

The Amateur Computerist

Webpage: <http://www.ais.org/~jrh/acn/>

Spring/Summer2000

Celebrating 25th Anniversary of TCP/IP

Volume 10 No. 1

Table of Contents

Welcoming the Millennium.	Page 1
Who Can Watch the Watchdog?.....	Page 5
Internet Pioneers Panel.	Page 8
Citizens' Agenda 2000 Forum.	Page 14
Cleveland Freenet Closed.....	Page 19
From the Internet.	Page 21
Oral History of the Internet.....	Page 26
30 Years of RFCs.....	Page 33
Principles of the Internet.....	Page 51
ARPANet Mailing Lists.....	Page 61

Welcoming the New Millennium

With this issue of the *Amateur Computerist*, we want to welcome the new millennium. Such an event happens rarely and when it does, it gives one reason to pause and consider its significance and the promise it represents.

The arrival of a new millennium happens at a propitious time in the plans of the *Amateur Computerist*. The current issue was delayed several months, and now it turns out to be an appropriate way to welcome in a new era.

This issue was to be a 25th anniversary issue for celebrating the publication in May 1974 of the paper describing the philosophy and design of the protocol for an internetting of diverse networks. We are a

little late. The paper “A Protocol for Packet Network Intercommunication” by Robert Kahn and Vinton Cerf appeared in the *IEEE Transactions on Communications*. This paper marks a significant change both in the development of packet switching networks as they were developed up to its publication and in the notion of what would make possible a global, ubiquitous computer communications infrastructure for the future.

There was a challenge facing society at the time the paper was written in summer of 1973. There were no personal computers at this time. The earliest kit version for a personal computer, the Mark 8, would not be announced in the magazine *Radio Electronics* until over a year later, in September 1974. Already by the summer of 1973, there were a number of timesharing systems and much interest in creating computer networks in countries around the world. The research documenting the development of the ARPAnet had been broadly disseminated. It led to widespread interest in setting up such computer networks for diverse uses such as research purposes, commercial purposes as for banks and airlines, for educational purposes and for other uses. Already the National Physics Laboratory (NPL) in the United Kingdom was developing a packet switching network, as was Louis Pouzin in France, who was creating Cyclades. And there were commercial networks beginning like TYMNET, and soon TELENET in the U.S. Also there were plans for creating a European Informatics Network (EIN). How would people or computers on any of the growing number of packet switching networks be able to communicate with those on other networks? Recognizing the need to be able to interconnect these diverse networks, Robert E. Kahn, who was the system designer of the ARPAnet and had worked at Bolt Beranek and Newman on the early development of the ARPAnet wrote:

“If separate data networks are jointly planned before development, at least at the interconnection level, they may be connected at a later date and viewed together as a single network that evolved by way of separate networks.”

“Resource-Sharing Computer Communications Networks” in *Proceedings of the IEEE*, vol 60, no. 11, November 1972, p. 1407.

The problem to be solved was more difficult than apparent. How

would it be possible for diverse networks using different technologies, under different forms of ownership and under different administrations, to interconnect? To do so, it would be necessary to recognize and provide for this diversity. It would also be necessary to identify the generality of what the networks had in common, and how they might differ, and to be able to accommodate these differences.

That is the task that Kahn found himself exploring in early 1973. Considering the general problem, he also had the advantage of having a particular problem to solve that was related to the general problem. He had come to work at ARPA/IPTO in November 1972 after arranging a successful demonstration of the ARPAnet for over 1000 people attending the International Conference on Computer Communications in Washington DC the previous month. At ARPA/IPTO he found there was a desire to have research in the area of developing a ground packet radio (PRNET) and a packet satellite (SATNET) network. While there had been research about a single node packet radio network called AlohaNet, there was not the kind of ground packet radio networking developed that Kahn decided to create. Money had already been appropriated, Kahn explains describing the situation at IPTO in early 1973. With the general problem in mind of how to link up diverse packet switching networks, Kahn had the particular problem of connecting PRNET to the ARPAnet. Also he had in mind connecting SATNET to PRNET and to the ARPAnet. In considering the particular problem in a general way, Kahn identified a conceptual framework for an architecture to solve the problem.

He calls this conceptual framework the Open Architecture Network Environment. Briefly, Kahn recognized that diverse packet switching networks will be created by different entities, and that their interconnection could not require any internal changes in the networks. Kahn's concept was for a meta-level system that would be independent of any networking technology or operation. It would make it possible to have these networks interconnect and intercommunicate.

Developing the ground packet radio network and packet satellite network in a general way so that they could be linked peer to peer, rather than becoming embedded in one big network like the ARPAnet, Kahn clarified the architectural principles that would make a global internet

a reality.

Realizing the need to embed the new protocol into the operating system and part of the gateways of the component networks to make this internetworking possible, Kahn recognized the need to understand how to interface the protocol to diverse operating systems. Vinton Cerf, who had recently joined the faculty at Stanford, had been part of the Network Working Group (NWG) and had had experience with operating systems. While a graduate student at UCLA, Cerf had helped Kahn to test the ARPAnet. Kahn invited Cerf to work with him developing the design for the protocol. In Spring 1973, Kahn described ideas about the new protocol to Cerf and they worked together on the details for the design of the protocol during the summer of 1973. They presented their draft paper describing the new protocol at the same time as a NATO sponsored meeting of international networking researchers at the University of Sussex in Brighton, U.K. in September, 1973. The final paper appeared in the journal *IEEE Transactions on Communications* in the May 1974 issue.

The Open Architecture Network Environment welcomes diversity and makes communication possible among different networks. The goal of making resource sharing possible not only in a network, but among diverse networks a goal at the foundation of the Internet. The Internet heralds in a new era and appropriately symbolizes the promise of the new millennium.

This issue of the *Amateur Computerist* includes an excerpt from testimony given by Robert Kahn to a U.S. Congressional subcommittee describing the early development of the Internet. It also contains RFC 2555 about the early development of RFCs. RFC 2555 was issued to celebrate the 30th anniversary of the RFCs and as a tribute to the work of Jon Postel who died in Fall, 1998. The RFC includes comments by Joyce Reynolds, Steve Crocker, Vint Cerf, Jake Feinler. Also, in the issue is the article "Some Principles of the Internet" describing in more detail the early technical issues of the Internet. There is a report on the Internet pioneers panel held at the ACM SIGCOMM99 in August 1999. Also, there is a report on a conference in Finland held by the EU on how citizens can participate more in decision making by those in government. A proposal for an oral history of the Internet follows. Along with a note

about the closing of the Cleveland Freenet, the issue ends with the continued serialization of an article on the early mailing list on ARPAnet, the MsgGroup mailing list.

Who Can Watch the Watchdog? The GAO Report on ICANN is Issued

by Ronda Hauben
ronda@ais.org

On Friday, July 7, 2000 the U.S. General Accounting Office (GAO-ICANN) posted their report¹ of their investigation of ICANN.²

The GAO-ICANN report was requested by Senator Judd Gregg of the U.S. Senate Committee Appropriations in a House-Senate conference report in October, 1999. The GAO was asked to review the relationship between the U.S. Department of Commerce (DOC) and the Internet Corporation for Assigned Names and Numbers (ICANN) and report back about to the U.S. Congress about the legality of the activities of ICANN and of the DOC's activities with ICANN.

The report is an interesting report, both in what it does and what it doesn't do. One of the essential issues the GAO raises is whether the U.S. government has the authority to transfer government property or functions to a private non profit corporation. This is an important question in the view of the U.S. government plan to transfer key assets of the Internet infrastructure to a private corporation. The GAO Report notes that the Department of Commerce "states that no government functions or property have been transferred" under its agreements with ICANN.

The GAO is supposedly an oversight body over the Department of Commerce. To report what the Department says and to leave that as the official statement of the situation is a breakdown of the independent role that an oversight body must have from the subject of its investigation. The GAO claims to have reviewed government documents as part of their process. However, a crucial document that they fail to mention is

the Report³ from the Office of Inspector General of the National Science Foundation (OIG-NSF) which was issued in February 1997.

During the interview I had with the GAO, I asked if they knew of this document. They said they did. As opposed to the GAO report, the OIG-NSF report identifies the public nature of the Internet's infrastructure. It notes that because the Internet is so crucial in the daily activities of the public, government must retain a responsible role to protect that the Internet will be continually available to the public. The OIG-NSF report also cites the factual history of the public administration of the IP numbers and domain names of the Internet, and states that the "public administration of this unique public resource should continue." (OIG-NSF, p. 8)

Instead of any reference to the factual history or government obligation with regard to the Internet's infrastructure, the GAO-ICANN report claims that it is unclear whether the "transition" of the domain name system and root server system "will involve the transfer of government property to a private entity." (p. 26)

And the Report claims an inability to know whether there is "government property" involved in the transfer of the management of the DNS to ICANN. Instead the GAO states that it is not the intentions of the Department of Commerce to transfer government property to ICANN.

There is a difference between wondering if there is "government" property to be transferred from public oversight, as opposed to whether there is a public interest requiring the protection of the domain name system, IP numbers and protocol standards process.

Such a public interest is a general public interest, involving the public in the U.S. and around the world. The U.S. government considered this obligation when it was administering the domain name functions as a public administration. But with the transfer to a private entity, the need to support the public's right to access to the Internet's is being challenged by those with a commercial self interest. The importance of considering the public interest regarding the Internet's infrastructure is an issue that crosses national boundaries, while the commercial self interest and bickering which is the basis for the creation of ICANN poses a barrier to the ability of the Internet to be protected as

an international public treasure.

In this context, there is an interesting discussion in the GAO-ICANN report about whether the U.S. Department of Commerce is the appropriate entity to be involved with the administration of ICANN and of the Internet. The discussion by the GAO was limited to whether the Department of Commerce or State Department was the appropriate entity to represent the U.S. government in its role of developing ICANN.

The OIG-NSF Report, however, had recommended the creation of a new scientific and public research commission to oversee and administer the Internet's infrastructure and to provide for the scaling of the Internet's infrastructure.

Such an administration must determine how to be an international administration, as the Internet is international. The starting point for such an endeavor, however, is not a commercially based private entity like ICANN, but a public and scientific entity like the one that was proposed in the OIG-NSF report.

The GAO-ICANN Report didn't discuss this recommendation.

Also, the GAO Report failed to recognize that there are procedures for input and determining consensus that have grown up within the Internet community, such as the RFC procedure. Yet the interviewers talking with me said they were familiar with the RFC procedure.

While the GAO Report on ICANN noted the concern of many who they interviewed that there be continued U.S. government oversight over ICANN until an appropriate international mechanism for oversight can be created, the Report fears that there will be international objection to maintaining U.S. government responsibility even if no appropriate international mechanism has been created.

While the Report fails to take up the hard questions, a careful reading of it will show the serious nature of the problem that the world is faced with regarding ICANN, if a government oversight body like the GAO is not able to do the needed factual and legal investigation to determine ICANN's legitimacy.

Who will be able to counter the vested interests who have pressured for the creation and development of ICANN? That is the important question and one that the GAO Report is not able to answer unfortu-

nately.

Footnotes:

1. <http://www.gao.gov/new.items/og00033r.pdf>
 2. <http://www.icann.org/>
 3. <http://www.ais.org/~ronda/new.papers/gao-icann/oig-nsf.txt>
-

Copyright © 1996-2000 All Rights Reserved. Alle Rechte vorbehalten Verlag Heinz Heise, Hannover

Internet Pioneers Panel Discusses Challenges Facing the Internet

by Ronda Hauben
ronda@ais.org

The conference: ACM SIGCOMM99 ¹

The place: Sanders Theater, Harvard University,

The date: Tuesday, August 31, 1999.

The time: 17:15.

The moderator: Bob Metcalfe, inventor of Ethernet.

Among the panelists:

Louis Pouzin and Hubert Zimmerman who created the Cyclades packet switching network in France in the early 1970s; Larry Roberts and Len Kleinrock pioneers of the ARPAnet, the earliest packet switching network; Bob Kahn and Vint Cerf, also ARPAnet pioneers, who went on to design the TCP/IP protocol for the internetworking of diverse packet switching networks; and Paul Baran, whose research helped to pioneer the development of packet switching technology.

The occasion: the 10th anniversary of the creation of the award honoring lifelong contributions to the field of computer communications research which has had an important impact on the work of others. On the stage is a panel of those who have won the award over the past 10 years.

For his first question, Metcalfe asks what the panelists think of the new protocol IPv6 which has been created to replace IPv4, the current version of the protocol that makes the Internet possible. Surprisingly, almost all the panelists say that they see IPv6 as problematic. Vint Cerf notes that there isn't any pressure from users or vendors to make the change. He adds, however, that lots of devices are being planned which will need IP numbers and thus, justify converting to the new protocol.

Len Kleinrock asks why there has not been any attention given in the new protocol version to make future changes easier. Sandy Fraser, another of the panelists, points to expenditures by big business which he believes will lead to a large installed base making it expensive to make future changes.

Bob Kahn reminds the other panelists about how the earliest version of the TCP/IP protocol only anticipated that there would be a few networks as part of the Internet. They thought they would never need addresses for more than ½ dozen or dozen national networks. But soon the maverick invention of the ethernet spawned local area networks, changing the landscape. In hindsight, Kahn notes, what should be the future is clear, but when going forward, it is hard to see ahead. "We are going to get it wrong again," he warns, if there isn't adequate thought put into what will be needed. Instead of going by the principle "ready, aim, fire," when tackling such problems, he observes, there's the tendency to "ready, fire, aim" or "fire, aim, ready." He suggests the need to get ready first and then to take aim at a problem, in order to be able to recognize whether or not the correct problem has been identified.

Kahn proposes that the problem which led to the creation of IPv6 may have been looked at in a way that is wrong or that there was a need for a different approach to the problem. He notes that this is an example of where the research community has not done a good job of thinking through the future and what is needed.

The discussion moves on to what the problem is that causes long delays for some in accessing the world wide web. The panelists consider whether there are technical causes of the delays which the appropriate research efforts could identify.

Changing the focus of the discussion, Metcalfe asks Louis Pouzin what his view is of the creation of ICANN. He asks Pouzin if ICANN

will blow the whole thing apart.

Pouzin responds, that since he lives in France, he is not sure what the issues are in the United States, but that you can not give up on the concerns of people around the world. The task of assigning unique IP numbers is not a real problem, Pouzin explains. But there has been a warp in handling it at the international level. Pouzin asks: What is wanted? Is there a desire for a situation whereby a few years from now because of ICANN a number of countries will be up in arms and decide after all they could just as well organize their own Internet? That it may be problematic, but that they can handle it among themselves.

Pouzin explains that the ITU is in charge of allocating virtual international resources of communication. That such sensitive issues must be handled by an international committee. There is no other way as they have the experience and the relationship and the habits of diplomacy.

Vint Cerf comments that he can't believe he is hearing Pouzin say that the ITU is better. Cerf disagrees that the ITU would be appropriate to solve the problems.

Pouzin admits that they are difficult problems but that this is the way to handle such difficult problems because there are so many items that are national obligations that in the end, an international body is how it is done.

