REMARKS AT THE DIGITAL BROADBAND MIGRATION: THE DYNAMICS OF DISRUPTIVE INNOVATION

INTERNET SPECULATIONS

VINT CERF*

Let me start out by warning you that this is an entirely unrehearsed speech and you probably can tell from the beginning that this must be the case. Second, what I hope to do today is provide no answers but provoke a whole lot more debate during this panel discussion.

So, this is where it starts, four nodes at the University of California, Los Angeles, experimenting with packet switching way back in 1969.[1] And then, we get here, more or less. That's a picture 10 years ago of what the inside of the Internet looked like with all the interconnections of many thousands of independent Internet service providers interconnected.[2] This view is generated automatically by looking at the Internet's global routing tables. Each color is a different operator and of course, as you know, this is all a big collaboration. And then of course this is where we're going. Hopefully over the next several decades we will be able to extend the Internet operations across the solar system. It's not a joke; there is actually stuff going on at JPL where we're looking at trying to build a rich network for space exploration.

Let me just remind you of a couple things: the growth of the number of machines on the Internet used to be easily tracked when most of the machines were visible, but now they're not. A lot of them are hiding behind firewalls. Many devices that are on the Internet are just episodically connected like laptops and desktops and increasingly mobiles. So, this number of nearly 800,000,000 machines is the lower

---

* Vint Cerf is a Vice President and the Chief Internet Evangelist at Google, Inc. Mr. Cerf holds a B.S. in Computer Science from Stanford University and an M.S. and a Ph.D. in Computer Science from UCLA. These remarks were given on February 13, 2011, at the University of Colorado Law School in Boulder, Colorado at the Silicon Flatirons Digital Broadband Migration Conference. The powerpoint slides that accompanied Mr. Cerf's speech are *available at* http://www.4shared.com/document/OGHLRYvH/FlatIrons_Feb_2011_Cerf.html.

1. Vint Cerf, *Internet Speculations*, Silicon Flatirons Conference, Slide 2 (Feb. 2011), *available at* http://www.4shared.com/document/OGHLRYvH/FlatIrons_Feb_2011_Cerf.html.

2. *See* Map of the Internet, THE INTERNET MAPPING PROJECT (June 28, 1999), http://www.cheswick.com/ches/map/gallery/isp-ss.gif (last visited Dec. 11, 2011).

limit of the number of devices that are on the Internet; the real numbers are much bigger.[3] The number of visible machines, publicly visible, with domain names and fixed addresses is on the order of 780 million; the number of users on the Internet as of the middle of last year was slightly under two billion; the number of mobiles in operation, not on the Internet necessarily but in operation, exceeds five billion and about 15 to 20 percent of those are Internet enabled.[4] So, that will be the biggest wave of Internet capable devices in this decade.

Here's where the users are, and I think it's notable that the statistics show that Asia is the largest population of users, even though the penetration rate is only 20 percent.[5] It used to be that North America was at the top of the list, but we aren't anymore; we will never catch up! We don't intend to reproduce at the rate that would be necessary to have as many people as there are in Asia, so, we should be attentive to the fact that what we do in the United States, as important as it has been in the past, is not necessarily going to be entirely definitive. Many, many users will be from different cultures, speaking different languages, and come from different traditions, and that will influence the Internet's evolution.

One thing that's very important is that we have run out of the original IP version 4 (IPv4) address space and we have officially allocated—or I should say IANA, the Internet Assigned Numbers Authority—has allocated the last of the IPv4 addresses, to the regional Internet Registries.[6] So it's time to shift into IPv6, where there is a 128 bit address space or 340 trillion, trillion, trillion addresses.[7] I used to say it was enough for every electron to have its own webpage until somebody pointed out that there were 10 to the 88th electrons in the universe and I was off by 50 orders of magnitude, so I don't say that anymore. It will all be enough to last until after I'm dead; then it's somebody else's problem.

What is important, though, is that we have to implement it. A lot of the software is already there. Most of the laptops and desktops and routers have the software, the servers have the software, but the Internet service providers have not turned it on. They have various excuses for this, one of which is "nobody's asking for it" and you know I would say that the average user doesn't care whether they are running IPv4 or IPv6.

