer worm and unleashed it pseudo-accidentally on the world in the late eighties. He actually got in a lot of trouble for it as well. So that book has kept coming back and those concepts have kept on coming back to me. So it's interesting.

A lot of people foresaw what was going on, but it took a few key innovations. If you look at history and you look at the confluence of events that come together to create sort of the great event, World War II, the invention of the Internet, the Manhattan Project, the invention of the Internet, things like that... the creation of the World Wide Web and the connected world as we know it today had to have a bunch of

things come together at the right time and it was a really fascinating time for me. I remember-- we're not rolling yet, we are?

Hancock: We are.

Ralston: I don't know if this is useful to you...

Hancock: It's great.

Ralston: I remember so-- my quick individually story if you're interested...

Hancock: Yes, actually, that's where I'd like to start is you're kind of laying this framework of this confluence of events, this sort of combustible time that created a spark that changed the world. So let's tie in your story to that.

Ralston: So, I tried to start this worldwide shopping network with my fellow students from this business school in Europe and we failed. And I wasn't feeling very good about that and I came back to the United States not really feeling very good about anything because I was an engineer. I was a coder. And I went to get a business degree, because I really liked business and I liked the idea and I wanted to start a company. I had always wanted to start a company. I had been thinking about starting a company forever and hadn't done it for like a decade because I went back to school twice and I got good opportunities to go abroad and I didn't do it.

And I kind of went back to the United States with my tail between my legs and went back to work at HP in of all things a business development job. And of all of the things network management which I hated. And it was a very frustrating kind of unhappy time in my life. And then an old friend of mine, I can remember this like it was yesterday although it was 21 years ago now. It was 1993. An old friend called me up and we were just talking about like catching up because I had been overseas for five years and we hadn't really seen that much of each other. And he said, "By the way, have you seen Mosaic?" And I said, no, I have no idea what you're talking about. What is it? And he said, "I'm not really sure but it's like it allows you to look at the network." And I said, you mean like a protocol analyzer? And he said, "No, not like that. But you have to check it out." I was like okay, whatever.

And it turns out that getting Mosaic which was the browser that Marc Andreessen created at Illinois. The interesting thing about Mosaic that was kind of unique and different was that for the first time it took text and images and put them together. And that was one thing. And the second thing that was really interesting about Mosaic, especially in light of what's happened since then, is that they distributed it completely for free. Here it is. Take it. Freeware, shareware. So it was not that hard to get it. I found it and downloaded it but it was actually pretty hard to get it to work, because I was working inside a corporation and you couldn't connect to the outside world. You had to setup this crazy proxy server and you had to do a bunch of stuff.

But in a week or two I set it up. And all of a sudden I was sort of stunned. I quit my job. I can't remember exactly, but I think I quit it the next week. I said this is it. This is finally it. And I proceeded to make every possible mistake you can make as a startup founder who really doesn't know a lot about what you're going to do. My boss at the time at HP looked at me and said, "What are you doing? If you want to go do something but don't leave without an idea. Don't just go. You're crazy." And I was like, but this is it. This is what I've been waiting for. It's it. I'm sure.

And that was an interesting. I think you had to be a true believer then because I remember I would be driving around Palo Alto with my then girlfriend, now wife, and this is like '93, '94 and every once in a while you'd drive past a sign and there would be a URL on it. And I would go, "but look it's real. This isn't just—" She would just say, "Yeah, whatever, sure it is."

It's hard to imagine back then, but it's like Marguerite were talking earlier about what Gordon Moore was saying. It just seemed like a nice little networking thing, and people are going to put it together -- and who knows. But there's a series of events, and a series of events, and it wasn't that long before we were buying our first filer -- I'm talking about NetApp -- to build a web-based email product that ended up being the reason why Four11 was purchased by Yahoo!, although we were originally created as a directory service.

Mike had mentioned the interesting thing for him about everything that's been created is the way it connects human beings. When we talk about the connected world, we do sort of talk about companies connected to companies, or products connected to reviews, some atom connected to some atom or some concept connected to some concept. But, of course in the end what we're really talking about is connecting people with people. And we saw that. We said, what if we created a way for people to more easily connect to people, created a directory of people so they could find people in ways that they could never find them before? Remember back then, we had these big books that got delivered by the phone company. Remember that? They would drop off "boom" for free, too. "Boom!" Here's your phone book and that was the phone books. Still like that line...

Hancock: Only the local. You had to go to the library...