Cerf again disagrees, noting that ICANN has had a difficult birth but believes that ICANN is needed because industry is the only game in town with regard to who can oversee the Internet names and numbers.

Another of the panelists, David Farber, interjects his view that ICANN got stuck in quick sand. It should have gotten an interim board to get the actual board, and only then taken on the difficult issues. He advises, first set up the infrastructure and then take on the functions. Farber disagrees that ITU's process would be appropriate for the Internet.

In response to a question from Metcalfe about some of the technical obstacles to the further development of the Internet, Len Kleinrock points to feature shock, or the problem for the user of absorbing new interfaces that require hours and hours of new learning.

Opposing the tendency toward the "dumbing down" of the network,

Kahn describes how the potential of the Internet will be lost if people can not themselves interact with the Net. He proposes that there should be ways to help people learn to become programmers so that every citizen would be able to get the Internet to do what he or she wants. Kahn also suggests that speech understanding research could help by making it possible to give computers verbal rather than typed instructions.

Further discussion on the desirability of voice activation leads Kleinrock to warn against voice activated agents. He worries that the agents will misunderstand the task because of the difficulty of human precision in understanding what the computer has understood and then in being able to control the computer.

Farber refers to the fear of people that the computing environment will expose everything a person does online to the observation of government. Others on the panel recognize this as a problem but also that there is the problem of corporations using the Internet to gather information on people. Another problem presented is the need to gather data to diagnose networking problems.

Metcalf asks if the Internet architecture will continue to scale making it possible for many more networks and computers and people to be connected. Sandy Fraser warns that the vision and the architecture that has made such scaling possible is being eroded. He wonders where the leadership will come from to continue to sustain the Catenet concept, the concept of a diversity of networks being able to interconnect and communicate. Also Fraser urges the need to reestablish the importance of basic concepts that are at the foundation of the Internet, such as the datagram.²

Another panelist, Paul Green, reminds the audience of the two cultures concept introduced by the British writer C. P. Snow. Technology, Green proposes, can be used for enslavement or empowerment, and there is a need for understanding and exchange between those involved with the liberal arts and the sciences. He points to the one sided portrayal of technology in George Orwell's book "1984." Instead of such a frightening scenario, the diffusion of communication can bring international peace and help solve the problems of the disadvantaged, Green adds.

Kahn questions how one can determine the issue of the ability to scale the Internet from current knowledge. He compares this to trying to predict the ability of the world economy to scale. Both are hard to predict, he contends, because we don't know what will be invented.

Metcalf asks if there are any silver bullets that will solve current problems. Kahn replies that it is much easier to build something reliable than to debug problems. He explains that it is crucial to have records of what happens on the Internet to be able to solve the technical problems. And though this may fly in the face of concerns with privacy, it is crucial to come to grips with this problem. It isn't possible to keep the Internet reliable without keeping certain kinds of records. He explains that the telephone system had found a way to monitor the workings of the system and so has been able to solve this problem.

The big expense that AT&T had incurred in buying a cable company is presented by Sandy Fraser as the motivation for the company to encourage customers to buy as much video and audio as possible to pay for AT&T's investment.

Others like Dave Clark and Bob Kahn raise the importance of exploring what users will want to do using the Internet, rather than deciding for users what they will do.

As his final challenge to the panel, Metcalfe asks what would be the most interesting questions to pose to a graduate student contemplating research in the data communications field.

Kleinrock proposes studying the field of nomadic computing. Also he suggests exploring issues such as: If one gives up all privacy, how much security could one achieve on the Internet? That the relationship between security and privacy requires study. Kahn proposes exploring the relationship of design theory with engineering practice. That there is the need to do good design and to have a way to have measurable results to compare with the theory. Paul Green urges students not to take a micro problem, but to work on something that they would be proud of, and to make it count.

The panel ends after two and ½ hours. The acoustics in the Harvard building known as Memorial Hall where the panelists had been seated were poor, making it often difficult for the panelist to hear each other or for some in the audience to hear. Also most had been through a tiring

day before the panel gathered at 17:15.

Despite the difficulties, however, something had been achieved. Some of the panelists challenged the fads in Internet research and development, urging that the problems need more effort to be understood. Several of the panelists freely disagreed with each other, yet often did so without any hostility or animosity. This led to a discussion where different views were presented so that the issues could be explored in a broader way than often happens in technical conversations. Also the issues examined were for the most part either social, or the discussion of the technical issues included social concerns or considerations. This, too, was quite different from the narrow technical discussion that is often proposed as the model for technical issues.

The panel discussion helped to present a view of the field of data communication that contributed to the foundation of the Internet and to its early development. Including discussion of social concerns as a part of the discussion of the field, helps to establish the fact that the user is part of the data network and the needs and interests and concerns of the user are an area to be included in the field of research and study. This then presents a glimpse into the future when the user and the interests of the user interacting with the hardware and software are recognized as a vital part of the Internet. Instead of viewing the user as customers or as victims of commercial firms vying for market share, users will be viewed as citizens of an online collaborative and participatory networking society, or more simply as netizens.

The panel did not, however, grapple with the most important issues of the continued development of the Internet. Such issues have to do with the way that, at least in the United States, there is an effort that the academic and government and other public or educational forms of Internet development be subsumed within a commercial sphere. Any broader vision of the user as netizen or of the need to connect all users would be ceded to industry who only view users as customers. J. C. R. Licklider who promoted much of the early vision for the development of computer networking, maintained that network access must be seen as a right, rather than as a privilege. This view required that all the population have the ability to have access to the developing computer network.³ And that the network be interactive, encouraging the users to

participate online and in developing it into something that would help meet the needs and desires of its users. This vision had the user participating in creating the ever developing vision for the future of the Internet. That is the challenge that users need to take up, taking the torch from the pioneers and carrying it forward.

1. The event was the opening session of Sigcomm'99, sponsored by the Special Interest Group (SIG) on Data Communication of the Association of Computer Machinery (ACM). The conference was held from Tuesday, August 31, 1999 through Friday, September 3, 1999. See: <http://www.acm.org/sigcomm/sigcomm99>

2. The datagram was one of the early conceptual and technical advances which made it possible to have an internet. A datagram is a packet containing source and destination information in addition to the data being transported. It doesn't contain information about the path for reaching the destination.

3. See for example "The Computer as a Communication Device," by J. C. R. Licklider and Robert Taylor, "Science and Technology: for the Technical Men in Management" 76 (April 1968): 21-31. Reprinted in "In Memoriam: J. C. R. Licklider: 1915-1990," 21-41, Palo Alto, Calif. Digital Systems Research Center, 1990. See: <http://memex.org/lick.html> and <http://www.columbia.edu/~hauben/netbook>

Summary of Seminar E2 Civic Participation, Virtual Democracy and the Net Citizens' Agenda NGO Forum 2000 (3-5 Dec., Tampere, Finland)

[Editor's Note: One of our editors, Ronda Hauben, was an invited participant in the Citizens' Agenda NGO 2000 Forum. The following is a summary of the session she participated in.]

The Internet provides citizens a channel where it is rather cheap and

fast to discuss and have at least some kind of impact on the society. The hard part is that you can talk and write as much as you want but does it mean that anyone listens? Is there real interaction or just monologues on the net?

According to the theme seminar Civic Participation, Virtual Democracy and the Net there are quite many people already on the net trying to participate and yet at least as many people making the decisions elsewhere. Most of the time nothing happens between these two groups! The decision makers are eager to refer to the silent majority while making decisions and not to the active participants.

Is it then worth sending e-mails and publishing websites? Can one encourage civic participation and create an active net community? Obviously, just providing the tools is not enough.

There is need for suitable attractive applications, training and political willingness. Intersectorial co-working has an important meaning, too, since together people are stronger. The same goes with NGOs.

Finally, one must state that the Internet is a great tool for building places where citizens can raise their voice and point out important issues. It also offers a new opportunity to mobilize people.

However, the big I is still just a tool. It is the user who makes the difference!

Since there is still a lot to do and quite a range of issues to talk about, the participants of this theme seminar decided to continue the discussion. If you want to join the group, send an e-mail to Mr. David Smith, wfa@hospitalitywales.demon.co.uk.

Presentation 1: Net participation: What can the City offer?

Jari Seppä

10 years experience of work as a news reporter in local newspapers and national tv news, 12 years Head of Information of the City of Tampere, Finland. He has acted as the chairman for two committees founded by the Association of Finnish Local Authorities, one creating the good practice for municipal information and the other one guidelines for municipal services presented over the Internet.

Mr. Seppä introduced some practical examples how the city of

Tampere (www.tampere.fi) has developed civic participation via the Internet. The city has for example a service where citizens can ask anything about the municipality and an official will reply to him or her as soon as possible. The residents have had also an opportunity to participate in financial planning by giving their comments on the budget for the year 2000.

In short, Seppälä explained how the Internet enables plan presentation, dialogue and lobbying, combined into the visual and functional opportunities provided by new media. According to him the full utilization of the electronic services is still held back by limited access to the Internet and lack of computer skills both among the residents and city employees in Tampere.

Presentation 2: Citizen forums, virtual publicness and practices of local democracy:

Lasse Peltonen & Seija Ridell

Researchers in University of Tampere, Department of Regional Studies and Environmental Policy, Journalism and Mass communication. The case of Tampere-foorumi (Lasse Peltonen)

Tampere-foorumi on the net (Seija Ridell)

The willingness of political and other powerful (local) actors to participate in open and equal dialogue with citizens and grassroots civic groups is a prerequisite for virtual democracy (at the local level). The refusal of the powerful to interact prevents the utilization of the ICTs for democratic purposes. However, people must try to overcome the obstacles. One way is to build sites like Tampere-foorumi (<http://mansefoorumi.uta.fi/>). It aims to support citizens possibilities to take initiative and contribute to the local government. It also provides continuity in civic discussions.

Presentation 3: Information Technology and the Possibility for the Production of New Democratic Ethos: the Philippine Case

Myrna J. Alejo

Research Associate, Democracy Watch Program

Institute for Popular Democracy (a research institute serving social movement groups and non-government organizations in the Philippines)

and overseas. IPD conducts policy studies and discourse analysis of the factors and issues that promote or retard democracy and development in the Philippines.)

Lecturer, Department of Political Science, De La Salle University

In the Philippines, the Internet account costs about 50 U.S.D per month, and the computers are expensive as well. The minimum wage is on the average 5 dollars a day. Who has the access to the Internet with this kind of salary? Not too many. Fortunately, there are web cafes where one can connect to the net for less than a dollar per hour. That makes it a bit easier.

According to Alejo only 1% of the population in the Philippines uses the Internet. These young urban professionals represent the upper middle class and they are the only ones who can reach out for the new ideas the Internet is full of. The country itself has no funds to allocate and there is also lack of institutional co-operation. What about the NGOs? Alejo says that over 68% of the NGOs in the Philippines are connected to the net and they use the net for networking and building partnership. However, the net could be in more effective use if the organizations were more skilled. Alejo believes better policies and decisions can be achieved only by having access to the information.

Presentation 4: Is the Internet a Laboratory for Democracy? The Vision of the Netizens vs The E-Commerce Agenda

Ronda Hauben

Founding editor and writer for the *Amateur Computerist* newsletter, co-author of *Netizens: On the History and Impact of Usenet and the Internet* published by IEEE Computer Society Press.

Main issue: Why it is important for Netizens to participate in the contest being waged (as for instance: ICANN) over which strata of society will gain the benefit of the Internet and how the Internet provides the means for such participation?

The Internet can be helpful in making it possible for citizens to be able to contribute their voices to the important policy decisions governments are making about the future of the Internet. The vision of early computer pioneers is that users participate in determining the

future of the developing network.

The vision includes a commitment to explore how the Internet can make possible a new form of citizenship and of an online citizenship or netizenship. Hauben described efforts made to challenge the privatization of the Internet and its essential functions, and the lessons from this contest toward determining what the role of government and of the public should be in the decisions about the future of the Internet.

Presentation 5: Networking for democracy: the digital future?

Steven Lenos

Specialist on New Media, Public and Politics Institute for Citizenship, Participation and Politics,
Organizer of several digital debates.

Main issue: How organizations can use the Internet for international networking and how they are able to organize successful digital public debates.

With the help of the Internet one can take major steps toward interactive policymaking especially when the decision makers join the discussion.

If NGOs want to be effective, they should use almost all means available: send e-mails, publish websites, keep up news mailing lists, make phone calls etc. Face to face contacts are still valuable, too. In short, best results can be achieved by combining different kinds of media in a way that best suits the organization.

Lenos told for example about a net debate which received more publicity with the help of a regional newspaper. The paper published weekly reports about the debate and gave sort of a quality mark to the debate held on net.

Presentation 6: Net or Trap – Urban Planning on an Internet-based Neighborhood Forum

Aija Staffans

Architect, Manager of the Laboratory of Urban Planning and Design, Helsinki University of Technology, Department of Architecture. Also a

PhD student in urban planning with a topic concerning interactive processes and tools between the residents and the municipalities in city planning.

Practical experience of housing design and of applying participative methods on different development processes of old housing areas.

The main issue was whether a digital neighbourhood forum is able to bring together the municipality and local stakeholders (like inhabitants, citizen organizations, schools, kindergarten, shopkeepers etc.) in order to develop urban environment.

The City of Helsinki has a centralized organization with strong sectorised offices. Recently, the necessity of wider collaboration between citizens and the municipality has come up on several occasions. At the same time, the use of the Internet has explosively increased. Each school and library in Helsinki is on the net and a growing number of households are connected to the web. The development of a digital neighborhood forum, called the Home Street, offers new opportunities to the management of cities in the information age. The Home Street Project has developed the Internet as a participatory channel in urban processes.

Cleveland Freenet Closed on October 1, 1999

Ronda Hauben
ronda@ais.org

Long Live the Goal of Access for All of the Cleveland Freenet

The Cleveland Freenet was something very special in the history of the development of the Internet as it made access to the Internet available to all in the community.

It made access available to school children in Cleveland as I learned when I gave a talk at a conference in Cleveland in 1988. The teacher introducing me told me how her students loved being online and com-

municating with other students.

It made access available in special new forms. Unsung pioneers like Dr. William Bohl of the St. Silicon Sports Medicine Clinic on the Cleveland Freenet would respond to questions from users with sports medicine problems from the earliest days of St. Silicon Hospital till the closing of the Freenet on October 1, 1999.

Dr. Bohl would post the questions sent to him as anonymous posts and would provide a helpful response that was available for all who looked in on the clinic newsgroup. One user had an experience where an injury that more than 20 doctors in the Detroit and Ann Arbor areas of Michigan were not able to diagnosis and treat was identified by Dr. Bohl. From the e-mail the user wrote to him, he provided information about what the problem was likely to be, along with the proviso that this was general information not a particular diagnosis. Because of his online clinic it was possible to get the needed treatment to cure the injury, and then to even correspond with the doctor via e-mail in an early use of e-mail between patient and doctor. Also all who looked in on the online clinic newsgroup would be able to learn about the nature of sports medicine injuries and the varieties of their treatment from the helpful responses to individual questions posted on the newsgroup.

