

Network Working Group
Request for Comments: 1607
Category: Informational

V. Cerf
Internet Society
1 April 1994

A VIEW FROM THE 21ST CENTURY

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

A NOTE TO THE READER

The letters below were discovered in September 1993 in a reverse time-capsule apparently sent from 2023. The author of this paper cannot vouch for the accuracy of the letter contents, but spectral and radiation analysis are consistent with origin later than 2020. It is not known what, if any, effect will arise if readers take actions based on the future history contained in these documents. I trust you will be particularly careful with our collective futures!

THE LETTERS

To: "Jonathan Bradel" <jbradel@astro.luna.edu>
CC: "Therese Troisema" <ttroisema@inria.fr>
From: "David Kenter" <dkenter@xob.isea.mr>
Date: September 8, 2023 08:47.01 MT
Subject: Hello from the Exobiology Lab!

Hi Jonathan!

I just wanted to let you know that I have settled in my new offices at the Exobiology Lab at the Interplanetary Space Exploration Agency's base here on Mars. The trip out was uneventful and did let me get through an awful lot of reading in preparation for my three year term here. There is an excellent library of material here at the lab and reasonable communications back home, thanks to the CommRing satellites that were put up last year here. The transfer rates are only a few terabits per second, but this is usually adequate for the most part.

We've been doing some simulation work to test various theories of bio-history on Mars and I have attached the output of one of the more interesting runs. The results are

best viewed with a model VR-95HR/OS headset with the peripheral glove adapter. I would recommend finding an outdoor location if you activate the olfactory simulator since some of the outputs are pretty rank! You'll notice that atmospheric outgassing seriously interfered with any potential complex life form development.

We tried a few runs to see what would happen if an atmospheric confinement/replenishment system had been in place, but the results are too speculative to be more than entertaining at this point. There has been some serious discussion of terra-forming options, but the economics are still very unclear, as are the time-frames for realizing any useful results.

I have also been trying out some new exercises to recover from the effects of the long trip out. I've attached a sample neuroscan clip which will give you some feeling for the kinds of gymnastics that are possible in this gravity field. My timing is still pretty lousy, but I hope it will improve with practice.

I'd appreciate it very much if you could track down the latest NanoConstructor ToolKit from MIT. I have need of some lab gear which isn't available here and which would be a lot easier to fabricate with the tool kit. The version I have is NTK-R5 (2020) and I know there has been a lot added since then.

Therese,

I wanted you to see the simulation runs, too. You may be able to coax better results from the EXAFLOP array at CERN, if you still have an account there. We're still limping along with the 50 PFLOP system that Danny Hillis donated to the agency a few years back.

The attached HD video clip shows the greenhouse efforts here to grow grapes from the cuttings that were brought out five years ago. We're still a long ways from '82 Beaucastel!

Gotta get ready for a sampling trip to Olympus Mons, so will send this off for now.

Warmest regards,

David

To: "David Kenter" <dkenter@xob.isea.mr>
CC: "Therese Troisema" <tttroisema@inria.fr>
From: "Jonathan Bradel" <jbradel@astro.luna.edu>
Date: September 10, 2023 12:30:14 LT
Subject: Re: Hello from the Exobiology Lab!

David,

Many thanks for your note and all its news and interesting data! Melanie and I are glad to know you are settled now and back at work. We've been making heavy use of the new darkside reflector telescope and, thanks to the new petabit fiber links that were introduced last year, we have very effective controls from Luna City. We've been able to run some really interesting synthetic aperture observations by linking the results from the darkside array and the Earth-orbiting telescopes, giving us an effective diameter of about 200,000 miles. I can hardly wait to see what we can make of some of the most distant Quasars with this set-up.

We had quite a scare last month when Melanie complained of a recurring vertigo. None of the usual treatments seemed to help so a molecular-level brain bioscan was done. An unexpectedly high level of localized neuro-transmitter synthesis was discovered but has now been corrected by auto-gene therapy.