Ralston: The yellow pages, too. The yellow pages. Remember the yellow pages the big-- and I still remember that great line from that Steve Martin movie "The Jerk", "The new phonebooks are in! The new phonebooks are in!" Remember that? That just seems so long ago like who would do that? But we said that seems anachronistic. Let's help people connect with people. But, of course, that's not really where it was. Where it was at was connecting people in a slightly deeper way which was with email, which of course, had been around since the seventies.

But with the creation of Rocketmail which created Four11 and the creation of Hotmail, and Rocketmail got bought by Yahoo! Hotmail got bought by Microsoft. Suddenly email went from curiosity, and then this corporate thing that was used internally, and then somewhat separately. And then in pockets that Prodigy

mail and AOL mail, it went to being universal. And that was, I think, certainly from a communications standpoint, one of the real fundamental shifts that happened back in those days.

And the thing about NetApp which was really interesting is that NetApp -- I think they were public when we started talking to them. I think they went public about that time. They were still a startup, and it felt like everyone was-- we used to talk about at Yahoo! about the fact that it was hard to actually get vendors for almost anything because we were paving the road as we were going. And where the road was going wasn't so clear. We just knew it was going. And it was going fast because the latent demand for sharing and for connections was unbelievable.

It's hard to imagine at this time because you see Facebook and there's a billion users, but at the time we used to say, oh my God a stadium full of people came in and signed up for Yahoo! Mail today. And then two stadiums full. It went from like oh my God, 1000 people came to 10,000. Like when was it going to stop? To 20,000 to 50,000 to 100,000 to 1 million every day. Literally, in a day a million people would come. And we've seen that again and again. And we will see that again and again. I still maintain that there will be services that pop up that within a week will have a billion people signed up because that's how quickly-- that's how tightly connected we are and that's how quickly ideas and memes spread around the world today.

Hancock: Let's go back to that moment where you're talking about you're paving the road. You're creating it, and you don't know where it's going. What were some of the hardest challenges that you were facing at that time? And you sort of have the advantage of 20/20 hindsight, but what were some of the technical, organizational, business connecting market, challenges at that time that you think were sort of essential for you to figure out that then enabled you to move forward?

Ralston: I guess the challenges were legion, right? It's interesting. Everything was hard. Everything. It's hard to choose where to start. And at the same time, we were incredibly lucky. It wasn't that hard because we were in the middle of this secular growth that was in a sense inevitable. And so, all you had to do was grab on and manage not to let go and just to hold on as long as you can-- until something happened because it was going to happen.

But it wasn't so easy to hang on. And it wasn't so easy to hang on for a lot of reasons. There was a lot of skepticism, a lot of skepticism. Even as the bubble expanded, no one believed, and it was like come on really? These things aren't that valuable. Like come on. Whatever. Yahoo! you're like a directory, how can this be valuable? You're giving everything away for free. You can't make that up in volume. It's not going to work.

At Y Combinator and Imagine K12, the two accelerators we work at now, people ask us all of the time, "Well, where should I base my company?" And we say, wherever you want -- where your network is the strongest. However, we'll say this, one of the real advantages of starting a company in Silicon Valley is when you go to the supermarket and you bump into someone and you say, oh Hi!

And they say, "What are you doing now?" And you say, you're doing a startup you're more likely to get, "Oh my gosh. That's cool." Or, "So am I. It's awesome." Whereas if you go to certain parts of the world where it's not perceived in the same light and you bump into someone and you say you're at a startup they say, "Why did you do that? Shouldn't you get a job?" And that wears on you. And in the early days what we were doing, we ran into all of the time. What are you doing? Why don't you work on something real? Enterprise software. Databases. Boxes. Hardware. At least that's what NetApp was doing, right. Now, NetApp was catching a very different wave, in a sense, and I think it's interesting.

Hancock: Let's talk about that, your perspective on NetApp.

Ralston: Let me come back to that because I think the idea of challenges-- because that led to another major challenge, which is the fundamental challenge for every startup -- that you run into first and have to solve, which is how do you get talent? So how do you describe to people when you're at a little puny company like Four11 and say, "We're creating the directory of the future." Or, "We're creating a Web-based email product." And how do you convince great engineers to come-- great anyone to come and work for you? It's hard. You have to spin a tale and they have to buy into it. And so luckily, I think there were other people like me who had read science fiction and believed in the possibility we were able to articulate -- sort of a dream of what this meant once you connect all of these things. But that was hard.