The Freenet made an e-mail mailbox available to each user so they could use and participate in e-mail. Shortly after I signed onto the Cleveland Freenet I had the thrill of receiving a New Year's greeting from a friend in Australia. This was January 1992.

One of the most important aspects of Cleveland Freenet was when it provided a free and helpful means for its users to explore and to post to Usenet newsgroups. After a post on Freenet I was soon receiving e-mail from numbers of people and also the posts generated interesting and sometimes prolonged discussion. It was only the fact that Cleveland Freenet provided totally free access that made it possible for me to participate in Usenet. And for years afterwards, Cleveland Freenet made it possible to have a connection to Usenet newsgroups.

When the green card lawyers wrote their infamous book advising on how to spam the Net, they advised spammers to stay away from the Freenets, warning them of the acceptable use policy of the Freenets which required responsible use from its users.

Sometime after I first got onto Cleveland Freenet, a U.S. government official from the Office of Technology Assessment (OTA) posted there requesting input on what users felt should be the role of the U.S. government in providing access to the Internet to citizens. Many people posted their responses. Several people responded that it was important that all have access, as citizens would be empowered by an ability to be online.

Again in 1994 the U.S. government, this time via the National Telecommunications Information Administration (NTIA), sponsored an online conference requesting input from users about their ideas on providing universal access to the Internet. On Cleveland Freenet this conference was carried as a local newsgroup making it easier to participate than in the mailing list form, as the volume of comments was very great.

Learning from the experience of the Cleveland Freenet, Canadian Freenets were started. The Freenet movement in Canada soon became a grassroots movement to make access available to all Canadians. Also Freenets were set up in some European countries, including Finland and Germany.

The development of the Cleveland Freenet provided a model for how the U.S. government could encourage and support a low cost means of access to the Internet for all. The U.S. government has missed this opportunity and both the U.S. government and the people of the U.S. have lost something very important.

The notion of a system of computer communications networks making e-mail and Usenet access available to all has provided an inspiring and important goal. The global communications that the Internet makes possible and affordable is a very precious treasure and a significant new development for our times. The Cleveland Freenet has provided a body of experience showing that such a goal is far from impossible. Those who recognize the importance of this goal need to redouble their efforts to make the vision of all having access to e-mail, Usenet newsgroups and a browser, a reality.

A special thank you to all who contributed to make the experience of the Cleveland Freenet such an important one in the development of the Internet.

From the Internet: Some Background*

By Robert Kahn

Editor's note: Robert Kahn is credited with being the system designer of the ARPAnet and the architect of the Internet. The following is an excerpt from the Supplemental Background information which he submitted with his testimony before the Congressional Subcommittee on Basic Research on March 31, 1998.]

In the early 1970s, DARPA was exploring radio and satellite-based packet networks along with the ARPAnet. Each network had different communication speeds, interfaces, packet sizes and internal operations. After joining DARPA, I became the principal architect of the packet radio network, a high-speed forerunner to today's CDMA cellular technology. I also assumed management responsibility for creating a packet satellite network, which was ultimately deployed on Intelsat IV and linked several European sites with a kind of "ethernet in the sky." The challenge, back then, was to connect these three different packet networks into a seamless whole whereby any computer on one of the three networks could talk to any computer connected to any one of the three networks without necessarily knowing the location of the other sites or the underlying network connectivity.

The Internet resulted from this effort to connect those three networks and their computers in such a way that other networks and computers could be easily connected in the future. At the time, there were no personal computers or workstations as we now know them. Local area networks (such as the ethernet and ring networks) were only in development within various research laboratories, but had not been deployed. By solving the network and computer connectivity problem in a generic way, we were able to ensure that new technological developments in the future could be accommodated.

The key technical contribution which enabled this "network of

networks” to be constructed was an architecture consisting of gateways (now called routers) which were placed between the networks, and a protocol, now known as TCP/IP, which was used by the computers and the routers. I collaborated with my colleague Vinton Cerf, then at Stanford University, on the development of this protocol which was presented publicly for the first time in September 1973 at a meeting in Sussex, England and published by the IEEE in May, 1974. Subsequently, I enlisted the help of BBN and University College London to work with Stanford in creating the initial implementations of the protocol (for different computers). With support from DARPA, BBN created the initial Internet gateway software for experimental use in the mid 1970s.

Until the early 1980s, the Internet was used primarily for experimental purposes. During that period, the protocols were steadily refined and tested. Other networks were connected during that period including many of the early local area networks; a few European research networks were also connected. During this period, the overall management of the Internet was handled by DARPA in the person of either myself or Dr. Cerf, who was with DARPA during the period 1976-1982. Many of the basic issues under consideration in this hearing can be traced to decisions we made during that period. However, since there were few commercial organizations participating at the time, and very little international involvement, decisions we made were largely determined on the basis of logically defensible criteria and fairly complete knowledge of all the relevant matters; fortunately, we were also in charge of the overall research program and, as a result, there was remarkably little controversy about the Internet within the research community.

One of the decisions we made during that period was to delegate responsibility for maintaining information about key Internet parameters to Jon Postel, currently a researcher at the University of Southern California (USC) Information Sciences Institute who had been carrying out similar functions for the ARPAnet. While DARPA retained the ultimate authority for decisions about policy and procedures, increasingly Jon Postel assumed primary responsibility for these functions, with DARPA retaining an oversight responsibility in the event this was necessary to

invoke. During that period, no occasion arose when there was a need to second guess his decisions (although we often would inquire as to how he came up with certain decisions). This function, performed by Jon Postel under USC's contract with DARPA, eventually became known as the Internet Assigned Numbers Authority (IANA) and included certain policy matters associated with domain names as well as IP addresses and protocol parameters. With DARPA's permission, Jon delegated certain clerical and operational functions to SRI International, while retaining other functions. Among the former were the maintenance of a database which mapped Internet names to Internet addresses and making this resource available on the Internet.

Moving ahead toward the present, the ARPAnet was phased out in 1990 and was effectively replaced by a higher-speed backbone known as NSFNET built by IBM, MCI and Merit under an award from the National Science Foundation (NSF). With encouragement and help from DARPA, NSF took over responsibility for maintaining most of the Internet management infrastructure from Defense, and recompeted the contract that the Defense Department had with SRI International. Network Solutions, Inc. (NSI) won the competition for providing the domain name registration services and has provided this service ever since, with a few exceptions, such as country codes.

When the Internet naming service known as the Domain Name Service (DNS) was first proposed in the 1980s by Paul Mockapetris (also from USC/ISI along with Jon Postel) most of the then existing sites could be characterized as educational (EDU), U.S. government (GOV & MIL) or other (this included network (NET), organization (ORG), some commercial sites that had first class research laboratories (COM) and a few special cases involving matters such as testing and multi-national experiments (ARPA and INT)). It was envisioned at the time that the overall database of names, which had previously been so small that it was trivial for a site to download the entire database from SRI daily, might become somewhat unwieldily if the number of hosts or networks increased significantly. Breaking the Internet names into categories such as EDU, COM, etc. would allow them to be managed separately and resolved into IP addresses separately, thus, affording an opportunity for efficiency and increased autonomy in the operation of

the Internet.

In addition, two letter country codes were introduced as domain names that could be managed by individual countries according to policies developed by the countries themselves. It is not necessary that all countries participate, and indeed not all have in the past. The IANA made the determination of who in a given country would be responsible for that country's domain, but gave deference to the legitimate government of the country if it chose to weigh in.

In the mid 1990s, the rapid commercial growth of the Internet was fueled in large measure by the success of the NSFNET, the introduction of many commercial Internet Service Providers, the Boucher bill which allowed NSF to open the NSFNET for commercial use (in addition to research and educational use), the continuing attraction of electronic mail and file transfer capabilities, and the subsequent introduction of the point-and-click browser for the World Wide Web. With competitive commercial service available for access to the Internet, NSF reduced its subsidy for the NSFNET and stopped subsidizing the services provided by NSI in order to put them on a pay-as-you-go-basis. NSI has continued to do an excellent job of providing such services for the Internet under a Cooperative Agreement with NSF that is currently due to expire later this year[1998]. However, with several million domain names in existence and the potential for many more in the future, the annual revenue derived from domain name registrations could easily exceed a hundred million dollars per year if the current level of fees were to be maintained. Although the fee for individual domain name registrations has been \$50 per year (it has since been announced that the fees will be reduced somewhat), many individuals and organizations have expressed strong feelings that the existing fee structure and organizational arrangements are untenable in the long term and should be rectified.

One proposed approach for domain name registration is to require the separation of service provider roles into registries and "registrars," although one party can provide both roles. In this approach, domain name registries would be placed on a not-for-profit basis, with the registrars offering competitive commercial services. I presume this need not imply that the organization running a registry must be non-profit, but only that the function must be based on cost recovery. In this model, NSI

and/or other competent organizations could provide this function. Oversight would still have to be provided from some appropriately constituted body. It is still unclear how best to introduce competition in this approach.

My view is that, in general, fewer separately managed gTLDs are better than more, but there is no obvious choice of the right number in a competitive environment unless, in principle, it can be arbitrarily large. Still, this general approach of increasing the number of gTLDs, at least as an interim approach, holds considerable appeal and almost all the parties are endorsing the principle but with considerable divergence of opinion about how to achieve it equitably and technically.

Another solution is for the U.S. Government to recompute the function, as it did for the InterNIC, according to a set of agreed principles (hopefully with broad community consensus) with a goal of enabling this function to operate in a stable and reliable fashion without direct U.S. government involvement in its operation. Others feel that this can be sorted out completely within the private sector. There would likely still be a need for an oversight role of some sort as there is for any critical societal function (even a competitive one) that cannot be allowed to fail. But even here, there is no consensus yet on what that oversight should be, who should provide it or even that it is needed.

More time is needed to reach a consensus on how best to proceed here.

*SUPPLEMENTARY BACKGROUND INFORMATION from Testimony before the Subcommittee on Basic Research of the Committee on Science on the subject of Internet Domain Names by Dr. Robert E. Kahn, President and CEO Corporation for National Research Initiatives March 31, 1998.

Notes toward an Oral History of the Internet

Considering the importance of the development of the Internet, and of the protocol suite TCP/IP that makes it possible, there are relatively few books or other forms of written historical accounts about it. The written documentation that does exist is in many cases scattered in technical literature either online in what are known as Requests for Comment (RFCs) or in journals of technical articles. And many of the RFCs from the early period of TCP/IP development are not yet readily available online.

The few currently existing accounts of this important networking development mainly focus on the earliest history of the ARPAnet, begun in 1969. The Internet, however, is a qualitatively different historical development from time-sharing and the early ARPAnet. This development grew out of work by researchers supported by ARPA's IPTO in the early 1970s. The ARPAnet, the pioneering packet switching network, was constructed along the concept of one central network that all would link up with if they wanted to be part of it. The Internet, however, grew from a different architectural concept – the concept of open architecture networking developed by Robert E. Kahn.

The concept of open architecture networking was built on a recognition that there would be diverse kinds of packet switching networks, but that they should all be able to interconnect and intercommunicate.

As a researcher at Bolt Beranek and Newman (BBN) in Cambridge, Kahn made a significant contribution to the development of the early ARPAnet. His research influenced the ARPAnet Request for Quotation (RFQ) issued by Larry Roberts of ARPA in 1968. With others at BBN, Kahn wrote BBN's proposal for the ARPAnet contract. Also he designed the IMP-host interface known as BBN Technical Report 1822. Along with Al Vezza, Kahn organized a demonstration showing the utility of a packet switching network. The demonstration took place at the International Conference on Computer Communication (ICCC72) in Washington D.C. in October, 1972, thrilling many of the participants

and convincing them that packet switching was a significant and functional new technology. As a result of the successful demonstration, several researchers from different countries met and formed the International Network Working Group (INWG) to collaborate while developing packet switching networks in their diverse countries.

In November, 1972, Kahn went to work at ARPA. He was interested in the multiple network problem of how to connect diverse packet switching networks. This problem had not been originally considered when the ARPAnet was designed. But with the growing interest in creating packet switching networks in the U.S. and abroad, this problem had become an urgent one to be solved. Becoming involved in ground packet radio network research (PRNet) and satellite packet radio network research (SATNET), when he joined ARPA/IPTO, Kahn was interested in how to internetwork these very different packet switching networks with the ARPAnet packet switching network. This was the beginning of the internetting project at ARPA and in time gave birth to the Internet.

By Spring of 1973, Kahn had identified the question that he felt had to be solved to make the interconnection of diverse packet switching networks possible: “How can I get a computer that is on a satellite net and a computer on a radio net and a computer on the ARPAnet to communicate uniformly with each other without recognizing what is going on in between?” (Hafner and Lyons, p. 223) He invited Vint Cerf, who had been part of the Network Working Group and the UCLA ARPAnet research, to collaborate with him in solving this generic problem. The two studied and struggled over the problem, finally creating a strategy and architectural design for the protocol that would solve the problem. They called the protocol the Transmission Control Program, and they presented it September 1973 at a gathering of those members of the INWG who were attending a conference at the University of Sussex, in Brighton, England. Several months later, their paper was published in the May 1974 issue of *IEEE Transactions on Communications*. The paper was titled: “A Protocol for Packet Network Intercommunication.”

A concern of researchers during this period, like Louis Pouzin who was developing the CYCLADES packet switching network in France,

was that there be a way to link up the diverse packet switching networks being developed in different countries. The development of TCP/IP would solve the problem and make possible the interconnection of a great diversity of packet switching networks into an Internet.

Research over the next ten years by many led to a series of implementations of TCP and its eventual split into TCP and IP. The internetworking protocols allowed the ARPANet to be interconnected with the satellite packet network SATNET and a mobile packet radio network. But the official adoption of TCP/IP by the U.S. Department of Defense did not occur until 1980. A cut over from the old ARPANet protocol of NCP to the internetwork protocol of TCP/IP was scheduled for January 1, 1983.

Several months after the cut over was successfully carried out, the ARPANet was split into MILNET, an operational packet switching network for the U.S. Department of Defense, and what remained of the ARPANet. The latter was continued as a research oriented packet switching network for university and other Department of Defense researcher contractors funded by ARPA.

Development work on the Internet continued during the 1983-86 period. In 1986 the National Science Foundation (NSF) began a networking project to link several supercomputer centers and to create a packet switching backbone network. By 1989, a number of ARPA Internet sites were transferred to the NSFNET. The NSFnet utilized a backbone model connecting diverse networks using TCP/IP. In 1995 the NSFNET was privatized, with the role of the U. S. government being replaced by commercial companies. Other countries and regions of the world have other forms of networking architecture. But TCP/IP makes it possible to interconnect a great variety of packet switching networks so that those on these networks can communicate with people around the world as part of an Internet.

Thus, the Internet as we know it today is the result not only of the pioneering packet switching research done on the early ARPANet and ground packet radio and satellite networks, but also of the internetworking research and development in the 1972-1987 period.

