Network Working Group
Request for Comments: 2555
Category: Informational

RFC Editor, et al.
USC/ISI

7 April 1999

30 Years of RFCs

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.

Copyright Notice

Table of Contents

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 innerworkings 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 unfailingly 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 eyeblink 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 email, 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 email 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 email 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 email, 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 email 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 contribute 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 requestors.

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. Email 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 email 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 2-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 email as well. Please send your email 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
EMail: 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
EMail: 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
EMail: crocker@mbl.edu

Vint Cerf
MCI

EMail: vcerf@mci.net

Jake Feinler
SRI Network Information Center
1972-1989

EMail: 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
EMail: 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