As you requested, I have attached the latest NanoConstructor ToolKit from MIT. This version integrates the Knowbot control subsystem which allows the NanoSystem to be fully linked to the Internet for control, data sharing and inter-system communication. By the way, the Internet Society has negotiated a nice discount for nano-fab services if you need something more elaborate than the ISEA folks have available at XOB. I could put the NanoSystem on the Solex Mars/Luna run and have it to you pretty quickly.

Keep in touch!

Jon and Melanie

To: "David Kenter" <dkenter@xob.isea.mr>
CC: "Jonathan Bradel" <jbradel@astro.luna.edu>
CC: "Troisema" <rm1023@geosync.hyatt.com>
From: "Therese Troisema" <ttroisema@inria.fr>
Date: September 10, 2023 12:30:14 UT
Subject: Re: Hello from the Exobiology Lab!

Bon Jour, David!

I am writing to you from the Hyatt Geosync where your email was forwarded to me from INRIA. Louis and I are here vacationing for two weeks. I have some time available and will set up a simulation run on my EXAFLOP account. They have the VR-95HR/OS headsets here for entertainment purposes, but they will work fine for examining the results of the simulation.

I have been taking time to do some research on the development of the Interplanetary Internet and have found some rather interesting results. I guess this counts as a kind of paleo-networking effort, since some of the early days reach back to the 1960s. It's hard to believe that anyone even knew what a computer network was back then!

Did you know that the original work on Internet was intended for military network use? One would never guess it from the current state of affairs, but a lot of the original packet switching work on ARPANET was done under the sponsorship of something called the Advanced Research Projects Agency of the US Department of Defense back in 1968. During the 1970s, a number of packet networks were built by ARPA and others (including work by the predecessor to INRIA, IRIA, which developed a packet network called CIGALE on which the CYCLADES network operating system was built). There was also work done by the French PTT on an experimental system called RCP that later became a commercial system called TRANSPAC. Some seminal work was done in the mid-late 1960s in England at the National Physical Laboratory on a single node switch that apparently served as the first local area network! It's very hard to believe that this all happened over 50 years ago.

A radio-based network was developed in the same 1960s/early 1970s time period called ALOHANET which featured use of a randomly-shared radio channel. This idea was later realized on a coaxial cable at XEROX PARC and called Ethernet. By 1978, the Internet research effort had produced 4 versions of a set of protocols called "TCP/IP" (Transmission Control

Protocol/Internet Protocol"). These were used in conjunction with devices called gateways, back then, but which became known as "routers". The gateways connected packet networks to each other. The combination of gateways and TCP/IP software was implemented on a lot of different operating systems, especially something called UNIX. There was enough confidence in the resulting implementations that all the computers on the ARPANET and any networks linked to the ARPANET by gateways were required to switch over to use TCP/IP at the beginning of 1983. For many historians, 1983 marks the start of global Internet growth although it had its origins in the research effort started at Stanford University in 1973, ten years earlier.

I am going to read more about this and, if you are interested, I can report on what happened after 1983.

I will leave any simulation results from the EXAFLOP runs in the private access directory in the CERN TERAFLX archive. It will be accessible using the JIT-ticket I have attached, protected with your public key.

Au revoir, mon ami, Therese

To: "Troisema" <rm1023@geosync.hyatt.com>
CC: "Jonathan Bradel" <jbradel@astro.luna.edu>
CC: "Therese Troisema" <ttroisema@inria.fr>
From: "David Kenter" <dkenter@xob.isea.mr>
Date: September 10, 2023 17:26:35 MT
Subject: Internet History

Dear Therese,

I am so glad you have had a chance to take a short vacation; you and Louis work too hard! I changed the subject line to reflect the new thread this discussion seems to be leading in. It sounds as if the whole system started pretty small. How did it ever get to the size it is now?