Though a few accounts have been written of the early ARPANet period, there is little public documentation of the activities of the

Internet researchers with the exception of the RFCs, journal articles, and a few articles written by networking pioneers. The one significant exception is the oral history project conducted by Dr. Arthur Norberg, along with researchers Judy O'Neill and William Aspray under the auspices of the Charles Babbage Institute. Funded by a grant from ARPA, they conducted an important set of oral histories of those working at ARPA/IPTO from its beginning in 1962 under J. C. R. Licklider to 1987, when the IPTO was ended and the research merged into another program, the Information Science and Technology Office (ISTO). The two components of the ISTO were the Basic Program and the Strategic Computing Program. In addition to the oral history interviews funded by the project, Dr. Norberg and Judy O'Neill produced two written documents. One was the report, "A History of the Information Processing Techniques Office of the Defense Advanced Research Projects Agency" (Minneapolis, Minn: Charles Babbage Institute, 1992) and the second, a book, *Transforming Computer Technology*. The focus of their study was the ARPA/IPTO contribution to the support of computer science research, and so the question of the development of the Internet received attention within that broader framework.

Other book length accounts are few and include the following: Michael Hauben and Ronda Hauben, *Netizens: On the History and Impact of Usenet and the Internet*, IEEE Computer Society Press, 1997. An online draft of the book was available via ftp in January 1994 and individual articles were posted on Usenet and available at ftp sites from 1992 on.

It contains chapters on the vision for the Net, the development of time-sharing leading up to development of the early ARPAnet, the early development of UNIX and of early Usenet. The book also contains chapters regarding the debate about the future of the Net.

<http://www.columbia.edu/~hauben/netbook>

Peter Salus, *Casting the Net: From ARPAnet to Internet and Beyond*, Addison and Wesley, 1995.

It contains quotes from RFCs of the period, some of which are not currently available online. This book describes some aspects of the

development of the ARPAnet or Internet, including opinions and views from some participants in the events of the period.

Katie Hafner and Matthew Lyon, *Where Wizards Stay Up Late: The Origins of the Internet*, Simon and Schuster, 1996.

The book presents the development of ARPA and ARPAnet research, some of the developments on MsgGroup mailing list, and a brief account of the origin of the Internet.

Stephen Segaller, *Nerds 2.0.1 : A Brief History of the Internet*, TV Books, 1998.

It has a few chapters that briefly describe the developments of the ARPAnet toward an Internet, including quotes from a number of the ARPAnet or Internet pioneers, and then focuses on the pioneers of the personal computer.

And the related book:

Arthur L. Norberg and Judy E. O'Neill, *Transforming Computer Technology: Information Processing for the Pentagon 1962-1986*, The John Hopkins University Press, 1996.

It presents the history of the Information Processing Techniques Office (IPTO) at the Advanced Research Projects Agency (ARPA) and its development of computer science, which includes support for Artificial Intelligence (AI), time-sharing, networking and graphics research.

There is a thesis on the topic. Janet Abbate, *From ARPAnet to Internet: A History of ARPA-sponsored Computer Networks, 1966-1988*. University of Pennsylvania, 1994. Abbate's thesis focuses on the 1966-1988 ARPA packet-switching network development with much emphasis on the earliest development of the ARPAnet. It provides some documentation of the SATNET and PRN developments and interconnection with the ARPAnet to create an internetwork. Her thesis mainly utilizes interviews done by researchers at the Babbage Institute and refers to a few articles in technical journals. The period of Internet development is presented as a transition to the later NSF backbone.

Abbate presents ARPAnet and Internet developments as stages in network development.

Abbate has also published a more recent book *Inventing the Internet*, MIT Press, 1999. This book describes the building of the ARPAnet and then gives some description of early research in building the Internet. She includes some discussion of the efforts to create PRNET and SATNET. Her book describes some of what has been included in the Babbage Institute interviews.

These books document the earliest development of the ARPAnet, which began in 1969. Some of the above accounts include some description or developments that were part of the period of Internet development, but are limited to a few comments from people involved at the time or to references to RFCs or technical articles.

None of the books currently available, however, provide the kind of study of the early and important events in the development of the Internet that will be helpful for those trying to understand its past so as to understand the current and future needs.

In addition to these books, there are journal articles or online articles that treat some aspect of Internet history and development. These include “A Brief History of the Internet” by Barry Leiner et al, <http://www.isoc.org/> Peter Kirstein, “Early Experiences with the ARPAnet and Internet in the United Kingdom,” in *Annals of the History of Computing*, Vol 21, No. 1, January-March 1999. Ronda Hauben, “From the ARPAnet to the Internet: A Study of the ARPAnet TCP/IP Digest and the Role of Online Communication in the Transition from the ARPAnet to the Internet,”

<http://www.ais.org/~ronda/new.papers/tcpdraft.txt>, John Adam, “Architects of the Net of Nets,” in *IEEE Spectrum*, September 1996, pp. 57-63, Vint Cerf, “How the Internet Came to Be, as told to Bernard Aboba,” in *The Online User’s Encyclopedia: Bulletin Boards and Beyond*, April 1994, pp. 527-534. There is a serious need for books and articles which document and analyze the nature of the developments that have given birth to the Internet and made it possible for it to grow and flourish. Oral history interviews with those who have contributed to this early history of the Internet, similar to those done by the Babbage Institute of those who were part of the IPTO, would help to make such

needed research and writing possible.

30 Years of RFCs

[Editor's note: The following RFC includes recollections by several Internet pioneers.]

Network Working Group

RFC Editor, et al.

Request for Comments: 2555

Category: Informational

USC/ISI

7 April 1999

Status of this Memo

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

Table of Contents

1. Introduction
2. Reflections
3. The First Pebble: Publication of RFC 1
4. RFCs – The Great Conversation
5. Reflecting on 30 years of RFCs
6. Favorite RFCs – The First 30 Years
7. Security Considerations
8. Acknowledgments
9. Authors' Addresses
10. APPENDIX - RFC 1
11. Full Copyright Statement

1. Introduction - Robert Braden

Thirty years ago today, the first Request for Comments document,

RFC 1, was published at UCLA (<ftp://ftp.isi.edu/in-notes/rfc1.txt>). This was the first of a series that currently contains more than 2500 documents on computer networking, collected, archived, and edited by Jon Postel for 28 years. Jon has left us, but this 30th anniversary tribute to the RFC series is assembled in grateful admiration for his massive contribution.

The rest of this document contains a brief recollection from the present RFC Editor Joyce K. Reynolds, followed by recollections from three pioneers: Steve Crocker who wrote RFC 1, Vint Cerf whose long-range vision continues to guide us, and Jake Feinler who played a key role in the middle years of the RFC series.

2. Reflections - Joyce K. Reynolds

A very long time ago when I was dabbling in IP network number and protocol parameter assignments with Jon Postel, gateways were still “dumb,” the Exterior Gateway Protocol (EGP) was in its infancy and TOPS-20 was in its heyday. I was aware of the Request for Comments (RFCs) document series, with Jon as the RFC Editor. I really didn’t know much of the inner workings of what the task entailed. It was Jon’s job and he quietly went about publishing documents for the ARPAnet community.

Meanwhile, Jon and I would have meetings in his office to go over our specific tasks of the day. One day, I began to notice that a pile of folders sitting to one side of his desk seemed to be growing. A few weeks later the pile had turned into two stacks of folders. I asked him what they were. Apparently, they contained documents for RFC publication. Jon was trying to keep up with the increasing quantity of submissions for RFC publication.

I mentioned to him one day that he should learn to let go of some of his work load and task it on to other people. He listened intently, but didn’t comment. The very next day, Jon wheeled a computer stand into my office which was stacked with those documents from his desk intended for RFC publication. He had a big Cheshire cat grin on his face and stated, “I’m letting go!,” and walked away.

At the top of the stack was a big red three ring notebook. Inside contained the “NLS Textbook,” which was prepared at ISI by Jon,

Lynne Sims and Linda Sato for use on ISI's TENEX and TOPS-20 systems. Upon reading its contents, I learned that the NLS system was designed to help people work with information on a computer. It included a wide range of tools, from a simple set of commands for writing, reading and printing documents to sophisticated methods for retrieving and communication information. NLS was the system Jon used to write, edit and create the RFCs. Thus, began my indoctrination to the RFC publication series.

Operating systems and computers have changed over the years, but Jon's perseverance about the consistency of the RFC style and quality of the documents remained true. Unfortunately, Jon did not live to see the 30th Anniversary of this series that he unflinchingly nurtured. Yet, the spirit of the RFC publication series continues as we approach the new millennium. Jon would be proud.

3. The First Pebble: Publication of RFC 1 - Steve Crocker

RFC 1, "Host Software," issued thirty years ago on April 7, 1969 outlined some thoughts and initial experiments. It was a modest and entirely forgettable memo, but it has significance because it was part of a broad initiative whose impact is still with us today.

At the time RFC 1 was written, the ARPAnet was still under design. Bolt, Beranek and Newman had won the all-important contract to build and operate the Interface Message Processors or "IMPs," the forerunners of the modern routers. They were each the size of a refrigerator and cost about \$100,000 in 1969 dollars.

The network was scheduled to be deployed among the research sites supported by ARPA's Information Processing Techniques Office (IPTO). The first four nodes were to be at UCLA, SRI, University of California, Santa Barbara and University of Utah. The first installation, at UCLA, was set for September 1, 1969.

Although there had been considerable planning of the topology, leased lines, modems and IMPs, there was little organization or planning regarding network applications. It was assumed the research sites would figure it out. This turned out to be a brilliant management decision at

ARPA.

Previously, in the summer of 1968, a handful of graduate students and staff members from the four sites were called together to discuss the forthcoming network. There was only a basic outline. BBN had not yet won the contract, and there was no technical specification for the network's operation. At the first meeting, we scheduled future meetings at each of the other laboratories, thus, setting the stage for today's thrice yearly movable feast. Over the next couple of years, the group grew substantially and we found ourselves with overflow crowds of fifty to a hundred people at Network Working Group meetings. Compared to modern IETF meetings all over the world with attendance in excess of 1,000 people and several dozen active working groups, the early Network Working Groups were small and tame, but they seemed large and only barely manageable at the time. One tradition that doesn't seem to have changed at all is the spirit of unrestrained participation in working group meetings.

Our initial group met a handful of times in the summer and fall of 1968 and winter 1969. Our earliest meetings were unhampered by knowledge of what the network would look like or how it would interact with the hosts. Depending on your point of view, this either allowed us or forced us to think about broader and grander topics. We recognized we would eventually have to get around to dealing with message formats and other specific details of low-level protocols, but our first thoughts focused on what applications the network might support. In our view, the 50 kilobit per second communication lines being used for the ARPAnet seemed slow, and we worried that it might be hard to provide high-quality interactive service across the network. I wish we had not been so accurate!

When BBN issued its Host-IMP specification in spring 1969, our freedom to wander over broad and grand topics ended. Before then, however, we tried to consider the most general designs and the most exciting applications. One thought that captured our imagination was the idea of downloading a small interpretative program at the beginning of a session. The downloaded program could then control the interactions and make efficient use of the narrow bandwidth between the user's local machine and the back-end system the user was interacting with. Jeff

Rulifson at SRI was the prime mover of this line of thinking, and he took a crack at designing a Decode-Encode Language (DEL) [RFC 5]. Michel Elie, visiting at UCLA from France, worked on this idea further and published Proposal for a Network Interchange Language (NIL) [RFC 51]. The emergence of Java and ActiveX in the last few years finally brings those early ideas to fruition, and we're not done yet. I think we will continue to see striking advances in combining communication and computing.

I have already suggested that the early RFCs and the associated Network Working Group laid the foundation for the Internet Engineering Task Force. Two all-important aspects of the early work deserve mention, although they're completely evident to anyone who participates in the process today. First, the technical direction we chose from the beginning was an open architecture based on multiple layers of protocol. We were frankly too scared to imagine that we could define an all-inclusive set of protocols that would serve indefinitely. We envisioned a continual process of evolution and addition, and obviously this is what's happened.

The RFCs themselves also represented a certain sense of fear. After several months of meetings, we felt obliged to write down our thoughts. We parceled out the work and wrote the initial batch of memos. In addition to participating in the technical design, I took on the administrative function of setting up a simple scheme for numbering and distributing the notes. Mindful that our group was informal, junior and unchartered, I wanted to emphasize these notes were the beginning of a dialog and not an assertion of control.

It's now been thirty years since the first RFCs were issued. At the time, I believed the notes were temporary and the entire series would die off in a year or so once the network was running. Thanks to the spectacular efforts of the entire community and the perseverance and dedication of Jon Postel, Joyce Reynolds and their crew, the humble series of Requests for Comments evolved and thrived. It became the mainstay for sharing technical designs in the Internet community and the archetype for other communities as well. Like the Sorcerer's Apprentice, we succeeded beyond our wildest dreams and our worst fears.

4. RFCs – The Great Conversation – Vint Cerf

A long time ago, in a network far, far away...

Considering the movement of planet Earth around the Sun and the Sun around the Milky Way galaxy, that first network IS far away in the relativistic sense. It takes 200 million years for the Sun to make its way around the galaxy, so thirty years is only an eye-blink on the galactic clock. But what a marvelous thirty years it has been! The RFCs document the odyssey of the ARPAnet and, later, the Internet, as its creators and Netizens explore, discover, build, re-build, argue and resolve questions of design, concepts and applications of computer networking.

It has been ultimately fascinating to watch the transformation of the RFCs themselves from their earliest, tentative dialog form to today's much more structured character. The growth of applications such as e-mail, bulletin boards and the world wide web have had much to do with that transformation, but so has the scale and impact of the Internet on our social and economic fabric. As the Internet has taken on greater economic importance, the standards documented in the RFCs have become more important and the RFCs more formal. The dialog has moved to other venues as technology has changed and the working styles have adapted.

Hiding in the history of the RFCs is the history of human institutions for achieving cooperative work. And also hiding in that history are some heroes that haven't been acknowledged. On this thirtieth anniversary, I am grateful for the opportunity to acknowledge some of them. It would be possible to fill a book with such names – mostly of the authors of the RFCs, but as this must be a brief contribution, I want to mention four of them in particular: Steve Crocker, Jon Postel, Joyce K. Reynolds and Bob Braden.

Steve Crocker is a modest man and would likely never make the observation that while the contents of RFC 1 might have been entirely forgettable, the act of writing RFC 1 was indicative of the brave and ultimately clear-visioned leadership that he brought to a journey into the unknown. There were no guides in those days – computer networking was new and few historical milestones prepared us for what lay ahead. Steve's ability to accommodate a diversity of views, to synthesize them

into coherence and, like Tom Sawyer, to persuade others that they wanted to devote their time to working on the problems that lay in the path of progress can be found in the early RFCs and in the Network Working Group meetings that Steve led.

In the later work on Internet, I did my best to emulate the framework that Steve invented: the International Network Working Group (INWG) and its INWG Notes, the Internet Working Group and its Internet Experiment Notes (IENs) were brazen knock-offs of Steve's organizational vision and style.