And then financing. Before the bubble especially, venture capitalists were leery of financing these things. Now, of course, that changed very quickly. Venture capitalists are nothing else but not quite clearheaded in seeing when opportunity strikes, and they saw it. They saw it hard and some of them made a lot of money because of that. Others didn't when the bubble burst and it was painful. But it was difficult in those very, very early years. It was difficult.

We were talking about Gordon Moore-- Moore's Law was in effect in so many ways as this revolution happened. Computers were getting cheaper and cheaper and cheaper every year. And your ability to serve pages very cheaply, and to store data more and more cheaply, was so obvious that we kind of missed it. It was so obvious.

There is this extraordinary change in the way we think about how we store data, and how we access data, and what that data is for. NetApp was sort of right at the very, very beginning of where there was this explosion and the need to store data, to access data. They were at the beginning-- like before people talked about big data that was what they were doing. They said, you need access to all of your data all of the time. I think the evolution of that is that that's now plumbing and now that we have all of that data, and we think we can keep on storing it, because we went from a single disk being a gigabyte to a terabyte, which is like a thousand fold in not that much time, right. Our first filer, by the way, I believe was 50 gigabytes. And the first filer we-- that was the first one we started testing with. The first filer we launched with was an F540 maybe. It might have been. I might have the number wrong. But it was 200 gigabytes. Not 200 gigabytes useable but because there was a bunch of OS stuff. It was probably more like 170 or 180 or something like that.

There are users of Yahoo! Mail today who use more data than that for a single filer. And by the way, our goal was, I can't remember exactly how close we got to this goal, but our goal was to have 500,000 users be able to use that first filer. So they were at the very beginning of this explosion, this incredible change in the way that data is thought about and stored. Those are the days when like you'd throw stuff away. Who cares? You'd just throw it away, right.

And then fundamentally accessing your data reliably, rapidly, became fundamental to almost everything we do now, whether it's your banking data, when you log into a bank or your order history on Amazon. Or their data which is the products and the product changes that they showed you ever day became so fundamental and pretty clearly NetApp hit that right on and they grew incredibly quickly because of that. In fact it's interesting, if you look at the growth of NetApp revenues and Yahoo! revenues I think they almost exactly parallel one another including in 2000 they came down together, too. Interesting.

Hancock: So that partnership or that relationship between NetApp and Yahoo! is interesting. You have these two rapidly growing companies, meteoric growth, doubling every year. And the people at NetApp kind of credit partners for helping evolve the business model.

Ralston: Oh yeah, and it hurt. It hurt a lot.

Hancock: From your perspective, how did that relationship work from your perspective? How did it evolve? What sort of real product or business implications came out of that relationship?

Ralston: We had a pretty tight relationship with NetApp in a variety of ways. I think we had three connections into NetApp. Dan Warmenhoven was my ex-boss at HP. Dave Hitz was a friend of Dave Nakiyama, and I'm pretty sure Larry Drebes one of the co-founders of Four11 knew one of the engineers, one of the key engineers on the NetApp team. So we were buddies. We knew NetApp. And when we thought about our architecture it was pretty clear to us—NFS.

The thing about NFS you have to understand is it pushes any worries about storage away because you know how to access files. It just makes it simple. So it's a nice shortcut. It's transparent and it's easy. And we didn't really want to worry about that. That's sort of the idea, right, is that if you think about like you create an API where we'll take that. You guys deal with the mail part because the storage part isn't really part of the mail part.

So, for better or worse Google took a very different approach with Gmail. Again, for better or worse. But we said we could focus on is-- we were a tiny company at the time. We were like 25 to 30 people, and mail itself is a fairly complex beast to build. So we needed to be able to take storage off the table and that's what we hoped NetApp would do, and that philosophy continued. Yahoo! started a relationship with NetApp for an equally critical part of the infrastructure which was our user database, where we stored the registration information for each user who joined Yahoo! And that grew but, of course, mail in the end dominated the relationship.

And I like to think that the pain we went through was good for both companies. But the pain we went through was that no one could have imagined how mail would grow and the performance demand that it would put on the storage device and the growth demands that it would put on the device. And frankly, filers were not ready for that. We weren't ready for it for a variety of reasons, but the good news is we had a tight relationship, and so it became tense because there were many, many sleepless nights where filers would break. And when a filer broke it was-- we all carried pagers all of the time.

