

# **Protocol**

**How Control Exists after Decentralization**

*Alexander R. Galloway*

The MIT Press Cambridge, Massachusetts London, England

© 2004 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

This book was set in Garamond and Bell Gothic by Graphic Composition, Inc., and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Galloway, Alexander R., 1974—

Protocol : how control exists after decentralization / Alexander R. Galloway.

p. cm.—(Leonardo)

Includes bibliographical references and index.

ISBN 0-262-07247-5 (alk. paper)

1. Computer networks—Security measures. 2. Computer networks—Management. 3. Computer network protocols. 4. Electronic data processing—Distributed processing. I. Title II. Leonardo (Series) (Cambridge, Mass.)

TK5105.59.G35 2004

005.8—dc22

2003060361

10 9 8 7 6 5 4 3 2 1

## II

# Failures of Protocol

## Institutionalization

In the Internet, there is no central node, and only a minimal centralized management structure, limited to a few housekeeping functions such as standards setting.

— PAUL BARAN, “Is the UHF Frequency Shortage a Self Made Problem?”

We define mechanism, not policy.

— TIM BERNERS-LEE, *Weaving the Web*

On April 12, 1994, the protocological organization of the Internet suffered a major setback. On that black Tuesday, an unsolicited commercial email message was sent systematically to each and every newsgroup in the Usenet system, violating the informational network's customary prohibition against such commercial advertisements.<sup>1</sup>

Spam was born. The perpetrators, Arizona lawyers Laurence Canter and Martha Seigel,<sup>2</sup> had effectively transformed a democratic, protocological system for exchange of ideas into a unilateral, homogenous tool for commercial solicitation.

A quick description of Usenet is as follows:

Usenet has evolved some of the best examples of decentralized control structures on the Net. There is no central authority that controls the news system. The addition of new newsgroups to the main topic hierarchy is controlled by a rigorous democratic process, using the Usenet group *news.admin* to propose and discuss the creation of new groups. After a new group is proposed and discussed for a set period of time, anyone with an email address may submit an email vote for or against the proposal. If a newsgroup vote passes, a new group message is sent and propagated through the Usenet network.<sup>3</sup>

This protocological covenant outlining open channels for Usenet's growth and governance, hitherto cultivated and observed by its large, diverse com-

---

Epigraphs: Paul Baran, "Is the UHF Frequency Shortage a Self Made Problem?" Paper presented at the Marconi Centennial Symposium, Bologna, Italy, June 23, 1995. Tim Berners-Lee, *Weaving the Web* (New York: HarperCollins, 1999), p. 124.

1. This standard of etiquette is articulated in Sally Hambridge and Albert Lunde's RFC on the topic, "DON'T SPEW: A Set of Guidelines for Mass Unsolicited Mailings and Postings (spam)," RFC 2635, FYI 35, June 1999. See also Sally Hambridge, "Netiquette Guidelines," RFC 1855, FYI 28, October 1995. Stopgap technical solutions for reducing the amount of spam are outlined in Gunnar Lindberg's "Anti-Spam Recommendations for SMTP MTAs," RFC 2505, BCP 30, February 1999.

2. The two document this and other questionable practices in their book *How to Make a Fortune on the Information Superhighway: Everyone's Guerrilla Guide to Marketing on the Internet and Other On-Line Services* (New York: HarperCollins, 1995).

3. Nelson Minar and Marc Hedlund, "A Network of Peers," in *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, ed. Andy Oram (Sebastopol: O'Reilly, 2001), p. 6.

munity of scientists and hobbyists, was sullied in the spam incident by the infraction of a few. The diversity of the many groups on Usenet was erased and covered by a direct-mail blanket with a thoroughness only computers can accomplish. As I stated earlier, protocol requires universal adoption. As a protocological product, Usenet is vulnerable because of this. Even a single party can exploit a weakness and, like a virus, propagate through the system with logical ferocity.