It is doubtful that the RFCs would be the quality body of material they are today were it not for Jonathan Postel's devotion to them from the start. Somehow, Jon knew, even thirty years ago that it might be important to document what was done and why, to say nothing of trying to capture the debate for the benefit of future networkers wondering how we'd reached some of the conclusions we did (and probably shake their heads...).

Jon was the network's Boswell, but it was his devotion to quality and his remarkable mix of technical and editing skills that permeate many of the more monumental RFCs that dealt with what we now consider the TCP/IP standards. Many bad design decisions were re-worked thanks to Jon's stubborn determination that we all get it "right" – as the editor, he simply would not let something go out that didn't meet his personal quality filter. There were times when we moaned and complained, hollered and harangued, but in the end, most of the time, Jon was right and we knew it.

Joyce K. Reynolds was at Jon's side for much of the time that Jon was the RFC editor and as has been observed, they functioned in unison like a matched pair of superconducting electrons – and superconductors they were of the RFC series. For all practical purposes, it was impossible to tell which of the two had edited any particular RFC. Joyce's passion for quality has matched Jon's and continues to this day. And she has the same subtle, puckish sense of humor that emerged at unexpected moments in Jon's stewardship. One example that affected me personally was Joyce's assignment of number 2468 to the RFC written to remember Jon. I never would have thought of that, and it was done so subtly that it didn't even ring a bell until someone sent me an e-mail asking whether

this was a coincidence. In analog to classical mystery stories, the editor did it.

Another unsung hero in the RFC saga is Bob Braden – another man whose modesty belies contributions of long-standing and monumental proportions. It is my speculation that much of the quality of the RFCs can be traced to consultations among the USC/ISI team, including Jon, Joyce and Bob among others. Of course, RFC 1122 and 1123 stand as two enormous contributions to the clarity of the Internet standards. For that task alone, Bob deserves tremendous appreciation, but he has led the End-to-End Research Group for many years out of which has come some of the most important RFCs that refine our understanding of optimal implementation of the protocols, especially TCP.

When the RFCs were first produced, they had an almost 19th century character to them – letters exchanged in public debating the merits of various design choices for protocols in the ARPAnet. As e-mail and bulletin boards emerged from the fertile fabric of the network, the far-flung participants in this historic dialog began to make increasing use of the online medium to carry out the discussion – reducing the need for documenting the debate in the RFCs and, in some respects, leaving historians somewhat impoverished in the process. RFCs slowly became conclusions rather than debates.

Jon permitted publication of items other than purely technical documents in this series. Hence one finds poetry, humor (especially the April 1 RFCs which are as funny today as they were when they were published), and reprints of valuable reference material mixed into the documents prepared by the network working groups.

In the early 1970s, the Advanced Research Projects Agency was conducting several parallel research programs into packet switching technology, after the stunning success of this idea in the ARPAnet. Among these were the Packet Radio Network, the Atlantic Packet Satellite Network and the Internet projects. These each spawned note series akin to but parallel to the RFCs. PRNET Notes, ARPA Satellite System Notes (bearing the obvious and unfortunate acronym...), Internet Experiment Notes (IENs), and so on. After the Internet protocols were mandated to be used on the ARPAnet and other DARPA-sponsored networks in January 1983 (SATNET actually converted before that),

Internet-related notes were merged into the RFC series. For a time, after the Internet project seemed destined to bear fruit, IENs were published in parallel with RFCs. A few voices, Danny Cohen's in particular (who was then at USC/ISI with Jon Postel) suggested that separate series were a mistake and that it would be a lot easier to maintain and to search a single series. Hindsight seems to have proven Danny right as the RFC series, with its dedicated editors, seems to have borne the test of time far better than its more ephemeral counterparts.

As the organizations associated with Internet continued to evolve, one sees the RFCs adapting to changed circumstances. Perhaps the most powerful influence can be seen from the evolution of the Internet Engineering Task Force from just one of several task forces whose chairpersons formed the Internet Activities Board to the dominant, global Internet Standards development organization, managed by its Internet Engineering Steering Group and operating under the auspices of the Internet Society. The process of producing "standards-track" RFCs is now far more rigorous than it once was, carries far more impact on a burgeoning industry, and has spawned its own, relatively informal "Internet Drafts" series of short-lived documents forming the working set of the IETF working groups.

The dialogue that once characterized the early RFCs has given way to thrice-annual face-to-face meetings of the IETF and enormous quantities of e-mail, as well as a growing amount of group-interactive work through chat rooms, shared white boards and even more elaborate multicast conferences. The parallelism and the increasing quantity of transient dialogue surrounding the evolution of the Internet has made the task of technology historians considerably more difficult, although one can sense a counter-balancing through the phenomenal amount of information accumulating in the World Wide Web. Even casual searches often turn up some surprising and sometimes embarrassing old memoranda – a number of which were once paper but which have been rendered into bits by some enterprising volunteer.

The RFCs, begun so tentatively thirty years ago, and persistently edited and maintained by Jon Postel and his colleagues at USC/ISI, tell a remarkable story of exploration, achievement, and dedication by a growing mass of internauts who will not sleep until the Internet truly is

for everyone. It is in that spirit that this remembrance is offered, and in particular, in memory of our much loved colleague, Jon Postel, without whose personal commitment to this archive, the story might have been vastly different and not nearly as remarkable.

5. Reflecting on 30 years of RFCs – Jake Feinler

By now we know that the first RFC was published on April 7, 1969 by Steve Crocker. It was entitled “Host Software.” The second RFC was published on April 9, 1969 by Bill Duvall of SRI International (then called Stanford Research Institute or SRI), and it too was entitled “Host Software.” RFC 2 was a response to suggestions made in RFC 1 – and so the dialog began.

Steve proposed 2 experiments in RFC 1:

“1) SRI is currently modifying their on-line retrieval system which will be the major software component of the Network Documentation Center [or The SRI NIC as it soon came to be known] so that it can be modified with Model 35 teletypes. The control of the teletypes will be written in DEL [Decode-Encode Language]. All sites will write DEL compilers and use NLS [SRI Doug Engelbart’s oNLine System] through the DEL program.”

“2) SRI will write a DEL front end for full NLS, graphics included. UCLA and UTAH will use NLS with graphics.”

RFC 2, issued 2 days later, proposed detailed procedures for connecting to the NLS documentation system across the network. Steve may think RFC 1 was an “entirely forgettable” document; however, as an information person, I beg to differ with him. The concepts presented in this first dialog were mind boggling, and eventually led to the kind of network interchange we are all using on the web today. (Fortunately, we have graduated beyond DEL and Model 35 teletypes!)

RFC 1 was, I believe, a paper document. RFC 2 was produced online via the SRI NLS system and was entered into the online SRI NLS Journal. However, it was probably mailed to each recipient via snail mail by the NIC, as e-mail and the File Transfer Protocol (FTP) had not yet been invented.

RFC 3, again by Steve Crocker, was entitled, “Documentation Conventions;” and we see that already the need for a few ground rules

was surfacing. More ground-breaking concepts were introduced in this RFC. It stated that:

“The Network Working Group (NWG) is concerned with the HOST software, the strategies for using the network, and the initial experiments with the network. Documentation of the NWG’s effort is through notes such as this. Notes may be produced at any site by anybody and included in this series.”

It goes on to say: “The content of a NWG note may be any thought, suggestion, etc. related to the Host software or other aspect of the network. Notes are encouraged to be timely rather than polished. Philosophical positions without examples or other specifics, specific suggestions or implementation techniques without introductory or background explanation, and explicit questions without any attempted answers are all acceptable. The minimum length for a NWG note is one sentence”.

“These standards (or lack of them) are stated explicitly for two reasons. First, there is a tendency to view a written statement as discussion of considerably less than authoritative ideas. Second, there is a natural hesitancy to publish something unpolished, and we hope to ease this inhibition.”

Steve asked that this RFC be sent to a distribution list consisting of:

Bob Kahn, BBN
Larry Roberts, ARPA
Steve Carr, UCLA
Jeff Rulifson, UTAH
Ron Stoughton, UCSB
Steve Crocker, UCLA

Thus, by the time the third RFC was published, many of the concepts of how to do business in this new networking environment had been established – there would be a working group of implementers (NWG) actually discussing and trying things out; ideas were to be free-wheeling; communications would be informal; documents would be deposited (online when possible) at the NIC and distributed freely to members of the working group; and anyone with something to contrib-

ute could come to the party. With this one document a swath was instantly cut through miles of red tape and pedantic process. Was this radical for the times or what! And we were only up to RFC 3!

Many more RFCs followed and the SRI NLS Journal became the bibliographic search service of the ARPAnet. It differed from other search services of the time in one important respect: when you got a “hit” searching the journal online, not only did you get a citation telling you such things as the author and title; you got an associated little string of text called a “link.” If you used a command called “jump to link,” voila! you got the full text of the document. You did not have to go to the library, or send an order off to an issuing agency to get a copy of the document, as was the custom with other search services of the time. The whole document itself was right there immediately!

Also, any document submitted to the journal could not be changed. New versions could be submitted, and these superceded old versions, but again the new versions could not be changed. Each document was given a unique identifying number, so it was easy to track. These features were useful in a fast-moving environment. Documents often went through several drafts before they were finally issued as an RFC or other official document, and being able to track versions was very useful.

The SRI NLS Journal was revolutionary for the time; however, access to it online presented several operational problems. Host computers were small and crowded, and the network was growing by leaps and bounds; so connections had to be timed out and broken to give everyone a chance at access. Also, the rest of the world was still a paper world (and there were no scanners or laser printers, folks!), so the NIC still did a brisk business sending out paper documents to requesters.

By 1972 when I became Principal Investigator for the NIC project, the ARPAnet was growing rapidly, and more and more hosts were being attached to it. Each host was required to have a technical contact known as the Technical Liaison, and most of the Liaison were also members of the NWG. Each Liaison was sent a set of documents by the NIC called “functional documents” which included the Protocol Handbook (first issued by BBN and later published by the NIC.) The content of the Protocol Handbook was made up of key RFCs and a document called “BBN 1822” which specified the Host-to-Imp protocol.

The NWG informed the NIC as to which documents should be included in the handbook; and the NIC assembled, published, and distributed the book. Alex McKenzie of BBN helped the NIC with the first version of the handbook, but soon a young fellow, newly out of grad school, named Jon Postel joined the NWG and became the NIC's contact and ARPA's spokesperson for what should be issued in the Protocol Handbook.

No one who is familiar with the RFCs can think of them without thinking of Dr. Jonathan Postel. He was "Mister RFC" to most of us. Jon worked at SRI in the seventies and had the office next to mine. We were both members of Doug Engelbart's Augmentation Research Center. Not only was Jon a brilliant computer scientist, he also cared deeply about the process of disseminating information and establishing a methodology for working in a networking environment. We often had conversations way into the wee hours talking about ways to do this "right." The network owes Jon a debt of gratitude for his dedication to the perpetuation of the RFCs. His work, along with that of his staff, the NWG, the IETF, the various NICs, and CNRI to keep this set of documents viable over the years was, and continues to be, a labor of love.

Jon left SRI in 1976 to join USC-ISI, but by that time the die was cast, and the RFCs, NWG, Liaison, and the NIC were part of the network's way of doing business. However, the SRI NLS Journal system was becoming too big for its host computer and could not handle the number of users trying to access it. E-mail and FTP had been implemented by now, so the NIC developed methodology for delivering information to users via distributed information servers across the network. A user could request an RFC by e-mail from his host computer and have it automatically delivered to his mailbox. Users could also purchase hardcopy subscriptions to the RFCs and copies of the Protocol Handbook, if they did not have network access.

The NIC worked with Jon, ARPA, DCA, NSF, other NICs, and other agencies to have secondary reference sets of RFCs easily accessible to implementers throughout the world. The RFCs were also shared freely with official standards bodies, manufacturers and vendors, other working groups, and universities. None of the RFCs were ever restricted or classified. This was no mean feat when you consider that they were

being funded by DoD during the height of the Cold War.

Many of us worked very hard in the early days to establish the RFCs as the official set of technical notes for the development of the Internet. This was not an easy job. There were suggestions for many parallel efforts and splinter groups. There were naysayers all along the way because this was a new way of doing things, and the ARPAnet was “coloring outside the lines” so to speak. Jon, as Editor-in-Chief was criticized because the RFCs were not issued by an “official” standards body, and the NIC was criticized because it was not an “official” document issuing agency. We both strived to marry the new way of doing business with the old, and fortunately were usually supported by our government sponsors, who themselves were breaking new ground.

Many RFCs were the end result of months of heated discussion and implementation. Authoring one of them was not for the faint of heart. Feelings often ran high as to what was the “right” way to go. Heated arguments sometimes ensued. Usually they were confined to substance, but sometimes they got personal. Jon would often step in and arbitrate. Eventually the NWG or the Sponsors had to say, “It’s a wrap. Issue a final RFC.” Jon, as Editor-in-Chief of the RFCs, often took merciless flak from those who wanted to continue discussing and implementing, or those whose ideas were left on the cutting room floor. Somehow he always managed to get past these controversies with style and grace and move on. We owe him and others, who served on the NWG or authored RFCs, an extreme debt of gratitude for their contributions and dedication.

At no time was the controversy worse than it was when DoD adopted TCP/IP as its official host-to-host protocols for communications networks. In March 1982, a military directive was issued by the Under Secretary of Defense, Richard DeLauer. It simply stated that the use of TCP and IP was mandatory for DoD communications networks. Bear in mind that a military directive is not something you discuss – the time for discussion is long over when one is issued. Rather a military directive is something you DO. The ARPAnet and its successor, the Defense Data Network, were military networks, so the gauntlet was down and the race was on to prove whether the new technology could do the job on a real operational network. You have no idea what chaos and controversy that

little two-page directive caused on the network. (But that's a story for another time.) However, that directive, along with RFCs 791 and 793 (IP and TCP) gave the RFCs as a group of technical documents stature and recognition throughout the world. (And yes, TCP/IP certainly did do the job!)

Jon and I were both government contractors, so of course followed the directions of our contracting officers. He was mainly under contract to ARPA, whereas the NIC was mainly under contract to DCA. BBN was another key contractor. For the most part we all worked as a team. However, there was frequent turnover in military personnel assigned to both the ARPAnet and the DDN, and we all collaborated to try to get all the new participants informed as to what was available to them when they joined the network. We also tried to foster collaboration rather than duplication of effort, when it was appropriate. The NWG (or IETF as it is now known) and the RFCs became the main vehicles for interagency collaboration as the DoD protocols began to be used on other government, academic, and commercial networks.