---

3. *See* The ISC Domain Survey, INTERNET SYSTEMS CONSORTIUM, http://www.isc.org/solutions/survey (last visited Dec. 11, 2011).

4. *See id.*; *see also* World Internet Users and Population Stats, INTERNET WORLD STATS, http://www.internetworldstats.com/stats.htm (last visited Dec. 11, 2011).

5. For more recent statistics, *see id.*

6. Cerf, *supra* note 1, at Slide 8.

7. S. Deering & R. Hinden, *Request for Comments: Internet Protocol Version 6 (IPv6) Specification*, NETWORK WORKING GROUP, Dec. 1998, *available at* http://www.faqs.org/rfcs/rfc2460.html.

They don't even know what that is. What they know is domain names. So that's not a good excuse. We really need to get this implemented and uniformly tested. There's going to be a major demonstration planned on the 8th of June called World Internet IPv6 Day. We at Google and many others are participating in that. It's a complicated problem. The old protocol and the new one are not directly inter-operable. So, you have to run both of them in parallel and eventually there will be a time when you can only get a v6 address and not a v4 address, in which case you won't be able to talk directly to something speaking IPv4.[8]

There are new domain names in the Internet that have been allocated in character sets other than Latin; that is another innovation that has occurred in the last couple of years.[9] I know that's an eye chart for folks way in the back, but the point here is that we will be seeing top-level domains that are not written in Latin characters, and that is another evolution of the Internet. There are more that have been assigned, and before anybody asks what are these funny little characters here,[10] that's because my machine, when I made the slide, didn't have the Sinhala fonts available to put up so it put them up as little squares. My guess is that you may experience the same thing when you're surfing the Internet and you come to a website that has a font and a character set that you didn't happen to have in your own machines. So we need to be prepared for the kind of variations that will look different than they have in the past.

There are lots and lots of issues associated with security in the Internet,[11] and the way people approach using the Internet. Many people are quite concerned, for good reason, whether it's an individual user or businesses or government. There are lots of bad things that can happen on the Internet. You can see a long list of things that have been done or actions that are being taken in order to mitigate at least some of the risk factors and some of the vulnerabilities associated with the Internet. Some people will say "why didn't you put security in when you designed it in the first place, you idiot?" And of course, the answer in part is that the technologies that are in fact used now, or could be used now, were classified then.[12] And although I knew about those and worked with NSA

8. The IPv4 address is 32 bits long and there is no way to express a 128 bit IPv6 address in a field that is only 32 bits in length.

9. Cerf, *supra* note 1, at Slides 9, 10; *see also* John Yunker, *A Whole New Way of Looking at the World (Wide Web)*, GLOBAL BY DESIGN (June 17, 2007), http://www.globalbydesign.com/blog/2007/06/17/a-whole-new-way-of-looking-at-the-world-wide-web/.

10. There is a row of little squares on the graphic where non-Latin characters should be. Cerf, *supra* note 1, at Slide 10.

11. *Id.* at Slide 11.

12. Notably, packet mode cryptography was developed by Bolt Beranek and Newman in cooperation with NSA in the mid-1970s.

on a secured version of the Internet, I wasn't allowed to share that information with anyone who didn't have a security clearance, and that meant most of the people at the universities who were participating in the design and implementation and testing of the Internet. But the other thing that I want to point out here is that if we had started out to build a fully secured system, all of the operational aspects of that probably would have inhibited the Internet's spread and use. So, that's not really an excuse, but it's an observation that some things like this don't start out in fully secured mode.

It's fair to say that if we were doing this all over again, certainly if I were involved, I would want to introduce security mechanisms from the get-go in the core of the network, including routers that could verify that the other router they are talking to is what they expect it to be, or users could verify that they are talking to the right website, or users could identify themselves to each other if they chose to do so in some very strong way. So you're seeing lots of mechanisms which I don't think I need to go through point-by-point. Many of you, I'm sure, are very much involved in implementing some of these ideas and may be creating new ones. I would draw your attention to the value and importance of making the Internet a safer place.