At the time I was running a group in engineering at Yahoo! which had Mail amongst other things. And when Mail broke, when users couldn't get at their email, we were up all night. And we were waiting for the NetApp to essentially rebuild its file system to put new disks in, if that was a problem. If there was a bug, try to recover. The almost amazing news-- I talked to some of the guys -- we can't really remember ever losing data because of NetApp. We actually lost data way in the early days because of engineering mistakes, but almost never because of NetApp. But we lost a lot of time and suffered a lot of pain, and so did NetApp.

And we had some hard meetings where we're like either this device works for us or we have to go off separately. To Dan's credit, we had meetings with Dan -- like the CEO. Come on in and sit down and say, "Okay, this is bad. I'm tired and your shit doesn't work, and it's got to." And we had had a long relationship at the time that this was happening. These are sort of in the first couple of years of the growth of Yahoo! Mail as it started to go like this, and then like this and then all of a sudden we were ordering filers once a week. It began you'd order them once every two or three months. And pretty soon it was essentially a daily deployment of a new filer. And I think that forced NetApp to rethink a lot about how it built its filers, how it QA'd them, how it supported them, how the entire infrastructure worked. If they were going to be-- when a customer is driving you to the limits you have to ask yourself okay is this typical? Do we need to build a device that responds to this? Is this the space we're in? And they very clearly said, "Yes."

So, we need to be able to take our filers from 99 percent reliability to 99.9 and 99.99 percent reliability. How do we build an organization that can do that? I think they had to completely rethink how they did that. And within a couple of years we were a huge customer. We went from being an important customer to one of the two biggest customers, because it worked and there was a lot of innovation that had to do with-- I didn't see all of it clearly but business processes inside of NetApp.

Because if you think about what you have to do to get that kind of reliability it's actually hard. You actually have to rethink how you build and how you ship product to actually make that work. You have to think about how problems that come in, how that gets connected to engineering, and how engineering responds and how you think about when you're building something what kind of tolerances you have, when you test something what kind of tolerances do you have? How do you get it so that it can run for a year with no problem? And my perspective is that that was incredibly valuable to NetApp because once you get there, then a whole series of applications opens to you that wouldn't have been open otherwise.

And actually could have reduced the size of the company or even brought it down much more quickly. And I think that was-- I don't think NetApp competition, Auspex and the rest, had that sort of advantage.

Hancock: It's so interesting to hear this relationship in those early days. I'd like to pull back and move up to 10,000 foot level and think about Yahoo! and NetApp in Silicon Valley in sort of the larger ecosystem. And would you describe what were the attributes of Silicon Valley at that time, this rapid growth time? And then fast forward now to your role in Y Combinator and other as we're seeing a different kind of ecosystem building different kinds of companies. Can you lay that perspective out over the evolution of the last few years?

Ralston: Yeah. I've been thinking about that. I fear that ... You know, we were kind of in the middle of a tornado. It's that little term for when you're in a crazy growth business. And it's hard to see outside of a tornado. And interestingly for me I kind of left the entrepreneurial world I'd been in the middle of for about four years from '93 to '97 when we got bought by Yahoo! And because it was this amazing thing, amazing place to be, I stayed there and didn't get back into startups for another nine years. So I don't know that I have the best perspective on that. But I will say that in an interesting way, the time right around sort of the craziness, when John Doerr talked about the greatest creation of wealth ever seen, because there was this incredible disruptive change and there were so many people saw it.

All of a sudden, there was this change in behavior, where people went away from comfortable, long term jobs at places like HP, or Goldman Sachs or McKinsey, whether it was business school or engineering school and came and started companies. Because it was clearly the beginning of something so big. And in a sense for a little while Silicon Valley where... It used to be what you would do is you would work for four to ten years, get your experience and then go start your company. And submit a business plan to the guys on Sand Hill Road, the guys in suits over on Sand Hill Road. And you'd get dressed up in a suit and you'd go and you'd pitch your thing. And if you had good experience and good track record, usually selling enterprise software, because that's mostly what software was. There wasn't that much of financing of consumer products, games and the like, a little bit but not that much. Then you might get a first round of funding in exchange for 40 percent of your company, et cetera, et cetera, et cetera.