I left SRI and the NIC project in 1989. At that time there were about 30,000 hosts on what was becoming known as the Internet, and just over a 1000 RFCs had been issued. Today there are millions of hosts on the Internet, and we are well past the 3000 mark for RFCs. It was great fun to be a part of what turned out to be a technological revolution. It is heartwarming to see that the RFCs are still being issued by the IETF, and that they are still largely based on ideas that have been discussed and implemented; that the concepts of online working groups and distributed information servers are a way of life; that those little "links" (officially known as hypertext) have revolutionized the delivery of documents; and that the government, academia, and business are now all playing the same game for fun and profit. (Oh yes, I'm happy to see that Steve's idea for integrated text and graphics has finally come to fruition, although that work took a little longer than 2 days.)

6. Favorite RFCs – The First 30 Years – Celeste Anderson

Five years ago, Jon Postel and I had wanted to publish a 25th RFC

anniversary book, but, alas, we were both too busy working on other projects. We determined then that we should commemorate the thirtieth anniversary by collecting together thirty “RFC Editors’ Choice” RFCs based on original ideas expressed throughout the first 30 years of their existence.

Jon’s untimely death in October 1998 prevented us from completing this goal. We did, however, start to put online some of the early RFCs, including RFC 1. We weren’t sure whether we were going to try to make them look as close to the typewritten originals as possible, or to make a few adjustments and format them according to the latest RFC style. Those of you who still have your copies of RFC 1 will note the concessions we made to NROFF the online version. The hand-drawn diagrams of the early RFCs also present interesting challenges for conversion into ASCII format.

There are still opportunities to assist the RFC Editor to put many of the early RFCs online. Check the URL: <http://www.rfc-editor.org/rfc-online.html> for more information on this project.

In memory of Jon, we are compiling a book for publication next year of “Favorite RFCs – The First 30 Years.”

We have set up a web interface at: <http://www.rfc-editor.org/voterfc.html> for tabulating votes and recording the responses. We will accept e-mail as well. Please send your e-mail responses to: voterfc@isi.edu. We prefer votes accompanied by explanations for the vote choice.

We reserve the right to add to the list several RFCs that Jon Postel had already selected for the collection. Voting closes December 31, 1999.

7. Security Considerations

Security issues are not discussed in this commemorative RFC.

8. Acknowledgments

Thank you to all the authors who contributed to this RFC on short notice. Thanks also to Fred Baker and Eve Schooler who goaded us into

action. A special acknowledgment to Eitetsu Baumgardner, a student at USC, who NROFFed this document and who assisted in the formatting of RFCs 1, 54, and 62, converting hand-drawn diagrams into ASCII format.

9. Authors' Addresses

Robert Braden
USC/Information Sciences Institute
4676 Admiralty Way #1001
Marina del Rey, CA 90292
Phone: +1 310-822-1511
Fax: +1 310 823 6714
E-Mail: braden@isi.edu

Joyce K. Reynolds
USC/Information Sciences Institute
4676 Admiralty Way #1001
Marina del Rey, CA 90292
Phone: 1 310-822-1511
Fax: +1 310-823-6714
E-Mail: jkrey@isi.edu

Steve Crocker
Steve Crocker Associates, LLC
5110 Edgemoor Lane
Bethesda, MD 20814
Phone: +1 301-654-4569
Fax: +1 202-478-0458
E-Mail: crocker@mbl.edu

Vint Cerf
MCI
E-MAIL: vcerf@mci.net

Jake Feinler
SRI Network Information Center

1972-1989

E-MAIL: feinler@juno.com

Celeste Anderson

USC/Information Sciences Institute

4676 Admiralty Way #1001

Marina del Rey, CA 90292

Phone: +1 310-822-1511

Fax: +1 310-823-6714

E-MAIL: celeste@isi.edu

10. APPENDIX - RFC 1

The cover page said at the top:

“Network Working Group Request for Comments”

and then came an internal UCLA distribution list:

V. Cerf, S. Crocker, M. Elie, G. Estrin, G. Fultz, A. Gomez, D. Karas,
L. Kleinrock, J. Postel, M. Wingfield, R. Braden, and W. Kehl.

followed by an “Off Campus” distribution list:

A. Bhushan (MIT), S. Carr (Utah), G. Cole (SDC), W. English (SRI), K.
Fry (Mitre), J. Heafner (Rand), R. Kahn (BBN), L. Roberts (ARPA), P.
Rovner (MIT), and R. Stoughton (UCSB).

The following title page had:

“Network Working Group Request for Comments: 1”

at the top, and then:

HOST SOFTWARE

STEVE CROCKER

7 APRIL 1969

11. Full Copyright Statement

Copyright © The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished

to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an “AS IS” basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some Principles of the Internet

by Jay Hauben
jrh@ais.org

Introduction

Before 1973, the idea of an Internet had not yet been proposed. As of 2000, over 200 million people worldwide actively use this communications system. The Internet consists of these people and more than one million packet switching networks with very many different characteristics, interconnecting more than 100 million computers as nodes. Yet the Internet is still young. With care taken about its scientific basis and technological principles it has the potential to keep expanding for many years to come. But there are many aspects of Internet technology that

must be protected for further growth to occur. In particular, the acceptable unreliability at the internetwork level is unique and differentiable from other telecommunications network technologies such as the telephone system. Also the scaling of the Internet to meet the expected increase in demand for its use is in no way assured. To better understand this new means of communication and help protect its growth, it is worthwhile to look for the principles upon which it has been developed.

History

High speed, digital computers first appeared in the U.S. in any numbers in the early 1950s. By the end of that decade, such computers were so large and so expensive that they were operated almost exclusively for machine efficiency. That meant that few people used the computers interactively. Instead programmers and users submitted their programs on punched cards or tapes to computer centers. There, operators lined the jobs up and feed them to the computer, job after job. Sometime later the output was available for the user or programmer to retrieve and examine. This mode of use was called “batch processing.” It may have seemed efficient in terms of machine utilization. But, from the point of view of human users who waited time intervals of the order of hours or days for responses to their jobs, it was woefully inefficient. Batch processing was particularly frustrating for programmers trying to correct errors in their work.

In the early 1960s, a major improvement in efficiency for both humans and computers was achieved by the development of the “time-sharing” mode of computer operation. Taking advantage of the great processing speed of transistorized computers, this new mode allowed a set of users to simultaneously access the same computer. Computer processing time was now parceled out in very small intervals and made available to the users in round-robin fashion. Each user was offered his or her short time slots in turn so rapidly that each had the illusion of being the sole user of the computer. Time-sharing made possible wide spread interactive use of computers.

As time-sharing systems began to become more available, some computer pioneers realized that two time-sharing computers could be connected, each appearing to the other as just another user. By so doing,

all the users on both systems shared the resources available on the two computers which included the other users. To test how large a system might eventually be possible, a cross country hookup of such systems was attempted in 1965 using long distance telephone lines. The result was a success for long distance time-sharing computer networking but the call set ups and tear downs of the telephone system created time delays that were unacceptable for actual use of such a network.

Telephone switching requires the setup of a complete and dedicated path or circuit before actual end-to-end communication starts to take place. Such communication technology is known as “circuit switching.” The problem is that computer data is often bursty or a message of minimal size as when a single key stroke is sent to solicit a response. Therefore computer data communication over normal switched telephone lines requires frequent call setups or wasteful quiet times. A solution suggested by queuing theory and other lines of reasoning was “packet switching” as opposed to circuit switching. Data to be communicated from a number of sessions could be broken into small packets which would be transmitted interspersed, each routed to its destination separately without setting up a path for each packet. Once at the destination, the packets are reassembled to create an exact copy of the original message. Experimentation with packet switching technology was initiated in Europe and the U.S. starting in 1969 which confirmed the prediction of great efficiency.

Best known of the early packet switching computer networks were the ARPAnet in the U.S., Cyclades in France, and the National Physical Laboratory network in the U.K. The ARPAnet designers and researchers succeeded in achieving resource sharing among time-shared computers manufactured by different vendors and using different operating systems, character sets, and so forth. The computers were located at universities and military related research laboratories. The ARPAnet was funded and encouraged by the Information Processing Techniques Office (IPTO) of the Advanced Research Projects Agency (ARPA), a civilian agency within the U.S. Department of Defense. ARPA/IPTO also funded and encouraged packet switching experimentation using ground based radio receivers and transmitters and using satellites.

Internetting and Catenet

In the early 1970s, in Europe, a number of packet switching network experiments were undertaken. In Hawaii, a successful packet radio network, the ALOHANET, was developed. In the U.S., encouraged by the success of the ARPANet, commercial networks like Tymnet and Telenet were attempted. Just as isolated time-shared computers suggested networking, the existence of isolated packet switching networks, suggested the possibility of some sort of inter-connectivity. Robert Kahn in the U.S. and Louis Pouzin in France were among the first to consider what needed to be done to create such a meta network or internet of networks. Kahn at ARPA/IPTO developed the Internetting Project and Pouzin developed the concept of a Catenet.

The goal of the Internetting and Catenet project and concept was to develop an effective technology to interconnect the packet switching data networks that were beginning to emerge from the experimental stage. Both rejected the alternative of integrating all networks into one single unified network. The later might have produced better integration and performance but would have limited the autonomy and continued experimental development of new network technologies. Also, the developing networks were under different political and economic administrations and it is not likely they could have been enticed to give up their autonomy to voluntarily join together as part of a single network.

Kahn had been involved trying to solve a problem of great complexity: could a ground based packet radio network be developed that would even allow mobile transmitters and receivers? The complexity was that radio communication is prone to fading, interference, obstruction of line-of-site by local terrain or blackout such as when traveling through a tunnel. A radio signal link is by its nature inherently unreliable in itself for data communication. Crucial therefore to the success of such a packet radio network would be an end-to-end mechanism that could arrange for retransmissions and employ other techniques so that a reliable communication service could be provided despite the unreliability of the underlying link level.

Pouzin had worked on the time-sharing experiments at MIT in the 1960s. He was impressed by the successful way individual users were

‘networked’ on a single time sharing computer and then how these computers themselves were networked. He looked for the essence of packet switching networks to give the clue how they could be interconnected. He saw many features which were not mandatory to packet switching such as virtual circuits, end-to-end acknowledgments, large buffer allocations, and so forth. He felt that any end-to-end function which users might desire could be implemented at the user interface. The Catenet need only provide a basic service, packet transport.

Principles for an Internet

How then to achieve an effective interconnection of packet switching networks? If the interconnection was to include packet radio networks the resulting internet would have at least some unreliable links. Should packet radio networks and others that could not offer reliable network service be excluded? Kahn’s answer was that the new interconnection should be open to all packet switching networks. That was the first principle of the Internet that was to emerge: open architecture networking – the interconnection of as many current and future networks as possible by requiring the least commonality possible from each (Leiner, et al, 1998). Each network would be based on the network technology dictated by its own purpose and achieved via its own architectural design. Networks would not be federated into circuits that formed a reliable end to end path, passing individual bits on a synchronous basis. Instead, the new “internetworking architecture” would view networks as peers in helping offer an end-to-end service independent of path or of the unreliability or failure of any links.

“Four ground rules were critical to Kahn’s early thinking:

- * Each distinct network would have to stand on its own and no internal changes could be required to any such network to connect it to the Internet.

- * Communications would be on a best effort basis. If a packet didn’t make it to the final destination, it would shortly be retransmitted from the source.

- * Black boxes would be used to connect the networks; these would later be called gateways and routers. There would be no information retained by the gateways about the individual flows of packets passing

through them, thereby keeping them simple and avoiding complicated adaptation and recovery from various failure modes.

* There would be no global control at the operations level.”

(Leiner et al, 1998)

Pouzin and his colleagues developed similar ground rules and applied them in the development of the Cyclades network and its interconnection with the National Physical Laboratory (NPL) in London in August 1974, with the European Space Agency (ESA) in Rome in October 1975 and with the European Informatics Network (EIN) in June 1976 (Pouzin, 1982). Pouzin’s team implemented a packet service which did not assume any interdependence between packets. Each packet was treated as a separate entity moving from source to destination according to the conditions prevalent at each moment of its travel. There would be dynamic updating of the routing at the gateways and retransmissions because of congestion or link or node failures. Sometimes the packets would arrive at their destinations out of order or duplicated or with some packets missing from a sequence. The gateways were programmed to make an effort to keep the packets moving toward the source but no guarantee of delivery service was built into them. Such a best effort transmission service is called a datagram service.

In the past, out of sequence packets, packet duplication and packet loss were considered at least a burden if not serious problems, so communication switches were designed to prevent them. Now producing end-to-end acknowledgment and retransmission mechanisms rectified these events. In this way substantial simplicity, cost reduction and generality of the service that gateways provided was achieved. By requiring gateways to provide only a datagram service, the interconnection of networks was reduced to its simplest, most universally applicable technology. This was a second Internet principle: as little demand as possible is put on the Internet gateway; or stated conversely, as much as possible be done above the internetwork level. This came to be called the “end-to-end principle” (Carpenter, 1996). It provided for successful communication under almost any condition except the total failure of the whole system. Another way to state this principle was that the information about a communication session (state information) would be at the end points. Intermediate failures could not destroy such information.

Disrupted communication resulting from such failures could be continued when the packets began to arrive again at the destination.

In October 1972, Kahn had organized a large public demonstration of the ARPAnet at the International Computer Communications Conference (ICCC72) in Washington, DC. This was the first international public demonstration of packet switching network technology. Researchers were there from Europe, Asia, and North America. At the meeting, an International Network Working Group (INWG) was established to share experiences and be a forum to help work out standards and protocols. In 1973-74, the INWG was adopted by the International networking professional organization, the International Federation of Information Processing (IFIP) as its Telecommunications Committee Working Group 6.1 (IFIP/TC 6.1). Researchers around the world knew of each other's work and the work of others who were considering these problems by attending and presenting papers at meetings of the IFIP/TC 6.1 and sharing their work with each other on a regular basis. This is an early example of openness and collaboration. This was to become a third principle of the Internet: open and public documentation and open and cooperative standards and protocol development. [See RFC 2555 in this issue.]

TCP/IP and the Internet

In 1973, Kahn brought Vinton Cerf into the work on internetting. Together they sought a general solution to the internetting problem. They aimed to set specifications for what was needed in common on the end computers and the gateways so that the interconnection would be successful. The set of such specifications is called a communication protocol. At first, this protocol was called Transmission Control Protocol (TCP). Cerf and Kahn first shared their thinking in a formal way at a meeting of the INWG members who were in Brighton, England in September, 1973 and then in the article "A Proposal for Packet Network Intercommunication" (Cerf and Kahn, 1974). What they envisioned was a reliable, sequenced, data stream delivery service provided at the end points despite any unreliability of the underlying internetwork level.

The first implementation of TCP only resulted in packets traveling in a circuit like internetwork service. For some network services such

virtual circuits were too restrictive. At the time it was argued by Danny Cohen who was working on packet voice delivery that TCP functionality should be split between what was required end-to-end, like reliability and flow control, and what was required hop-by-hop to get from one network to another via gateways. Cohen felt packet voice needed timeliness more than it needed reliable delivery. This led to the reorganization of the original TCP into two protocols, Internet Protocol (IP) and the Transport Control Protocol (TCP). The simple IP provided for addressing, fragmentation and forwarding of individual packets. The separate TCP provided for recovery from out of sequence and lost packets.

