The internet at 40
Pioneers: Vint Cerf

Vinton Cerf is vice president and chief internet evangelist for Google. He has been involved in internet technology right from the very start. Ten years ago, the Open University interviewed him for a series called TheWebStory.com

His involvement with internet technology started in the 1960s, when he was working on a computer networking project funded by the US Defense Advanced Research Project Agency, ARPA. He is widely known as the 'father of the internet'.

Well I think people have gotten a little carried away at least in our culture there seems to be a need to focus a lot of attention on just one or two people. That's not right. There are many fathers of the Internet depending on how far back you want to go and the evolution of the technology. At UCLA I was a graduate student along with Steve Parker, John Postel, Bob Braden, and a number of others, all working for Professor Len Kleinrock, who had the network measurement centre at UCLA. And so our job was to put the first computer up on the ARPANET and particularly the one that did network measurement.

The ARPANET was a network of computers set up in 1969 to link research departments in universities around the United States. The big challenge was to get the different computers to talk to each other.

No one was really in charge of the development of the host protocols to connect computers up to the underlying network that Bolt Beranek and Newman was building, and so the graduate students just sort of gravitated to try to work on that. We always expected that someone would come out who was professional and would run the show. Steve Parker, who was my good friend then and still is, ran the network working group and we always expected that somebody from the East Coast would show up to tell us what to do but they never did. So we just went on and did the best we could.

Protocol is of course it's a diplomatic term. It's something that you establish in agreement between countries and that's called protocol. It's also turns out to come from the Greek word 'protocollum', which was the table of contents of a scroll. Well we stole that word for computer communication conventions because packets of information that computers exchange have little headers on them to say where they're going, where they came from and how much there is in that piece. So we called the procedures computer protocols. They're simply conventions for communications between computers.

To get information between the various computers that formed the ARPANET, the data was chopped up into small 'packets'. The packet switching protocols Vint and his colleagues devised enabled the packets to be sent by different routes and recombined at the other end to recreate the original data.

When the first wide area packet switch network was being built the ARPANET there was some uncertainty whether it was going to work at all. That actually worked out quite well. It was a very powerful and useful tool for computer science departments that were part of that system. In fact packet switching was so successful that we at ARPA anyway decided it would explore using packet switching in radio and satellite communication. That led to the development of the mobile packet radio network and an Atlantic satellite net that linked the US to Europe using packet switching technology. Well once those projects were under way Bob Kahn who was at ARPA at the time realised that these networks would ultimately have to be inter connected to each other. And we didn't have any protocols, no procedures, no conventions that would allow computers that were on different networks to smoothly intercommunicate with each other. That led to the inter net project which Bob started at ARPA

around 1970 – late ’72 or early ’73 and he posed that problem to me when I was at Stanford in March of ’73 and we worked together on solving that problem: how could computers on different networks communicate with each other uniformly. And that led to the design of what is now called TCP/IP.

This really was a back of the envelope moment. I was sitting in San Francisco in a hotel lobby waiting for some session to start at a conference and had an envelope in my pocket and I pulled it out and I was just sketching what the implications were of the architecture that Bob and I had talked about, eventually leading to what we called gateways and are today called routers. And so in a sense the system’s basic architecture was forced on us because we weren't allowed to change any of the networks themselves we had to work outside of them and then figure out a way to achieve uniformity. So that little sketch, which is long since lost, I had no idea it was an important sketch at the time, it was just getting my thoughts in order, was the beginning of at least for me of understanding how the structure would work.

TCP/IP was one of the great technological breakthroughs of the twentieth century. It allowed the internet to become what it is today.

The easiest way to understand how the Internet works is to think of Internet packets as electronic postcards. Just like postcards: they have a 'to' address and a 'from' address and a finite amount of content on them. And the fact that they're electronic means that they go through the system about a hundred million times faster than the postcard that goes through the post office. But they behave just like postcards. They don’t necessarily arrive in the same order they were sent. They might not even arrive on the same day. Some of them get lost. That’s true of the Internet packets as well as postcards and so if you think about Internet packets as postcards you have a pretty good model. To understand TCP you need a little bit more thinking. Suppose that you were sending a novel to someone and the only way you could send it was by sending postcards so you cut the pages of the novel up, put them on postcards and then you realise that you have to number the postcards in order to let the party at the other end put them in the right order. Then you wonder you know if some of them got lost you'd have to re-transmit them so you keep copies to send. Then you realise that you need to find out whether you need to send any copies and you have acknowledgements coming back in the form of post cards, some of which might get lost. And so you have time outs that say if I haven't heard anything I'll start sending copies. It's basically the way the TCP works. It essentially allows us to send novels in sequenced order on top of postcards except of course we do it electronically and much faster.

So the internet is essentially a giant game of pass the packet, using a set of rules or protocols called TCP/IP. When the protocol was first used to create an inter network from three separate networks, it was a milestone in internet history.

Interestingly enough there are two big milestones, neither one of which were very noted at the time by anybody except those of us deeply involved. In 1977, late in the year we actually got all three, packet radio, sat net and ARPANET networks to function together using the Internet protocols and gateways in between and that was very exciting for a few of us who were a part of that, but not noted anywhere. We didn’t call a press conference or issue a press release or anything. We just breathed a great sigh of relief. In 1983 in January we actually insisted on the deployment of those protocols by all the computers that were part of the ARPANET and satellite net and packet radio net. And that was a big moment for the people who had to get their machines up but there were only about two hundred and at most four hundred computers involved. Today there are fifty million computers on the Internet. And so any such similar kind of transformation can't be done in what's called a flash cut. It wasn’t even a flash cut in the Internet’s case. It took several months to get everybody up and running on the new protocols. But that was a forced change. Today you can't force that change.

The only real regret I have is that I didn’t argue that we should have a larger address space for the Internet than we decided on. In 1977 I picked a thirty-two byte address space which was enough to identify up to four billion things. It's now very clear that there will be many hundreds of billions of things on the Internet in the future and we should have picked a bigger

address space but at the time it was an experiment and it never dawned on me or most everyone else that we needed anything like the scale that we will ultimately require.

Although he resists being called the father of the internet, Vint Cerf is one of its greatest pioneers. When he first sketched out his ideas for the TCP/IP protocol, did he imagine that the internetworking project would turn out the way it did?

Certainly not in the form which it has ultimately materialised in. Tim Berners-Lee's World Wide Web is something that's truly phenomenal at the rate at which it has been absorbed and adopted. We knew however that we were working with very powerful technology. We knew that computing and the distributed programmes that are around the network would be very, very powerful engines. We just didn't know exactly what they were going to do at the time. Software is sort of the ultimate clay – you can make anything you want to out of software if you can figure out how to programme it and so the Internet simply underscores the possibilities by creating an endless frontier of software that sits on top of the computers in the networks that communicate. So although in detail we didn't know how this would all evolve and economics has played a big role in the evolution. Lower cost of networking, lower cost of computing making it available to many more people. But I think we knew we were working with something that was very powerful and that ultimately might make a big difference.