In part I I described how protocol has succeeded as a dominant principle of organization for distributed networks. Yet at the same time the spam incident of April 12, 1994, illustrates that there have been numerous instances where protocol has, in a sense, failed. The openness of the network was wrenched away from its users and funneled toward a single commercial goal. What was multiple became singular. What was contingent and detached became directed and proprietary.

Failures of protocol occur in many places of contemporary life, from the dominance of international capitalism and the World Trade Organization, itself a power center that buckled under distributed, protocological protests against it in Seattle in 1999, to the monolithic Microsoft and its battle with the U.S. Justice Department (the anti-Microsoft action is, to be precise, a failure of a failure of protocol).

By failure I mean to point out not a failure on protocol's own terms (that's what part III of this book is for), but a failure for protocol to blossom fully as a management diagram. That is to say, this section is not about how protocol doesn't work—because it does, very well—but how protocol is not allowed to work purely on its own terms.

This chapter, then, covers how protocol has emerged historically within a context of bureaucratic and institutional interests, a reality that would seem to contradict protocol. And indeed it does. (Or, as I will put it at the end of this chapter, in a sense protocol has to fail in order to succeed, to fail tactically in order to succeed strategically.) While in Paul Baran's estimation these interests are a "minimal" management structure, they have exerted influence over the network in significant ways. Proprietary or otherwise commercial interests (from the spam incident to Microsoft and everything in between) also represent a grave threat to and failure of protocol.

To date, most of the literature relating to my topic has covered protocol through these issues of law, governance, corporate control, and so on. Lawrence Lessig is an important thinker in this capacity. So I do not cover that in de-

tail in this chapter. But in passing consider this heuristic: It is possible to think of bureaucratic interests as visiting protocol *from without* due to the imposition of a completely prior and foreign control diagram, while proprietary interests arrive *from within* as a coopting of protocol's own explosive architecture. Bureaucracy is protocol atrophied, while propriety is protocol reified. Both represent grave challenges to the effective functioning of protocol within digital computer networks.

Let me say also that this is the least significant section—and indeed because of that, the most significant—to read if one is to understand the true apparatus of protocol. The argument in this book is that bureaucratic and institutional forces (as well as proprietary interests) are together the inverse of protocol's control logic. This is why I have not yet, and will not, define protocol's power in terms of either commercial control, organizational control, juridical control, state control, or anything of the like. Protocol gains its authority from another place, from technology itself and how people program it.

To be precise, many believe that bureaucratic organizations such as ICANN (the Internet Corporation for Assigned Names and Numbers) are synonymous with protocol because they regulate and control the Net. But the opposite is true. Organizations like ICANN are the enemy of protocol because they limit the open, free development of technology. (It is for this reason that I have waited until this chapter to discuss the RFCs in detail, rather than talking about them in chapter 1.)

Likewise, the market monopoly of Intel in the field of microchips or of Microsoft in the field of personal computer software appears to many to constitute a type of protocol, a broad technical standard. But, again, market monopolies of proprietary technologies are the inverse, or enemy, of protocol, for they are imposed from without, are technically opaque, centrally controlled, deployed by commercial concerns, and so on.

As long-time RFC editor Jon Postel put it, "I think three factors contribute to the success of the Internet: (1) public documentation of the protocols, (2) free (or cheap) software for the popular machines, and (3) vendor

---

4. See Jon Postel's biographical entry in Gary Malkin's "Who's Who in the Internet: Biographies of IAB, IESG and IRSG Members," RFC 1336, FYI 9, May 1992.

independence.”<sup>4</sup> Commercial or regulatory interests have historically tended to impinge upon Postel’s three factors. Standards bodies like the Institute of Electrical and Electronics Engineers (IEEE) make a point of publishing standards that do not reference or favor any specific commercial vendor. (They accomplish this by describing how a technology should perform, not any specific design implementation, which may be linked to a specific commercial product or patented technology.) Hence, this chapter is nothing but a prophylactic. It addresses the negative influences that restrict protocol’s full potential.

In short, *protocol is a type of controlling logic that operates outside institutional, governmental, and corporate power*, although it has important ties to all three.