A major boost to the use of what became known as TCP/IP was its adoption by the U.S. Department of Defense (DoD). The DoD funded work in 1979-80 that incorporated TCP/IP into modifications of the Unix operating system being made at the University of California at Berkeley. When this version of Unix was distributed to the universities and other sites that had adopted the Unix operating system, much of the computer science community in the U.S. and around the world began to have TCP/IP capability built into its operating systems. This was a great boost for broad adoption of the Internet. It is also another example of the principle of free and open documentation, in this case open Unix source code. The DoD required all users of the ARPAnet to adopt TCP/IP by January 1, 1983, further insuring that it would be broadly implemented. The transition to TCP/IP was not easy but by April 1, 1983 was pretty much successful.

A key element of the design of IP is the capability at each gateway to break packets too large for the next network into fragments that will fit in the next network's network frames. These fragments then travel along as ordinary datagrams until they are reassembled at the destination host. By allowing for fragmentation, IP makes it possible for large packet handling and small packet handling networks to coexist on the same Internet. This is an example of applying the open architecture principle. Allowing fragmentation relieves the necessity of specifying a minimum or a maximum packet size (although in practice such limits do exist). Leaving the reassembly until the destination minimizes the requirements on the gateway/routers. Schemes that would eliminate

fragmentation from future versions of IP should be carefully scrutinized because they may cause the obsoleting of under resourced networks that could not adapt to the mandated packet sizes. That would violate the open architecture principle.

Conclusion

The highest order feature a communications system can provide is universal connectivity. This has been up until the present the guiding vision and goal of the Internet pioneers. Leonard Kleinrock has argued that “as the system resources grow in size to satisfy an ever increasing population of users” gains in efficiency occur (1976, p. 275). This is an example of the law of large numbers which suggests that the more resources and users there are, the more sharing there is. This results in a greater level of efficient utilization of resources without increased delays in delivery. So far the scaling of the Internet has conformed to the law of large numbers and provides a remarkably inexpensive, convenient and efficient communications system. Also the desire for connectivity grows with the Internet’s growth as does its value since with its growth comes more connectivity for those who have already been connected as well.

In its first 25 years (1973-1998) the Internet grew to provide communication to 2.5% of the world’s people. This is a spectacular technical and social accomplishment. But much of the connectivity is concentrated in a few parts of the world (North America, Europe and parts of Asia). The web of the Internet’s connectivity is also still sparse even in North America. Often, even though there is sufficient total bandwidth, there are too few alternative paths so that the communication service available has uncomfortably long delays.

Top priority for the Internet policy, computer science and technical communities is to find ways of continuing the growth and scaling of the connectivity provided by the Internet. But the principle of universal connectivity to a global Internet communications system is being challenged by those who want instead to convert the Internet into an e-commercenet. The Internet is a new and different technology, making possible a meta-level interconnection of independent networks. To achieve the necessary further scaling, the Internet will require a large

pool of well supported, talented and highly educated scientists and engineers who have studied the principles and unique features of the Internet and who are dedicated to its essence as a communications system. It will also require government policy makers who understand the technology and its social implications or who listen to advisors with such knowledge. All will need to work collaboratively online and off line to hold each other to the principles as they seek solutions to the current and future problems. Then the Internet has a chance of reaching the goal of its pioneers, universal connectivity.

Bibliography

- Carpenter, B. *RFC 1958: Architectural Principles of the Internet*. June, 1996
- Cerf, Vinton G. and Robert Kahn. "A Protocol for Packet Network Intercommunication." *IEEE Transactions on Communications*, Vol. Com-22, No 5. May, 1974.
- Cerf, Vinton G. *IEN48: The Catenet Model for Internetworking*. July, 1978. <http://lwp.ualg.pt/htbin/ien/ien48.html>
- Clark, David D. "The Design Principles of the DARPA Internet Protocols." *Proceedings SIGCOMM88 ACM CCR Vol 18 #4*. August, 1988.
- Comer, Douglas E. *Internetworking with TCP/IP Vol I: Principles, Protocols, and Architecture* 2nd Edition. Englewood Cliffs, NJ. Prentice Hall. 1991.
- Comer, Douglas E. *The Internet Book: Everything You Need to Know about Computer Networking and How the Internet Works*. Englewood Cliffs, NJ. Prentice Hall. 1995.
- Davies, D. W., D. L. A. Barber, W. L. Price C. M. Solomonides. *Computer Networks and Their Protocols*. Chichester. John Wiley & Sons. 1979.
- Hauben, Michael and Ronda Hauben. *Netizens: On the History and Impact of Usenet and the Internet*. Los Alamitos, CA. IEEE Computer Society Press. 1997
- Kleinrock, Leonard. *Queueing Systems Volume II: Computer Applications*. New York. John Wiley and Sons. 1976.
- Leiner, Barry M., et al. "A brief History of the Internet" at: <http://www.isoc.org/internet/history/brief.html>
- Lynch, Daniel C. and Marshall T. Rose. Editors. *Internet Systems Handbook*. Reading, MA. Addison-Wesley. 1993.
- Pouzin, Louis. "A Proposal for Interconnecting Packet Switching Networks." *Proceedings of EUROCOMP*. Brunel University. May, 1974. Pp. 1023-36.
- Pouzin, L. Ed. *The Cyclades Computer Network*. Amsterdam. North Holland. 1982.
- Stevens, W. Richard. *TCP/IP Illustrated, Vol 1 Protocols*. Reading, MA. Addison-Wesley. 1994.

ARPAnet Mailing Lists and Usenet Newsgroups Creating an Open and Scientific Process for Technology Development and Diffusion

by Ronda Hauben
ronda@ais.org

[Editor's Note: Following is the second installment of a longer article about the importance of MsgGroup mailing list and the kinds of lessons it can provide toward determining how to solve the problems of scaling the Internet. The first installment appeared in Vol 9 no 1, pp 38-44]

Part III Government Use at the FCC

While the ARPAnet was helping to research how ARPA would use online communication, other government entities found it helpful in broadening the mechanism of input into their work. Stephen Lukasik had been a director of ARPA from 1970 until 1975. After he left ARPA (then called DARPA), he spent some time at government contractors Xerox and RAND. By September 1979, he posts on MsgGroup:²⁴

“I recently assumed the position of Chief Scientist at the Federal Communications Commission in Washington.”

He notes that he is looking to fill the position of Deputy Chief Scientist/Engineer who will assist him in directing technical, scientific and engineering activities of his office at the FCC. He also announces that there will be positions in a new Technical Planning Staff within the agency. And he requests input from those on MsgGroup.

In October, 1979, Lukasik announced that he was to give the keynote at the December Computer Networking Workshop at the National Bureau of Standards (NBS).²⁵ “The topic will be regulation of

computer communication,” he wrote. And he asks for both questions and input into his talk. “I would be interested to know what questions and concerns you have in this area. Your viewpoints would also be welcome.” He signed his message Steve Lukasik, Chief Scientist, FCC, and his message included “reply to: LUKASIK @usc-isi” so that replies could be sent to him by e-mail.

In February 1981, Einar Stefferud posted an unofficial copy of an FCC Notice of Inquiry (NOI) to MsgGroup, though those interested in receiving an official copy were instructed to write MJMarcus@ISI.²⁶ “This copy is being circulated,” the message explained, “via MsgGroup to allow individuals with ARPAnet access to comment informally on the NOI. Interested parties may file comments on or before March 16, 1981,” Stefferud noted. “You may file informal comments by sending messages to MJMarcus@ISI. To be considered by the FCC, your informal comment should include your full name and U.S. Postal Service Address.”

Stefferd described how it was even possible to file informal comments via e-mail, “All such messages will be forwarded to the Secretary for filing in the Docket as stated in paragraph 23 of the NOI where informal comments are solicited from DEAF-NET users.” DEAF NET was a demonstration telecommunications network project for the deaf funded by Department of Health, Education and Welfare funds. Questions about procedure could be sent by e-mail to Mike Marcus or “with MsgGroup distribution so we may share your questions and answers.”

“Any discussion of this NOI in the regular manner of group discussion via MsgGroup distribution,” Stefferud noted, “will also be made available to the FCC as informal public comments in response to the NOI, and as such will be forwarded to the Secretary for filing in the Docket.”

“This is a new kind of activity for MsgGroup,” Stefferud wrote, “and we hope that it might afford some progress in the use of network facilities for the type of inquiry.” He went on to note that:

The use of MsgGroup is not sponsored by the FCC, though it is understood that FCC staff members are aware of our undertaking.

The text of the Notice of Inquiry in FCC 80-702 General Docket 80-756 followed as a message to MsgGroup. The issue involved Digital Communications protocol conversions between different networks.

E-mail comments to the U.S. Postal Service

Another example of government officials seeking input from MsgGroup participants involved United States Postal Service interface specifications for Electronic Computer-Originated Mail (ECOM). Richard Shuford posting from MIT-AI²⁷ in a message dated July 8, 1981, noted that there had been an announcement in the Federal Register on June 19, 1981 (p. 32111) of a public meeting for questions and comments on the proposed system. That meeting was then held at the Postal Services headquarters in Washington, D.C. However, as there seemed to be no press coverage that the meeting would happen, only “professional Federal Register readers” knew of the meeting to attend it. Shuford described how the result of this situation was that “the meeting was therefore attended only by representatives of large corporations that have some economic interest in what the Postal Service does with electronic mail.” However, a few days before this post on MsgGroup, Shuford had received a call from a Postal Service consultant who worked at SRI International. The consultant said that he wasn’t on the ARPAnet but wanted Shuford to send a message to those on the ARPAnet for him. “He feels very strongly,” wrote Shuford, “that comments on the proposed system should come from a wider variety of ‘stake-holders’ (as he calls them) in the future of electronic mail. In particular, he would like to hear comments from personal computer users and others who are not interested in electronic mail from a purely commercial point of view.” He related how the deadline was in 2 weeks on July 23, 1981 and that comments could be sent by regular mail to Charles Shaw, Director of Electronic-Mail Systems Development at the Postal Service Research and Development Laboratory in Maryland.

Shuford explained that the consultant was making his request in an unofficial capacity and that therefore comments sent should not mention his request.

In response, Pickers at SRI-UNIX observed:²⁸

“In a message which is sent to 100+ institutions, 200+

individuals and spanning both North America and Europe (5 million square miles), the suggestion to keep an individual's name in confidence seems a bit incongruous."

Steve Kudlak, at MIT-MC disagreed. He wrote:²⁹

"ACTUALLY THAT'S NOT TOO UNREASONABLE TO BELIEVE. We all know the ARPAnet is another world and I assume a very high percentage of us are nice enough to hold someone's name in confidence if they requested it."

Several messages later, on July 18, Shuford explains that Ron Newman at Parc-Maxc had located an e-mail address for the consultant and that it was possible to send him one's comments directly by e-mail.³⁰

"He will then have them printed and will pass them along to the proper people at the Postal Services. Please keep in mind," Shuford emphasized, "that any comments passed along in such a manner are officially regarded as 'informal' comments. And that to register 'official' opinion, traditional procedures had to be followed." Thus, a way to make input directly into a government proceeding was available via e-mail.

Debating the Focus of MsgGroup

Many different issues were discussed on MsgGroup and when some on the list suggested limiting what could be discussed, others on the list would invariably complain and encourage a broadness of subjects.

For example, Brian Reid at Carnegie Mellon University, objected to efforts to limit the discussion on MsgGroup. He wrote:³¹

"MsgGroup is the closest that we have to a nationwide Computer science community forum. MsgGroup is supposedly devoted to topics involving electronic mail. One of the many virtues of computer-based mail systems is their astounding ability to support conferencing. All of us are still learning a lot about the ways in which people communicate over these marvelous mail systems, and about the kinds of discussions that can and cannot be made to work over computer-based mail networks."

"Despite the large amount of supposed chitchat that passes over MsgGroup... I believe that such conferencing schemes are still very much at the research stage, and that ARPA and the

public will ultimately benefit from our experiences using MsgGroup as a nationwide community forum, no matter what the topic at hand.”

“Until such time as people start suggesting the overthrow of our government over MsgGroup, I don’t think any sensible topic should be off limits unless you decide that said topic falls outside the scope of MsgGroup. If you decide to restrict the topics that ought to be discussed in MsgGroup, then I submit that there ought to be a “Network-Forum” mailing list which could be a general-purpose forum.”

The crucial issue for the MsgGroup, however, was seen to be the discussion of message systems and eventually of office automation. In May of 1980, Stefferud announced that office automation should be a significant focus of the MsgGroup mailing list. He wrote:³²

“As the ‘Coordinator in Chief’ of MsgGroup, I would like to take this opportunity to ask whether we should shift our focus to office automation in general, as a natural expansion from the message systems orientation that we have had for the last five years? (Yes! Count them, five whole years!)”

“It is my opinion that the ARPAnet provides the best available prototypical office automation environment, one that contains all the required facilities, elements, functions, and features somewhere or other around the net. I use a wide variety of systems on different hosts to get my work done. I truly use the network as my electronic office, which is somewhat remarkable because I am working as a management consultant, rather than as a computer or network technician.”

“Unless we hear some serious dissent, we should consider this change of focus to be a fait accompli.”

Cheers - Stef

His proposal was greeted with support.³³

“I agree wholeheartedly with Stef that we should accept our destiny and let all office automation be within the MsgGroup purview. I, too, conduct large amounts of my work via various network facilities, and often describe the “office of the future” to groups as already existing within the net frame-

work. So by all means let's continue discussions such as the recent one on the Prime OA stuff."

[Howard]

But it was also greeted with an opposing view from Gaines at Rand.³⁴ He wrote:

"I think the term "office automation" is at once too broad and too narrow for the charter of MsgGroup. The MsgGroup ought to broadly focus on issues relevant to computer generation, manipulation, and transmission of messages.... But, there are nevertheless aspects of office automation that are pretty distant from issues related to messages. Taste and judgment rather than any sort of strict rules should be the determinant of whether something is appropriate for the MsgGroup, and we ought to take kindly to rather far removed discussions if somebody considers that they are worth presenting to the MsgGroup. However, I think we ought to still say that our focus is on issues related to computers and messages. The field of office automation is too narrow. Messages are used in other context than what people normally associate with the office environment.... Men communicate for a large variety of reasons in a wide variety of circumstances and we should not narrowly constrain ourselves to any one subset of that universe of communications."

"So here's a vote against a change of focus and a vote for a very wide latitude in interpreting what falls within the purview of MsgGroup," concluded Gaines.

Stefferd responded that his view of office automation was not a narrow one, but a broad one encompassing the broad scope that was being proposed by others. He wrote:³⁵

"Thanks...for your careful comments. I concur with your assessment and suggestion. I see the new focus as being wider as you propose it, but your clarification is very helpful."

"From my ARPAnet experience, I find that office automation should mean the application of computer networking and computer mail facilities to all kinds of work in all possible locations."