David

To: "David Kenter" <dkenter@xob.isea.mr>
CC: "Therese Troisema" <ttroisema@inria.fr>
CC: "Troisema" <rm1023@geosync.hyatt.com>
From: "Jonathan Bradel" <jbradel@astro.luna.edu>
Date: September 11, 2023 09:45:26 LT
Subject: Re: Internet History

Hello everyone! I have been following the discussion with great interest. I seem to remember that there was an effort to connect what people thought were "super computers" back in the mid-1980's and that had something to do with the way in which the system evolved. Therese, did your research tell you anything about that?

Jon

To: "Jonathan Bradel" <jbradel@astro.luna.edu>
CC: "David Kenter" <dkenter@xob.isea.mr>
CC: "Troisema" <rm1023@geosync.hyatt.com>
From: "Therese Troisema" <ttroisema@inria.fr>
Date: September 12, 2023 16:05:02 UT
Subject: Re: Internet History

Jon,

Yes, the US National Science Foundation (NSF) set up 5 super computer centers around the US and also provided some seed funding for what they called "intermediate level" packet networks which were, in turn, connected to a national backbone network they called "NSFNET." The intermediate level nets connected the user community networks (mostly in research labs and universities at that time) to the backbone to which the super computer sites were linked. According to my notes, NSF planned to reduce funding for the various networking activities over time on the presumption that they could become self-sustaining. Many of the intermediate level networks sought to create a larger market by turning to industry, which NSF permitted. There was a rapid growth in the equipment market during the last half of the 1980s, for routers (the new name for gateways), work stations, network servers, and local area networks. The penetration of the equipment market led to a new market in commercial Internet services. Some of the intermediate networks became commercial services, joining others that were created to meet a growing demand for Internet access.

By mid-1993, the system had grown to include over 15,000 networks, world-wide, and over 2 million computers. They must have thought this was a pretty big system, back then. Actually, it was, at the time, the largest collection of networks and computers ever interconnected. Looking back from our perspective, though, this sounds like a very modest beginning, doesn't it? Nobody knew, at the time, just how many users there were, but the system was doubling annually and that attracted a lot of attention in many different quarters.

There was an interesting report produced by the US National Academy of Science about something they called

"Collaboratories" which was intended to convey the idea that people and computers could carry out various kinds of collaborative work if they had the right kinds of networks to link their computer systems and the right kinds of applications to deal with distributed applications. Of course, we take that sort of thing for granted now, but it was new and often complicated 30 years ago.

I am going to try to find out how they dealt with the problem of explosive growth.

Louis and I will be leaving shortly for a three-day excursion to the new vari-grav habitat but I will let you know what I find out about the 1990s period in Internet history when we get back.

Therese

To: "Troisema" <rm1023@geosync.hyatt.com>
CC: "David Kenter" <dkenter@xob.isea.mr>
CC: "Therese Troisema" <ttroisema@inria.fr>
From: "Jonathan Bradel" <jbradel@astro.luna.edu>
Date: September 13, 2023 10:34:05 LT
Subject: Re: Internet History

Therese,

I sent a few Knowbot programs out looking for Internet background and found an interesting archive at the Postel Historical Institute in Pacific Palisades, California. These folks have an incredible collection of old documents, some of them actually still on paper, dating as far back as 1962! This stuff gets addicting after a while.

Postel apparently edited a series of reports called "Request for Comments" or "RFC" for short. These seem to be one of the principal means by which the technology of the Internet has been documented, and also, as nearly as I can tell, a lot of its culture. The Institute also has a phenomenal archive of electronic mail going back to about 1970 (do you believe it? Email from over 50 years ago!). I don't have time to set up a really good automatic analysis of the contents, but I did leave a couple of Knowbots running to find things related to growth, scaling, and

increased capacity of the Internet.