We hear a lot about cyber security, I'd like to suggest that we put another meme into the picture which is called "cyber safety." I think a lot of individuals think of the Internet as a place that's not necessarily safe to carry out transactions. That's not a good thing because until people feel comfortable we won't be able to grow e-commerce, we won't be able to use this as an engine for economic growth, and I think that's important to all of us. I'm anticipating that there will be many more devices on the Internet than there have been in the past and I confess to complete surprise when people put up things like refrigerators, and picture frames, and things that look like telephones, but they're really voice over IP gadgets. Many of you have seen this slide before; I am quite amazed at the guy that built an Internet-enabled surfboard. I guess he's Dutch, and he must've been out on the water waiting for the next wave thinking, "oh, if I had a laptop on my surfboard, I could be surfing the Internet while I'm waiting for the next wave," so he put a laptop in the surfboard and Wi-Fi service at the rescue shack, and now he sells this as a product.[13] So, if you have a desire to get out there and surf the Internet while you're on the water, this is the product for you.

There are things happening in the Internet today, trends that I think are very clear. One of which is an Internet of things, larger numbers of devices that are on the Internet. The second thing is sensor networks that

13. *See Intel's Surfboard*, TRAVELIZMO, http://www.travelizmo.com/archives/000325.html (last visited Dec. 11, 2011).

are increasingly becoming part of the Internet community of devices.[14] I have a sensor net at home.[15] Every room in the house that has a sensor that is picking up temperature, humidity, and light levels every five minutes, and that data is being stored in a server down in the basement. This is an IPv6, radio-based network. Each device is a sensor, and it's also a store-and-forward relay. So this is a mesh network which maintains itself. It is a commercial product made by a company called Arch Rock, which was acquired recently by Cisco. So the system has been remarkably robust. Each sensor runs on batteries, and for the hell of it, I ran the sensors for over a year on two little AA cells until they got down to 2.4 V at which point they finally pooped out. Certainly longer than the guys who built it expected it would continue to work.

One of the rooms in the house is a wine cellar, and as many of you may know, I'm a big fan of accumulating and consuming wines. I don't buy them to look at; I buy them to drink. So I'm very concerned that they stay at 60°F or lower. Humidity is another issue, it has to stay above about 30 or 40% humidity so the corks don't dry out. Every 5 minutes, the sampling takes place and if the temperature goes up above 60°F, I get an SMS on my mobile. This happened to me when I was walking into Argonne National Laboratory last year for a three-day visit. Just as I was walking into the building, the mobile goes off: it's the wine cellar calling, "You've just broken through the 60°F barrier." Every 5 minutes for the next three days I got a little message saying, "Your wine is warming up." Unfortunately, my wife was away on a two-week vacation and so she was not home to reset the cooling system. So I called the Arch Rock guys and said "do you make remote actuators" and they said "yes." So this sounds like a weekend project. It's obvious, though, that I'm going to need very strong authentication software, or the 15-year-old nextdoor is going to reprogram the house while I'm away. Again, security is important.

I thought, "You know, there's more that you can infer from looking at sensor net data." This is a privacy issue. If you can observe sensor data from every room in the house; if you can observe the lights in every room of the house every so often; if you can get information about the temperature in the garage, you can tell when the car has been driven because the temperature in the garage goes up in a noticeable way. Observing this type of data, you could tell a lot about the diurnal patterns of the people that live in a house. I think the emphasis here should be on the inferential nature of privacy; that information that looks like it

14. Gregory T. Huang, *Casting the Wireless Sensor Net*, TECHNOLOGY REVIEW, July/Aug. 2003, at 50, *available at* http://www.cens.ucla.edu/News/TechReview.pdf.

15. This graphic depicts the Cerf residential sensor network. Cerf, *supra* note 1, at Slide 13.

wouldn't be specifically of concern can turn out to be more sensitive because of what you can infer from it.