In this day and age, technical protocols and standards are established by a self-selected oligarchy of scientists consisting largely of electrical engineers and computer specialists. Composed of a patchwork of many professional bodies, working groups, committees, and subcommittees, this technocratic elite toils away, mostly voluntarily, in an effort to hammer out solutions to advancements in technology. Many of them are university professors. Most all of them either work in industry or have some connection to it.

Like the philosophy of protocol itself, membership in this technocratic ruling class is open. “Anyone with something to contribute could come to the party,”<sup>5</sup> wrote one early participant. But, to be sure, because of the technical sophistication needed to participate, this loose consortium of decision makers tends to fall into a relatively homogenous social class: highly educated, altruistic, liberal-minded science professionals from modernized societies around the globe.

And sometimes not so far around the globe. Of the twenty-five or so original protocol pioneers, three of them—Vint Cerf, Jon Postel, and Steve Crocker—all came from a single high school in Los Angeles’s San Fernando Valley.<sup>6</sup> Furthermore, during his long tenure as RFC editor, Postel was the single gatekeeper through whom all protocol RFCs passed before they could be published. Internet historians Katie Hafner and Matthew Lyon describe

---

5. Jake Feinler, “30 Years of RFCs,” RFC 2555, April 7, 1999.

6. See Vint Cerf’s memorial to Jon Postel’s life and work in “I Remember IANA,” RFC 2468, October 1988.

this group as “an ad-hocracy of intensely creative, sleep-deprived, idiosyncratic, well-meaning computer geniuses.”<sup>7</sup>

There are few outsiders in this community. Here the specialists run the show. To put it another way, while the Internet is used daily by vast swaths of diverse communities, the standards makers at the heart of this technology are a small entrenched group of techno-elite peers. The reasons for this are largely practical. “Most users are not interested in the details of Internet protocols,” Cerf observes. “They just want the system to work.”<sup>8</sup> Or as former IETF Chair Fred Baker reminds us: “The average user doesn’t write code. . . . If their needs are met, they don’t especially care how they were met.”<sup>9</sup>

So who actually writes these technical protocols, where did they come from, and how are they used in the real world? They are found in the fertile amalgamation of computers and software that constitutes the majority of servers, routers, and other Internet-enabled machines. A significant portion of these computers were, and still are, Unix-based systems. A significant portion of the software was, and still is, largely written in the C or C++ languages. All of these elements have enjoyed unique histories as protocological technologies.

The Unix operating system was developed at Bell Telephone Laboratories by Ken Thompson, Dennis Ritchie, and others beginning in 1969, and development continued into the early 1970s. After the operating system’s release, the lab’s parent company, AT&T, began to license and sell Unix as a commercial software product. But, for various legal reasons, the company admitted that it “had no intention of pursuing software as a business.”<sup>10</sup> Unix was indeed sold by AT&T, but simply “as is” with no advertising, technical support, or other fanfare. This contributed to its widespread adoption

---

7. Katie Hafner and Matthew Lyon, *Where Wizards Stay Up Late: The Origins of the Internet* (New York: Touchstone, 1996), p. 145. For biographies of two dozen protocol pioneers, see Gary Malkin’s “Who’s Who in the Internet: Biographies of IAB, IESG and IRSG Members,” RFC 1336, FYI 9, May 1992.

8. Vinton Cerf, personal correspondence, September 23, 2002.

9. Fred Baker, personal correspondence, December 12, 2002.

10. AT&T’s Otis Wilson, cited in Peter Salus, *A Quarter Century of Unix* (New York: Addison-Wesley, 1994), p. 59.

by universities who found in Unix a cheap but useful operating system that could be easily experimented with, modified, and improved.

In January 1974, Unix was installed at the University of California at Berkeley. Bill Joy and others began developing a spin-off of the operating system that became known as BSD (Berkeley Software Distribution).

Unix was particularly successful because of its close connection to networking and the adoption of basic interchange standards. “Perhaps the most important contribution to the proliferation of Unix was the growth of networking,”<sup>11</sup> writes Unix historian Peter Salus. By the early 1980s, the TCP/IP networking suite was included in BSD Unix.