It turns out that the technical committee called the Internet Engineering Task Force was very pre-occupied in the 1991-1994 period with the whole problem of accommodating exponential growth in the size of the Internet. They had a bunch of different options for replacing the then-existing IP layer with something that could support a larger address space. There were a lot of arguments about how soon they would run out of addresses and a lot of uncertainty about how much functionality to add on while solving the primary growth problem. Some folks thought the scaling problem was so critical that it should take priority while others thought there was still some time and that new functionality would help motivate the massive effort needed to replace the then-current version 4 IP.

As it happens, they were able to achieve multiple objectives, as we now know. They found a way to increase the space for identifying logical end-points in the system as well as increasing the address space needed to identify physical end-points. That gave them a hook on which to base the mobile, dynamic addressing capability that we now rely on so heavily in the Internet. According to the notes I have seen, they were also experimenting with new kinds of applications that required different kinds of service than the usual "best efforts" they were able to obtain from the conventional router systems.

I found an absolutely hilarious "packet video clip" in one of the archives. It's a black-and-white, 6 frame per second shot of some guy taking off his coat, shirt and tie at one of the engineering committee meetings. His T-shirt says "IP on everything" which must have been some kind of slogan for Internet expansion back then. Right at the end, some big bearded guy comes up and stuffs some paper money in the other guy's waistband. Apparently, there are quite a few other archives of the early packet video squirreled away at the PHI. I can't believe how primitive all this stuff looks. I have attached a sample for you to enjoy. They didn't have TDV back then, so you can't move the point of view around the room or anything. You just have to watch the figures move jerkily across the screen.

You can dig into this stuff if you send a Knowbot program to concierge@phi.pacpal.ca.us. This Postel character must have never thrown anything away!!

Jon

To: "Jonathan Bradel" <jbradel@astro.luna.edu>
CC: "David Kenter" <dkenter@xob.isea.mr>
CC: "Troisema" <rm1023@geosync.hyatt.com>
From: "Therese Troisema" <ttroisema@inria.fr>
Date: September 15, 2023 07:55:45 UT
Subject: Re: Internet History

Jon,

thanks for the pointer. I pulled up a lot of very useful material from PHI. You're right, they did manage to solve a lot of problems at once with the new IP. Once they got the bugs out of the prototype implementations, it spread very quickly from the transit service companies outward towards all the host computers in the system. I also discovered that they were doing research on primitive gigabit-per-second networks at that same general time. They had been relying on unbelievably slow transmission systems around 100 megabits-per-second and below. Can you imagine how long it would take to send a typical 3DV image at those glacial speeds?

According to the notes I found, a lot of the wide-area system was moved over to operate on top of something they called Asynchronous Transfer Mode Cell Switching or ATM for short. Towards the end of the decade, they managed to get end to end transfer rates on the order of a gigabyte per second which was fairly respectable, given the technology they had at the time. Of course, the telecommunications business had been turned totally upside down in the process of getting to that point.

It used to be the case that broadcast and cable television, telephone and publishing were different businesses. In some countries, television and telephone were monopolies operated by the government or operated in the private sector with government regulation. That started changing drastically as the 1990s unfolded, especially in the United States where telephone companies bought cable companies, publishers owned various communication companies and it got to be very hard to figure out just what kind of company it

was that should or could be regulated. There grew up an amazing number of competing ways to deliver information in digital form. The same company might offer a variety of information and communication services.

With regard to the Internet, it was possible to reach it through mobile digital radio, satellite, conventional wire line access (quaintly called "dial-up") using Integrated Services Digital Networking, specially-designed modems, special data services on television cable, and new fiber-based services that eventually made it even into residential settings. All the bulletin board systems got connected to the Internet and surprised everyone, including themselves, when the linkage created a new kind of publishing environment in which authors took direct responsibility for making their work accessible.