I thought about the light level thing, and I thought, "You know, I could actually tell if somebody's gone into the wine cellar and turned on the lights while I'm away." And I thought some more about that and somebody else in the discussions said, "Yeah, but you don't know what they did in there." So, I thought, "Okay maybe I should put an RFID chip on each bottle, and then I can do an instantaneous inventory, and that will tell me if any wine has left the wine cellar without my permission." This got debugged by another friend who said, "Yeah, but you can go into the wine cellar and drink the wine and leave the bottle!" So, we're going to have to put sensors in the cork, and as long as we're going to do that, we might as well sample the esters to see if the wine is ready to drink. So, before you open the bottle, you interrogate the cork. If that's the bottle that got up to 90°F when the cooling system failed, that's the bottle that you give to somebody who doesn't know the difference! It's a very practical thing to do.

These sensor networks are going to be part of our normal environment. I think those LEED buildings, the ones that are low-energy rated, will probably be heavily instrumented. They will know a lot about their internal conditions. They will know about whether people are present. That kind of information will be very important to us to improve our use of energy.

Now I want to switch gears for a minute. This is very speculative, and I would like very much to ask, if we can, during the panel discussion, explore some of this. Up until now, this is the kind of infrastructure that we have seen evolve around the Internet, and it's pretty clear that there are quite a wide range of different ways in which people have provided broadband Internet access.[16] This is a picture of New York City in 1889—it tells you something about the large number of telephone companies that were competing with each other, not necessarily interconnected either.[17] We have this history of infrastructure which is not necessarily coherent, and we're no better off now than we were 100 some odd years ago.

There is a variety of different ways in which broadband access can be provided, but an awful lot of the business models have tended to rely very heavily on more than just paying for the cost of access.[18] So, you see combined kinds of business models. In the case of local exchange carriers, they started off as telephone companies, and the business that they were in was telephone voice calls that extended into facsimile

16. Cerf, *supra* note 1, at Slide 15.
17. *Id.* at Slide 16.
18. *Id.* at Slide 17.

machines, and now it has expanded into DSL, and in some cases, fiber and Internet access. But the business models have been very combined, and of course recently wireless services have turned out to be a major component of the business models supporting the local exchange carriers. Cable carriers started out being television delivery businesses; they would either generate or, more typically, acquire access to content and then make it available to large numbers of subscribers. They would package this up into various groups of channels with various kinds of content and charge for that. So, the recovery of the cost of building the infrastructure was based on the product, mainly the entertainment, and not necessarily on simply recovering the cost of installing and operating the infrastructure.

The more recent extension of those companies into Internet access once again raises questions about the business model, and whether or not they could make a business out of just supplying Internet service, or whether you need to have all of the other components of revenue in order to support the enterprise. There are lots of wireless service providers, and they typically operate on a rather small physical range, and most of the time their models are based solely on recovering costs of access to the Internet.[19] So, they are much more affected by, for example, the cost of access to broadband service by wireline providers whose reach they extend by providing wireless access. Satellite communications started out, of course, as broadcast television but after a while also evolved from either point-to-point dedicated links or broadcast television, into Internet-based services as well.

But, the real economics here, about which I know only a little (I'm just an engineer, I am not an economist, and I'm not even a very good business guy), are complicated and you really have to pick apart all the various pieces of the enterprise to understand capital, operating expense, and requirements for interconnection. So, one of the experiments that Google is engaged in right now is figuring out whether or not it's possible to build a broadband infrastructure which will support itself solely on the basis of recovering the cost of capital and operation, and not necessarily requiring any additional revenues from value added services. Whether or not that works, we don't know. Our general sense is that it's worth finding out, so there's an effort going to build broadband infrastructure. The first pass at that is in the Stanford University area, around where the professors live, where gigabit fibers are intended to be installed and used to provide services to them.[20] So part of the panel

19. *Id.* at Slide 18.

20. *See* William Quade, *Imran Nazar's JavaScript Based Game Boy Emulator*, ZLSTUDIOS (Nov. 10, 2010), http://zlstudios.net/2010/11/imran-nazars-javascript-based-game-boy-emulator/.

discussion might be about that.