Unix was designed with openness in mind. The source code—written in C, which was also developed during 1971–1973—is easily accessible, meaning a higher degree of technical transparency.

The standardization of the C programming language began in 1983 with the establishment of an American National Standards Institute (ANSI) committee called “X3J11.” The ANSI report was finished in 1989 and subsequently accepted as a standard by the international consortium ISO in 1990.<sup>12</sup> Starting in 1979, Bjarne Stroustrup developed C++, which added the concept of classes to the original C language. (In fact, Stroustrup’s first nickname for his new language was “C with Classes.”) ANSI standardized the C++ language in 1990.

C++ has been tremendously successful as a language. “The spread was world-wide from the beginning,” recalled Stroustrup. “[I]t fit into more environments with less trouble than just about anything else.”<sup>13</sup> Just like a protocol.

It is not only computers that experience standardization and mass adoption. Over the years many technologies have followed this same trajectory. The process of standards creation is, in many ways, simply the recognition

---

11. Salus, *A Quarter Century of Unix*, p. 2.

12. See Dennis Ritchie, “The Development of the C Programming Language,” in *History of Programming Languages II*, ed. Thomas Bergin and Richard Gibson (New York: ACM, 1996), p. 681.

13. Bjarne Stroustrup, “Transcript of Presentation,” in *History of Programming Languages II*, ed. Thomas Bergin and Richard Gibson (New York: ACM, 1996), p. 761.

of technologies that *have* experienced success in the marketplace. One example is the VHS video format developed by JVC (with Matsushita), which edged out Sony's Betamax format in the consumer video market. Betamax was considered by some to be a superior technology (an urban myth, claim some engineers) because it stored video in a higher-quality format. But the trade-off was that Betamax tapes tended to be shorter in length. In the late 1970s when VHS launched, the VHS tape allowed for up to two hours of recording time, while Betamax provided only one hour. "By mid 1979 VHS was outselling Beta by more than 2 to 1 in the US."<sup>14</sup> When Betamax caught up in length (to three hours), it had already lost a foothold in the market. VHS would counter Betamax by increasing to four hours and later eight.

Some have suggested that it was the pornography industry, which favored VHS over Betamax, that provided it with legions of early adopters and proved the long-term viability of the format.<sup>15</sup> But perhaps the most convincing argument is the one that points out JVC's economic strategy that included aggressive licensing of the VHS format to competitors. JVC's behavior is pseudo-protocological. The company licensed the technical specifications for VHS to other vendors. It also immediately established manufacturing and distribution supply chains for VHS tape manufacturing and retail sales. In the meantime Sony tried to fortify its market position by keeping Betamax to itself. As one analyst writes:

Three contingent early differences in strategy were crucial. First, Sony decided to proceed without major co-sponsors for its Betamax system, while JVC shared VHS with several major competitors. Second, the VHS consortium quickly installed a

---

14. S. J. Liebowitz and Stephen E. Margolis, "Path Dependence, Lock-In and History," *Journal of Law, Economics and Organization* 11, April 1995. Available online at <http://wwwpub.utdallas.edu/~liebowit/paths.html>.

15. If not VHS then the VCR in general was aided greatly by the porn industry. David Morton writes that "many industry analysts credited the sales of erotic video tapes as one of the chief factors in the VCR's early success. They took the place of adult movie theaters, but also could be purchased in areas where they were legal and viewed at home." See Morton's *A History of Electronic Entertainment since 1945*, available online at [http://www.ieee.org/organizations/history\\_center/research\\_guides/entertainment](http://www.ieee.org/organizations/history_center/research_guides/entertainment), p. 56.

large manufacturing capacity. Third, Sony opted for a more compact cassette, while JVC chose a longer playing time for VHS, which proved more important to most customers.<sup>16</sup>

JVC deliberately sacrificed larger profit margins by keeping prices low and licensing to competitors. This was in order to grow its market share. The rationale was that establishing a standard was the most important thing, and as JVC approached that goal, it would create a positive feedback loop that would further beat out the competition.