“Office Automation does not belong exclusively to the Word Processing Industry any more than to the TWX Switching Industry or the ADP Systems Industry. It belongs to the integration of all these, which to this date has only been demonstrated in these hallowed ARPAnet halls.”

“And, to me, COMPUTER NETWORK MAIL is THE KEY ADDED INGREDIENT.”

So to further set our new context - Onward! Stef

While new and exploratory uses of the Net were tried out on MsgGroup, there was also discussion of the kinds of uses that had to be prevented. A post by Leonard Foner³⁶ explains that as a “tourist” on the ARPAnet he was able to get an account at MIT but had to sign and return an application form which detailed “good uses of MIT’s computer resources, as well as caveats about things that a tourist should not do. It is fairly simple at least to warn them about abusing the network,” he wrote, especially against using it for commercial purposes, which were forbidden. He recommended, “That all users of the net...should be informed as to its intended uses, and what is strictly forbidden (such as profit-making from the Net)... Discussion of funded research on the net seems fine,” he continued, noting that that was what the ARPAnet was created to support.

In 1977, a message from IPTO’s Steve Walker indicated that he would no longer be following MsgGroup in his old status, but that he had found the work done by those participating in MsgGroup very valuable. He wrote:³⁷

“It has been a long time since I have sent a message to this group but I have certainly enjoyed the dialog which has taken place here for the past two and a half years. In remembering all the things that have happened during that time, it is with a good bit of reluctance that I announce my departure from ARPA in late January for a position with the Undersecretary of Defense for Research and Engineering. In my new position I hope to be able to influence the acceptance by the Defense Dept of secure computer systems, interactive message systems and general networking capabilities. I plan to remain active on the ARPAnet and to maintain close contact with groups such as

yours.”

“– I am personally proud to have been associated with the collection of people on the ARPA network who got this whole message handling, electronic mail thing started. Keep up your excellent work.”

“Have a good holiday season,” his message ended.

The Need for Interneting

By 1979, Steve Crocker noted that he and others were working on a project to create a new distributed mail program MMDF, Multi-channel Memo Distribution Facility,³⁸ “to allow mail transmission between machines which have access to a variety of communication lines.” In particular, he wrote, “We want to allow Interneting and to eliminate the need for being attached to the ARPAnet.”

A report by the DCA (Defense Communications Agency) in July 1980 documented how the ARPAnet had grown to over 66 nodes and included 4000-5000 users.³⁹ The report explained how even though the ARPAnet was successful, there were problems. “The basic hardware and software are becoming obsolete,” it noted. It described how the nodes used minicomputers developed in the 1960s which no longer had sufficient memory and other capabilities to support technical components to the network. The ultimate goal, “of our planning,” the report explained, “is to provide for an ARPAnet II which will be a virtual network and will make use of several different networks.”

The report described how in the next 3 years the ARPAnet Host Protocols Network Control Program (NCP) would be replaced with a new DoD Standard Protocol Set. The new protocols were DoD Standard Transmission Control Protocol (TCP) and the Internet Protocol (IP). Also, new computers would replace the IMPs and TIPs that formed the IMP sub-network administered by BBN. All Honeywell equipment was to be replaced with the BBN C/30 costing \$20,000 - \$35,000 (depending on the configuration) if funding could be obtained, and the software would run in a virtual mode.

Unix and the Transition to TCP/IP

Other messages noted that there were many sites that wanted network connections, but that the ARPAnet couldn't accommodate them. It was during this 1979-80 period that Usenet was being introduced at Duke University and the University of North Carolina to provide an online network for those in the Unix community.⁴⁰

In a post on July 4, 1981, Mike Muuss at the Ballistic Research Laboratory noted that it was possible to run Unix on many of the computers being used by those who wanted network connectivity. He wrote:⁴¹

“Unix runs on everything these days.”

This would help facilitate the transition from the IMP with NCP sub-network to TCP/IP protocols that was being planned for January 1, 1983. “There exists AT LEAST one choice of software for UNIX systems,” writes Muuss in a post on the faddiest-p on January 14, 1982,⁴² “(all machines), T(w)agnizes, Multics, and IBES, so the majority of the ‘ordinary’ systems will at least be able to talk, even if non conveniently.” However, he noted that there was not a TCP/IP implementation for the ITS machines at MIT that archived and carried many of the ARPAnet Mailing lists.

By May 3, 1982, a post by Steve Hartwell noted, “Let's not forget, there are more Unix sites than ARPAnet sites.” And Usenet was helping to meet the goal of providing “interoperability among our differently hosted message systems.”⁴³

Also, the problem of large mailing lists had become clear on the ARPAnet. Lists that had several hundred participants like MsgGroup and others were sometimes a heavy load on the host machines that were used to send them out.

Mark Horton noted the superiority of Usenet to ARPAnet for mailing lists as it made it possible to send one copy to each site, rather than having to send out a copy to each person subscribing,⁴⁴ “Note that one of the big points of Usenet is that only one copy of each digest or article is sent to each site....”

Those sites using Unix as their operating system could connect to Usenet and thus, have access to some of the ARPAnet mailing lists. Mark Horton, posting on MsgGroup in 1983 wrote:⁴⁵

“I’ll repeat my invitation to any sites, ARPAnet or otherwise, who want to join Usenet – drop me a line and I’ll point you at a nearby contact. If you run UNIX, the code is all written; if you run something else, you’ll have some work to do....”

Also, by this period several of those who had participated in MsgGroup and the ARPAnet were participants in the discussions on Usenet. And the MsgGroup themes of supporting and exploring the development of communication using an online net-work were continued via Usenet and the ARPAnet mailing lists which were ported to Usenet by Horton at the University of California at Berkeley.

Part IV

The Early Days of Usenet

Usenet was created in 1979 by graduate students at the Duke University and the University of North Carolina who were trying to create a network to connect those who had access to the Unix operating system.⁴⁶ By the summer of 1980, Mark Horton, at the University of California at Berkeley had joined Usenet. Berkeley was also a site on the ARPAnet and Horton soon began to port the discussion from several ARPAnet mailing lists onto Usenet. At first those on Usenet could only read the discussion on the ARPAnet mailing lists, but by Fall 1980, contributions from Usenet participants began to be a part of the ARPAnet lists carried on Usenet.

Among the earliest ARPAnet mailing lists carried on Usenet were Sf-lovers and Human Nets. By Spring of 1981, however, a new mailing list was started to deal with office automation. That mailing list was made available on Usenet as FA.apollo. It was named after one of the workstations. In an early post to the mailing list, Roger Duffy wrote:⁴⁷

Hello,

Welcome to the APOLLO mailing list. APOLLO discusses personal work station computers, such as the APOLLO work station computer, the Three Rivers Corporation PERC, or the recently announced Xerox STAR. APOLLO provides a way for interested members of the ARPAnet community to discuss

what is wrong with these machines, compare notes on work in progress, and share useful insights about these kinds of systems. The list is managed by Hank Dreifus <Dreifus at WHARTON>.

He explained that “APOLLO is currently discussing initial reactions to the Xerox Star Workstation.” And he ended his message, “Lastly, welcome to APOLLO. I trust you will enjoy being part of these discussions.”

A flurry of discussion followed, and it soon began to center on the pros and cons of having a programming language available with the Xerox Star Workstation.

Summarizing responses from those on the mailing list and participating on the Usenet newsgroup, Hank Dreifus at the Wharton School in PA noted several generalizations he felt applied to the subject area.⁴⁸

- o Everyone’s view of Personal Workstations is different.
- o The machine(s) selected are wide ranged and apparently well suited for each application chosen.
- o There is no wrong Personal Workstation machine.
- o The technology of Personal Workstations is not well established as of yet.
- o There is a demonstrated need for this technology, it appears to be one year away from general use.

The summary listed the common characteristics of workstations and described the parts not yet available. “The intention is to educate ourselves about personal workstations,” explained the post, “They sound neat, but what they are under the surface is still a hot topic.”

Particular discussion in the list focused on the Xerox machines – the Xerox Star, their high end machine and the 820, a less expensive product.

Questions were raised as to whether the 820 could be networked to the Star. Others asked what software would be available with the Star⁴⁹ and particularly if there would be a programming system available. One response noted that the Star would come with a low power programming language, but that a more powerful programming environment called the Mesa development system had been developed at Xerox would not be made available.⁵⁰

Apparently, the poster noted, “the reasoning behind this involves consistency in system software.” The post explained that Xerox felt it would keep users from doing harm to the system by restricting access to the Mesa programming environment. Those who wanted new applications would have to ask Xerox to create them.

Another post explained that if Xerox wanted to succeed in selling the Star⁵¹ “it is essential that they provide a decent programming language with it. Otherwise,” the post continued, “it will be just a word processor or maybe a little more.” He went on to explain that those using the Star would need specific specialized applications and only if there was a programming language would it be possible to have those written.

A subsequent post noted that though the initial purchase of the Star was expensive, that would end up being a minimal cost compared to the cost of renting software. He wrote:⁵²

“You people seem to be concentrating on the hardware costs of STAR, which, from my reading of the information available is just the start-up. I think this is like worrying about Gillette’s pricing of the razor-blade holder. Most people will be renting software (blades) forever. This could get very expensive.”

Soon the moderator of the Apollo mailing list announced that the name of this office automation system mailing list would be changed. On Usenet it would become FA.works for personal workstations, as it wasn’t appropriate to name the list after one particular product.⁵³

The economics of buying a workstation was the subject of discussion. One post noted⁵⁴ that because workstations like the Star appeared expensive (\$10,000 per person) they would probably be attractive to managers rather than office peons. Another poster⁵⁵ responded pointing out that for an engineer earning \$30,000 a year, his or her time might cost the company \$60,000, when the cost of the technology being used was added to the salary paid. If having such a personal workstation like the Star made work more productive, it would save the company money and thus, be worth the investment. He wrote⁵⁶ “so if I do my work 10% faster, the company in some way, ‘saves’ 6,000 (the savings could be in hiring less engineers or by getting more work done per unit time or by getting the job done more effectively.)”

Another post cautioned that there was an interest cost to borrowing for capital investment.⁵⁷ “At today’s rates, \$10K capital investment costs the economy 20% interest, either directly because they had to borrow it, or indirectly because they don’t have it to invest elsewhere. So your increase in productivity,” he noted, “would have to be at least 20% to break even. He went on to discuss the difficulty of proving such “increases in productivity.”

One of the participants on the FA.apollo newsgroup, and on the successor newsgroup that followed it, FA.works, was Randy Ivanciw. He had also posted on the MsgGroup list. He became a regular contributor to the FA.Apollo and FA.works.⁵⁸

In his introduction, he wrote:

I am Randy Ivanciw, a computer specialist with the U.S. Army Development and Readiness Command (DARCOM). My major duties include long range and short range planning for office automation. I work at DARCOM headquarters (I am a civilian) as a member of a 7 person staff dealing with the use, planning, implementation and other nasties of office automation.

He explains how the installation at DARCOM benefitted from the discussion on the list, which helped to make possible a broad view of what they were trying to do. He wrote:

In reading the debates pro and con on big systems and little systems, where big systems are large mainframes and little systems are personal workstations.... Let me illustrate how we have attempted to incorporate both worlds in our OA plans. Describing the system he helped create, he writes:

DARCOM has a DEC 10 (DARCOMKA) on the ARPAnet which it uses to provide electronic mail and other OA services to a broad community of users throughout the command (the command is all over this country). Access is via ARPAnet. Advantages here are that for a relatively inexpensive yearly charge a remotely located single user can obtain OA service with a communications capability as powerful as the ARPAnet. This service is in such demand that we cannot supply services in large enough quantities (thus, the DEC 10 will soon be

replaced with a couple of 11/780s to provide more services).

Outlining a 3 level office automation system, he explains how it is used to encourage participation.

For example, let me paint a typical scenario of one of DARCOM's subordinate commands or activities just entering into the world of office automation:

The Commander or somebody at the command wants to try office automation. Now they are unsure of its benefits so they don't want to spend mucho money. They buy a mailbox on our DARCOM-KA (LARGE MAINFRAME). With this mailbox they can experiment with all the OA tools.

After a short while they want 5 or 10 other people at their command or activity to get mailboxes so that they can communicate via electronic mail. They buy more mailboxes on the large mainframe.

Then it is determined that office automation is good for the command. They make large scale plans to provide OA services to 100, or 200, or 300, or how-ever-many people. At this point the economies of scale move toward the LARGE CLUSTER machine. With a large cluster installed locally, the command is essentially running their own OA.

But soon they find that more and more users are demanding service. Enter the small cluster. As one division goes from one or two users (who were getting OA services on the large cluster) to a demand to provide services to 8 or 10 people in that particular division, a micro computer is installed in the division to provide those services (and offset the demand on the large cluster).⁵⁹

His post indicates a process within ARPA encouraging office automation. The discussion on FA.apollo and then FA.works mailing lists proved helpful to those like Ivanciw who were charged with such a task, but who did not find their questions were answered by the vendors. For example, Ivanciw, describes the difficulty he encountered during a sales event trying to get information about how successfully the Xerox 820 and Star Workstations could be connected to the Ethernet. He writes,⁶⁰ "So what it breaks down to is this: there are not too many folks

at Xerox that know how these things connect to the ethernet. The literature is written so that one can assume a lot.”

A response to his post described how the two different Xerox workstations had been developed and how there was ethernet capability really functioning on only one of them. Paul Karger, who had worked at Xerox, wrote:⁶¹

The key to getting through the Xerox propaganda is to realize that there is NOT one, but TWO office automation product lines which have been forcefully “merged.” These lines were developed by two competing groups and don’t really have much in common....

The two product lines evolved and were designed separately.... I hear that the Xerox sales force is claiming that they have an integrated product line for office automation. Low cost 820's up to the Star. Ah ... I don’t think I can agree with that. I believe they are undermining their credibility when they try to convince people of this.

Karger’s post included a diagram with two columns describing the origins of the two sets of products designs.⁶² In a postscript to his message, he wrote:

P.S. Randy – to answer your specific message, the products in column one all have the Ethernet designed and built in from the start. The products in column two have had the Ethernet added with chewing gum and bailing wire (if at all).

TO BE CONTINUED

(Note: The notes corresponding to the numbers in the above article are available from the author via e-mail.)

The opinions expressed in articles are those of their authors and not necessarily the opinions of the *Amateur Computerist* newsletter. We welcome submissions from a spectrum of viewpoints.

ELECTRONIC EDITION

ACN Webpage:

<http://www.ais.org/~jrh/acn/>

All issues of the *Amateur Computerist* are on-line.

Back issues of the *Amateur Computerist* are available at:

http://www.ais.org/~jrh/acn/Back_Issues/

All issues can be accessed from the Index at:

<http://www.ais.org/~jrh/acn/NewIndex.pdf>

EDITORIAL STAFF

Ronda Hauben

William Rohler

Norman O. Thompson

Michael Hauben (1973-2001)

Jay Hauben

The *Amateur Computerist* invites submissions.

Articles can be submitted via e-mail: jrh@ais.org

Permission is given to reprint articles from this issue in a non profit publication provided credit is given, with name of author and source of article cited.