Interestingly, this didn't do away either with the need for traditional publishers, who filter and evaluate material prior to publication, nor for a continuing interest in paper and CD-ROM. As display technology got better and more portable, though, paper became much more of a specialty item. Most documents were published on-line or on high-density digital storage media. The basic publishing process retained a heavy emphasis on editorial selection, but the mechanics shifted largely in the direction of the author - with help from experts in layout and accessibility. Of course, it helped to have a universal reference numbering plan which allowed authors to register documents in permanent archives. References could be made to these from any other on-line context and the documents retrieved readily, possibly at some cost for copying rights.

By the end of the decade, "multimedia" was no longer a buzz-word but a normal way of preparing and presenting information. One unexpected angle: multimedia had been thought to be confined to presentation in visual and audible forms for human consumption, but it turned out that including computers as senders and recipients of these messages allowed them to use the digital email medium as an enabling technology for deferred, inter-computer interaction.

Just based on what I have been reading, one of the toughest technical problems was finding good standards to represent all these different modalities. Copyright questions, which had been thought to be what they called "show-stoppers," turned out to be susceptible to largely-established case

law. Abusing access to digital information was impeded in large degree by wrapping publications in software shields, but in the end, abuses were still possible and abusers were prosecuted.

On the policy side, there was a strong need to apply cryptography for authentication and for privacy. This was a big struggle for many governments, including ours here in France, where there are very strong views and laws on this subject, but ultimately, the need for commonality on a global basis outweighed many of the considerations that inhibited the use of this valuable technology.

Well, that takes us up to about 20 years ago, which still seems a far cry from our current state of technology. With over a billion computers in the system and most of the populations of information-intensive countries fully linked, some of the more technically-astute back at the turn of the millennium may have had some inkling of what was in store for the next two decades.

Therese

To: "Therese Troisema" <ttroisema@inria.fr>
CC: "Jonathan Bradel" <jbradel@astro.luna.edu>
From: "David Kenter" <dkenter@xob.isea.mr>
Date: September 17, 2023 06:43:13 MT
Subject: Re: Internet History

Therese and Jon,

This is really fascinating! I found some more material, thanks to the Internet Society, which summarizes the technical developments over the last 20 years. Apparently one of the key events was the development of all-optical transmission, switching and computing in a cost-effective way. For a long time, this technology involved rather bulky equipment - some of the early 3DV clips from 2000-2005 showed rooms full of gear required to steer beams around. A very interesting combination of fiber optics and three-dimensional electro-optical integrated circuits collapsed a lot of this to sizes more like what we are accustomed to today. Using pico- and femto- molecular fabrication methods, it has been possible to build very compact, extremely high speed computing and communication

devices.

I guess those guys at Xerox PARC who imagined that there might be hundreds of millions of computers in the world, hundreds or even thousands of them for each person, would be pleased to see how clear their vision was. The only really bad thing, as I see it, is that those guys who were trying to figure out how to deal with Internet expansion really blew it when they picked a measly 64 bit address space. I hear we are running really tight again. I wonder why they didn't have enough sense just to allocate at least 1024 bits to make sure we'd have enough room for the obvious applications we can see we want, now?

David

Final Comments

The letters end here, so we are left to speculate about many of the loose ends not tied up in this informal exchange. Obviously, our current struggles ultimately will be resolved and a very different, information-intensive world will evolve from the present. There are a great many policy, technical and economic questions that remain to be answered to guide our progress towards the environment described in part in these messages. It will be an interesting two or three decades ahead!

Security Considerations

Security issues are not discussed in this memo.

Author's Address

Vinton Cerf
President, Internet Society
12020 Sunrise Valley Drive, Suite 270
Reston, VA 22091

EMail: +1 703 648 9888

Fax: +1 703 648 9887

EMail: vcerf@isoc.org

or

Vinton Cerf
Sr. VP Data Architecture
MCI Data Services Division
2100 Reston Parkway, Room 6001
Reston, VA 22091

Phone: +1 703 715 7432

Fax: +1 703 715 7436

EMail: vinton_cerf@mcimail.com