The VHS/Betamax story is a good example from the commercial sector of how one format can triumph over another format to become an industry standard. This example is interesting because it shows that protocological behavior (giving out your technology broadly even if it means giving it to your competitors) often wins out over proprietary behavior. The Internet protocols function in a similar way, to the degree that they have become industry standards not through a result of proprietary market forces, but due to broad open initiatives of free exchange and debate. This was not exactly the case with VHS, but the analogy is useful nevertheless.

This type of corporate squabbling over video formats has since been essentially erased from the world stage with the advent of DVD. This new format was reached through consensus from industry leaders and hence does not suffer from direct competition by any similar technology in the way that VHS and Betamax did. Such consensus characterizes the large majority of processes in place today around the world for determining technical standards.

Many of today's technical standards can be attributed to the IEEE (pronounced "eye triple e"). In 1963 IEEE was created through the merging of two professional societies. They were the American Institute of Electrical Engineers (AIEE) founded in New York on May 13, 1884 (by a group that included Thomas Edison) and the Institute of Radio Engineers (IRE) founded in 1912.<sup>17</sup> Today the IEEE has over 330,000 members in 150 countries. It is the world's largest professional society in any field. The IEEE works in con-

---

16. Douglas Puffert, "Path Dependence in Economic Theory." Available online at [http://www.vwl.uni-muenchen.de/lis\\_komlos/pathe.pdf](http://www.vwl.uni-muenchen.de/lis_komlos/pathe.pdf), p. 5.

17. *IEEE 2000 Annual Report*, available online at <http://www.ieee.org>.

junction with industry to circulate knowledge of technical advances, to recognize individual merit through the awarding of prizes, and to set technical standards for new technologies. In this sense the IEEE is the world's largest and most important protological society.

Composed of many chapters, subgroups, and committees, the IEEE's Communications Society is perhaps the most interesting area vis-à-vis computer networking. It establishes standards in many common areas of digital communication including digital subscriber lines (DSLs) and wireless telephony.

IEEE standards often become international standards. Examples include the "802" series of standards that govern network communications protocols. These include standards for Ethernet<sup>18</sup> (the most common local area networking protocol in use today), Bluetooth, Wi-Fi, and others.

"The IEEE," Paul Baran observed, "has been a major factor in the development of communications technology."<sup>19</sup> Indeed Baran's own theories, which eventually would spawn the Internet, were published within the IEEE community even as they were published by his own employer, the Rand Corporation.

Active within the United States are the National Institute for Standardization and Technology (NIST) and ANSI. The century-old NIST, formerly known as the National Bureau of Standards, is a federal agency that develops and promotes technological standards. Because it is a federal agency and not a professional society, it has no membership per se. It is also nonregulatory, meaning that it does not enforce laws or establish mandatory standards that must be adopted. Much of its budget goes into supporting NIST research laboratories as well as various outreach programs.

ANSI, formerly called the American Standards Association, is responsible for aggregating and coordinating the standards creation process in the

---

18. The IEEE prefers to avoid associating its standards with trademarked, commercial, or otherwise proprietary technologies. Hence the IEEE definition eschews the word "Ethernet," which is associated with Xerox PARC where it was named. The 1985 IEEE standard for Ethernet is instead titled "IEEE 802.3 Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications."

19. Paul Baran, Electrical Engineer, an oral history conducted in 1999 by David Hochfelder, IEEE History Center, Rutgers University, New Brunswick, NJ, USA.

United States. It is the private-sector counterpart to NIST. While it does not create any standards itself, it is a conduit for federally accredited organizations in the field who are developing technical standards. The accredited standards developers must follow certain rules designed to keep the process open and equitable for all interested parties. ANSI then verifies that the rules have been followed by the developing organization before the proposed standard is adopted.

ANSI is also responsible for articulating a national standards strategy for the United States. This strategy helps ANSI advocate in the international arena on behalf of U.S. interests. ANSI is the only organization that can approve standards as American national standards.