The original model of the Internet was that we would supply everyone with the details of how to build a piece of it and then it was your job not only to build a piece but then find somebody to connect to. The Internet actually grew in an organic way as a consequence of that. It was very much intentionally designed that way. It wasn't top-down control, there was no single authority that determined who could connect and who could not, who could build and who could not—it was completely open and the consequence of that was a great many participants showed up.

Some of you will recall that prior to broadband Internet access, there was only dial-up access. I can recall in the 1990s, early parts of the 1990s, there were on the order of 8,000 dial up Internet service providers. This had a very interesting effect because you could change Internet service provider by dialing a different number and, so, switching was really easy. In the current days when you're expecting to get broadband access, switching is harder because it often requires switching of physical infrastructure, installation of new equipment, maybe a truck to come out and reconfigure things. So there is a switching cost, which is substantial compared to what it once was—not to suggest that anyone would ever go back to dial-up right now. I've never seen anybody back away from higher bandwidth after they've had access to it.

Now, I wonder, I mentioned the speculative question: can you build the infrastructure to pay for itself, and we're about to try and find that out at Google. I am going to channel Bob Frankston for a minute.[21] He has often speculated that if we could buy our own pieces of Internet as users, and then hook it up to the rest of the world, that maybe we could afford to simply install the fiber, and own it and use it, in whatever capacity it can provide. So maybe one question is whether that is, in fact, a feasible thing to do and I don't know the answer to that. It probably would be hard for a community to do trenching unless it's at the time the housing is being built in the first place. But, it's conceivable if there are poles available, that you could rent poles, and then there are arguments over who gets to charge what for the poles.

We had an experience at Google a few years ago when we were invited to provide wireless Internet service in the city of San Francisco.[22] We partnered with Earthlink and worked out what we thought was a

21. Bob Frankston is the co-inventor of VisiCalc (the first spreadsheet application for personal computers).

22. Verne Kopytoff & Ryan Kim, *Google Offers S.F. Wi-Fi –For Free/ Company's Bid is One of Many in Response to Mayor's Call for Universal Online Access*, SFGATE.COM (Oct. 1, 2005), http://articles.sfgate.com/2005-10-01/news/17393876_1_wi-fi-network-free-wireless-internet-access-google.

reasonable model: we were going to provide a free service with a certain amount of capacity, and Earthlink was going to augment that with a service that had some charges associated with it. What I discovered is that after we "won," so to speak, all we had "won" was the right to negotiate with over 29 different jurisdictions over how much they were going to charge us for access to the light poles on which the wireless capability would be placed, and in the end I think it didn't work out. I don't believe we built that. We did build a wireless system in Mountain View as a test case, mostly to figure out what we were getting into. Google is very strong on data to make decisions, not intuition. We built the wireless system in Mountain View just to see what it would cost, and what problems would arise, but we didn't get to do the one in San Francisco.

The other thing which is an obvious possibility, which has been tried, is municipal networks. Those of you who are interested in the *Sturm und Drang* of legal rights and things like that would be interested to know that some competing carriers have objected to municipal networks building their own systems, and have even attempted to have legislation passed that would inhibit a municipality from building its own net.[23] The argument, as I understood it, goes something like, "Gee, a municipality is a governmental agency, it's the government competing with the private sector." I always wondered about that argument because it seemed to me that what would likely happen is that the citizens would agree to tax themselves for the capital cost of installing this wireless, or possibly wired, system. Since they agreed to do that, they would probably contract with some private sector entity to go build and operate it. So, in a sense, it would still be the private sector, it just might not be a local exchange carrier that would get the business. So maybe that's why they objected to the possibility of someone else getting the business than themselves. I don't know if I've offended anyone in the room, but if I have, it was intentional.

So, we have models where this has worked out well, and we have models where this has not worked out well. One of the interesting questions is to understand more deeply, why not? Why has this not worked? And I think it would be important for us, as we look at policy, to try and understand what the dynamics are here, not just the politics but more the economics and physics of it all. There's a long list of policy considerations that could affect how the Internet evolves.[24] This is clearly not a complete list, but it's intended once again to suggest that during the panel discussions that we might take a moment to consider some of these

23. Anthony Sciarra, *Municipal Broadband: The Rush to Legislate*, 17 ALB. L.J. SCI. & TECH. 233 (2007).

24. Cerf, *supra* note 1, at Slide 19.

things.