There was this fusion of founders, younger and younger, that kind of stopped for a while. Silicon Valley kind of exploded and then kind of contracted in the bursting of the bubble in 2000. And it sort of went back a few years and getting a real job in 2000 started to look a lot more attractive to people. But I think there was an unstoppable change. There was something fundamental that had changed in how you think about your career and how you think about the opportunities that are available to you. And I think that's what's sort of coming back in the valley now. And in some ways it doesn't feel like a bubble to me now. I don't think that's what's going on. I think there's a new perspective in how younger and younger people see the opportunities for them and their careers and the value in being in a younger company, an earlier company, whether they started themselves or they worked for a startup or worked for a newer company. And the impact they can have, the fun they can have, the education they can get.

It's interesting. For me it's kind of there was this big burst of that and it really sort of changed the way everyone looked at how they do things. Venture capitalists were funding things on a meeting, "Oh, wow, you're the globe. You're making websites. That must be good, here's \$1 million or \$2 million go." And then what do you know they went IPO and you made a lot of money if you go out of the stock fast. But it changed the way they thought of the world and the opportunities in front of them, they being investors. And it changed the way a lot of kids, who aren't kids now, those people now are the people who are bigger companies. But new kids coming in sort of see that, again, and see this really long term opportunity in change and disruption that has continued.

There's still immense change and disruption happening all of the time and that creates opportunities for new innovation and new startups. So I don't know if that answers your question, but for me what's more interesting is the parallel between the change that happened then and the more I think durable, sustainable changes happening now in the world of startups.

Hancock: I really like the way that you're sort of drawing on your experience and pulling out parallels. If you were to distill down you do this on a daily basis as you're working with young entrepreneurs ... lessons from history, whether it's from your personal experience, companies you've worked for in the past where there have been patterns, or whatever. If we were to distill some of those lessons that you think are important to pull from past companies or from past environment what are some of those that you would feel are vital and important for current entrepreneurs to understand for going forward?

Ralston: I think what's hard for me is to distinguish between lessons I've learned and working with the several hundred companies I've worked with over the last four years versus lessons from back then.

Hancock: Well, from any time period, it can be recent or...

Ralston: One thing I always tell kids when I talk to them, go to talk to a school, is I say don't do what I did. I went and worked for the man for ten years before I started my first company. And I say, are you kidding me, do you know from 22 to your early 30s is your most creative, energetic, productive time in your life, unencumbered time generally in your life. Why would you spend all of that energy, that ability to learn and grow working for the man? Do it for yourself or do it in a place where you can have maximum impact. Have maximum impact on yourself and on the world and do it now. And I say, that's what you should do now.

That doesn't mean you have to start a company, because there's a corollary, and I learned this hard when I did this, is that startups are so difficult. I talk to guys all of the time, you know, who are in this-- whether you've managed to raise a lot of money and then you're trying to get customers, you're trying to figure out how you're ever going to survive still. Or you're just trying to raise money, or you're just trying to get your act together, and find some kind of customer for the product that you think is going to be awesome, but you haven't quite gotten there. It's hard at every step. And there's this immense crazy rollercoaster that you're on.

And that's not necessarily for everybody, especially as the founder. Some of that is-- you're a little insulted from that if you're an employee, but it's still there, but it's really hard. And it takes sort of a grim determination to survive that.

But Paul Graham has this thing that he famously said, which I kind of paraphrased earlier, that back then we were in this immense period of change and growth. One of the things I always point out to people when the bubble burst in 2000. The thing that didn't change was the number of people using services like Yahoo! was offering. It wasn't like, you know, stock prices went like that. Use of Yahoo! Mail like that. It didn't change. That fundamental shift that we talked earlier about, that inflection point, was never going to stop. So all you had to do was grip on tight and hold on. And Paul has said, "You have two choices when you have a startup. You can either quit or get rich." And that's just a shorthand to say that if you survive, if you can keep on going, you'll find success. Big if! Because it's hard, because sometimes you can't or sometimes you run out of gas.

Although it's not always clear when you run out of gas, right. When exactly are you done in startup? I used to like to ask people so when does a startup die? When does a startup die? Well, one answer a lot of people give is when you run out of money! But that's not true! There's lots of occasions -- I mean Elon Musk was \$400 million in debt with Tesla and he had little hope of ever getting it back. That's not out of money. That's like well beyond even imagining that and that didn't die. <laughs> There are so many cases of people who struggle on. Companies die when founders give up. That's really the only answer I can ever think of.

Hancock: Geoff, I wish we had all morning to talk. Thank you so much. I really appreciated your insights into your own experience, what lead you to critical decisions as well as your perspective on the evolution about the technology and this place.

Ralston: Thanks, Marguerite. It's been fun.

END OF INTERVIEW