Many of ANSI's rules for maintaining integrity and quality in the standards development process revolve around principles of openness and transparency and hence conform with much of what I have already said about protocol. ANSI writes that

- Decisions are reached through *consensus* among those affected.
- Participation is *open* to all affected interests. . . .
- The process is *transparent*—information on the process and progress is directly available. . . .
- The process is *flexible*, allowing the use of different methodologies to meet the needs of different technology and product sectors.<sup>20</sup>

Besides being consensus-driven, open, transparent, and flexible, ANSI standards are also voluntary, which means that, like NIST, no one is bound by law to adopt them. Voluntary adoption in the marketplace is the ultimate test of a standard. Standards may disappear in the advent of a new superior technology or simply with the passage of time. Voluntary standards have many advantages. By not forcing industry to implement the standard, the burden of success lies in the marketplace. And in fact, proven success in the marketplace generally predates the creation of a standard. The behavior is emergent, not imposed.

---

20. ANSI, "National Standards Strategy for the United States." Available online at <http://www.ansi.org>, emphasis in original.

On the international stage several other standards bodies become important. The International Telecommunication Union (ITU) focuses on radio and telecommunications, including voice telephony, communications satellites, data networks, television, and, in the old days, the telegraph. Established in 1865, it is the world's oldest international organization.

The International Electrotechnical Commission (IEC) prepares and publishes international standards in the area of electrical technologies including magnetics, electronics, and energy production. They cover everything from screw threads to quality management systems. IEC is comprised of national committees. (The national committee representing the United States is administered by ANSI.)

Another important international organization is ISO, also known as the International Organization for Standardization.<sup>21</sup> Like IEC, ISO grew out of the electro-technical field and was formed after World War II to "facilitate the international coordination and unification of industrial standards."<sup>22</sup> Based in Geneva, but a federation of over 140 national standards bodies including the American ANSI and the British Standards Institution (BSI), its goal is to establish vendor-neutral technical standards. Like the other international bodies, standards adopted by the ISO are recognized worldwide.

Also like other standards bodies, ISO develops standards through a process of consensus-building. Its standards are based on voluntary participation, and thus the adoption of ISO standards is driven largely by market forces (as opposed to mandatory standards that are implemented in response to a governmental regulatory mandate). Once established, ISO standards can have massive market penetration. For example, the ISO standard for film speed (100, 200, 400, etc.) is used globally by millions of consumers.

Another ISO standard of far-reaching importance is the Open Systems Interconnection (OSI) Reference Model. Developed in 1978, the OSI Reference Model is a technique for classifying all networking activity into seven

---

21. The name ISO is in fact not an acronym, but derives from a Greek word for "equal." This way it avoids the problem of translating the organization's name into different languages, which would produce different acronyms. The name ISO, then, is a type of semantic standard in itself.

22. See <http://www.iso.ch> for more history of the ISO.

abstract layers. Each layer describes a different segment of the technology behind networked communication, as described in chapter 1.

Layer 7	Application
Layer 6	Presentation
Layer 5	Session
Layer 4	Transport
Layer 3	Network
Layer 2	Data link
Layer 1	Physical

This classification, which helps organize the process of standardization into distinct areas of activity, is relied on heavily by those creating data networking standards.

In 1987 ISO and IEC recognized that some of their efforts were beginning to overlap. They decided to establish an institutional framework to help coordinate their efforts and formed a joint committee to deal with information technology called the Joint Technical Committee 1 (JTC 1). ISO and IEC both participate in the JTC 1, as well as liaisons from Internet-oriented consortia such as the IETF. ITU members, IEEE members, and others from other standards bodies also participate here. Individuals may sit on several committees in several different standards bodies, or simply attend as ex officio members, to increase inter-organizational communication and reduce redundant initiatives between the various standards bodies. JTC 1 committees focus on everything from office equipment to computer graphics. One of the newest committees is devoted to biometrics.