One of them, of course, we have been talking about is broadband infrastructure. Another issue has gone by the term "net neutrality," which has become such a distorted debate that the term almost doesn't mean anything anymore. But the point, at least as I see it, has been to assure that consumers are not unduly constrained in their ability to reach anywhere on the Internet to get a product or a service that they are interested in. And, that business deals that are advanced by the provider of the underlying broadband access should not inhibit the consumer from making that choice. There are all kinds of other arguments having to do with what the consumer does with the access to the Internet, whether it overloads the network. I fully accept the argument that you need to manage the resources of the network to provide fair access to all the consumers, but, at the same time, I don't think there should be discrimination with regard to where you get to go, even if there's some discrimination as to how much you get to use and when.

I think there are also some major issues having to do with safe harbor, and we're seeing increasing problems where governments will say, "Well I don't want this provider of service to allow any of the traffic to leave my geographical authority." So, in Europe you might be concerned that you don't want any data that your citizens are generating to show up in any country other than the country of residence. Often, there have been safe harbor rules that allow actors not in the country to satisfy the requirement even though they have an international footprint. Here, I think we should pay attention to that because that's an important facilitator of international commerce. Again, I'm not going to go through all of these things.

I do want to emphasize this notion of "permissionless innovation." In the Internet environment, the applications on the Internet have grown very organically. Part of the reason for this is that the Internet wasn't designed to do anything in particular. When I teach engineering courses, one of the things that I emphasize is that if you design something too carefully to be too precisely attuned to a particular application, it may not work for anything else. The Internet on the other hand wasn't designed to do anything except move packets from point A to point B with some probability greater than zero and that's all that we ask and the rest of it is what happens at the top, at the edges. So, I think this has led to what we have been calling "permissionless innovation" because if you want to try something out, you just do it. The Yahoo! guys and the Google guys and the Skype guys didn't ask permission to build their products and services; they just put them up on the Internet and let people come and use them. If they were successful, more people would use them; and, if they weren't, we would never hear of them at all.

Let me skip down to multi-lateral legal frameworks. There are many

abuses that occur on the Internet, some of them are tort-like in nature, some of them are criminal. Some of them are socially abusive and, in many cases, we look at these as a society and say these are not good things. The problem is that we don't have the same legal frameworks from country to country, and even from jurisdiction to jurisdiction, about these things. And until we begin to address common views of what is socially acceptable and what is not, it's going to be very hard to maintain any kind of international, let me use the word, "discipline," in a network environment. Because if you are operating in country A, and having a negative impact on someone in country B, unless there is some reciprocity between the legal systems, there may be no way to reach the party who is causing harm. Once again, it's going to take some serious thinking and multilateral discussion. And we may end up at the beginning, with maybe small sets of countries agreeing on things, or maybe a broad set of countries on not very many things. But unless we begin, we won't be able to deal with some of the social and economic side effects of negative things that happen on the Internet.

On a more constructive side, we also could agree to the meaning of certain things that would enhance electronic commerce, for example: strong authentication mechanisms, the procedures by which certificates are granted for identification, the particular technologies that are used, so that you can strongly authenticate. I want to emphasize that this is not an argument in favor of losing anonymity on the Internet. A lot of people will argue, very persuasively, that anonymity leads to bad behavior, and therefore, we should do away with it, and you should be forcibly required to identify yourself when you are on the Internet. I resist this. I do think we should have really good tools for identifying ourselves when we wish to be identified in order to carry out a transaction. But, we should have the freedom to say, "No, I don't want to identify myself." The other party, in this case, could say, "Then this conversation is over," and that's perfectly okay. So, strong tools, but not necessarily forcibly applied, strike me as being important. And that is as much as I am going to say this morning. I am ten minutes over time, so I am going to stop here. I hope I have stimulated some thinking.