

April 7, 2009
Op-Ed Contributor

How the Internet Got Its Rules

By STEPHEN D. CROCKER

Bethesda, Md.

TODAY is an important date in the history of the Internet: the 40th anniversary of what is known as the Request for Comments. Outside the technical community, not many people know about the R.F.C.'s, but these humble documents shape the Internet's inner workings and have played a significant role in its success.

When the R.F.C.'s were born, there wasn't a World Wide Web. Even by the end of 1969, there was just a rudimentary network linking four computers at four research centers: the University of California, Los Angeles; the Stanford Research Institute; the University of California, Santa Barbara; and the [University of Utah](#) in Salt Lake City. The government financed the network and the hundred or fewer computer scientists who used it. It was such a small community that we all got to know one another.

A great deal of deliberation and planning had gone into the network's underlying technology, but no one had given a lot of thought to what we would actually do with it. So, in August 1968, a handful of graduate students and staff members from the four sites began meeting intermittently, in person, to try to figure it out. (I was lucky enough to be one of the U.C.L.A. students included in these wide-ranging discussions.) It wasn't until the next spring that we realized we should start writing down our thoughts. We thought maybe we'd put together a few temporary, informal memos on network protocols, the rules by which computers exchange information. I offered to organize our early notes.

What was supposed to be a simple chore turned out to be a nerve-racking project. Our intent was only to encourage others to chime in, but I worried we might sound as though we were making official decisions or asserting authority. In my mind, I was inciting the wrath of some prestigious professor at some phantom East Coast establishment. I was actually losing sleep over the whole thing, and when I finally tackled my first memo, which dealt with basic communication between two computers, it was in the wee hours of the morning. I had to work in a bathroom so as not to disturb the friends I was staying with, who were all asleep.

Still fearful of sounding presumptuous, I labeled the note a "Request for Comments." R.F.C. 1, written 40 years ago today, left many questions unanswered, and soon became obsolete. But the R.F.C.'s themselves took root and flourished. They became the formal method of publishing Internet protocol standards, and today there are more than 5,000, all readily available online.

But we started writing these notes before we had e-mail, or even before the network was really working, so we wrote our visions for the future on paper and sent them around via the postal service. We'd mail each research group one printout and they'd have to photocopy more themselves.

The early R.F.C.'s ranged from grand visions to mundane details, although the latter quickly became the most common. Less important than the content of those first documents was that they were available free of charge and anyone could write one. Instead of authority-based decision-making, we relied on a process we called "rough consensus and running code." Everyone was welcome to propose ideas, and if enough people liked it and used it, the design became a standard.

After all, everyone understood there was a practical value in choosing to do the same task in the same way. For example, if we wanted to move a file from one machine to another, and if you were to design the process one way, and I was to design it another, then anyone who wanted to talk to both of us would have to employ two distinct ways of doing the same thing. So there was plenty of natural pressure to avoid such hassles. It probably helped that in those days we

avoided patents and other restrictions; without any financial incentive to control the protocols, it was much easier to reach agreement.

This was the ultimate in openness in technical design and that culture of open processes was essential in enabling the Internet to grow and evolve as spectacularly as it has. In fact, we probably wouldn't have the Web without it. When [CERN](#) physicists wanted to publish a lot of information in a way that people could easily get to it and add to it, they simply built and tested their ideas. Because of the groundwork we'd laid in the R.F.C.'s, they did not have to ask permission, or make any changes to the core operations of the Internet. Others soon copied them — hundreds of thousands of computer users, then hundreds of millions, creating and sharing content and technology. That's the Web.

Put another way, we always tried to design each new protocol to be both useful in its own right and a building block available to others. We did not think of protocols as finished products, and we deliberately exposed the internal architecture to make it easy for others to gain a foothold. This was the antithesis of the attitude of the old telephone networks, which actively discouraged any additions or uses they had not sanctioned.

Of course, the process for both publishing ideas and for choosing standards eventually became more formal. Our loose, unnamed meetings grew larger and semi-organized into what we called the Network Working Group. In the four decades since, that group evolved and transformed a couple of times and is now the Internet Engineering Task Force. It has some hierarchy and formality but not much, and it remains free and accessible to anyone.

The R.F.C.'s have grown up, too. They really aren't requests for comments anymore because they are published only after a lot of vetting. But the culture that was built up in the beginning has continued to play a strong role in keeping things more open than they might have been. Ideas are accepted and sorted on their merits, with as many ideas rejected by peers as are accepted.

As we rebuild our economy, I do hope we keep in mind the value of openness, especially in industries that have rarely had it. Whether it's in health care reform or energy innovation, the largest payoffs will come not from what the stimulus package pays for directly, but from the huge vistas we open up for others to explore.

I was reminded of the power and vitality of the R.F.C.'s when I made my first trip to Bangalore, India, 15 years ago. I was invited to give a talk at the Indian Institute of Science, and as part of the visit I was introduced to a student who had built a fairly complex software system. Impressed, I asked where he had learned to do so much. He simply said, "I downloaded the R.F.C.'s and read them."

Stephen D. Crocker is the chief executive of a company that develops information-sharing technology.