ISO, ANSI, IEEE, and all the other standards bodies are well-established organizations with long histories and formidable bureaucracies. The Internet, on the other hand, has long been skeptical of such formalities and spawned a more ragtag, shoot-from-the-hip attitude about standard creation.<sup>23</sup> I fo-

---

23. The IETF takes pride in having such an ethos. Jeanette Hofmann writes: "The IETF has traditionally understood itself as an elite in the technical development of communication networks. Gestures of superiority and a dim view of other standardisation committees are matched by unmistakable impatience with incompetence in their own ranks." See Hofmann, "Govern-

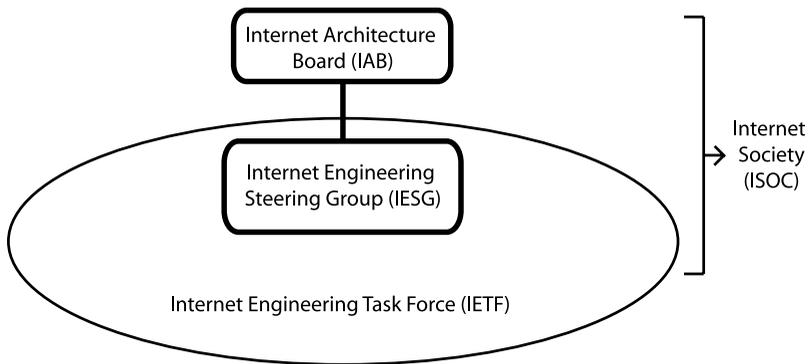


Figure 4.1  
ISOC chart

cus the rest of this chapter on those communities and the protocol documents that they produce.

Four groups make up the organizational hierarchy in charge of Internet standardization. They are the Internet Society, the Internet Architecture Board, the Internet Engineering Steering Group, and the Internet Engineering Task Force.<sup>24</sup>

The Internet Society (ISOC), founded in January 1992, is a professional membership society. It is the umbrella organization for the other three groups. Its mission is “to assure the open development, evolution and use of

---

ment Technologies and Techniques of Government: Politics on the Net.” Available online at <http://duplox.wz-berlin.de/final/jeanette.htm>.

24. Another important organization to mention is the Internet Corporation for Assigned Names and Numbers (ICANN). ICANN is a nonprofit organization that has control over the Internet’s DNS. Its board of directors has included Vinton Cerf, coinventor of the Internet Protocol and founder of the Internet Society, and author Esther Dyson. “It is ICANN’s objective to operate as an open, transparent, and consensus-based body that is broadly representative of the diverse stakeholder communities of the global Internet” (see “ICANN Fact Sheet,” available online at <http://www.icann.org>). Despite this rosy mission statement, ICANN has been the target of intense criticism in recent years. It is for many the central lightning rod for problems around issues of Internet governance. A close look at ICANN is unfortunately outside the scope of this book, but for an excellent examination of the organization, see Milton Mueller’s *Ruling the Root* (Cambridge: MIT Press, 2002).

the Internet for the benefit of all people throughout the world.”<sup>25</sup> It facilitates the development of Internet protocols and standards. ISOC also provides fiscal and legal independence for the standards-making process, separating this activity from its former U.S. government patronage.

The Internet Architecture Board (IAB), originally called the Internet Activities Board, is a core committee of thirteen, nominated by and consisting of members of the IETF.<sup>26</sup> The IAB reviews IESG appointments, provides oversight of the architecture of network protocols, oversees the standards creation process, hears appeals, oversees the RFC editor, and performs other chores. The IETF (as well as the Internet Research Task Force, which focuses on longer-term research topics) falls under the auspices of the IAB. The IAB is primarily an oversight board, since actually accepted protocols generally originate within the IETF (or in smaller design teams).

Underneath the IAB is the Internet Engineering Steering Group (IESG), a committee of the Internet Society that assists and manages the technical activities of the IETF. All of the directors of the various research areas in the IETF are part of this steering group.

The bedrock of this entire community is the IETF. The IETF is the core area where most protocol initiatives begin. Several thousand people are involved in the IETF, mostly through email lists, but also in face-to-face meetings. “The Internet Engineering Task Force is,” in its own words, “a loosely self-organized group of people who make technical and other contributions to the engineering and evolution of the Internet and its technologies.”<sup>27</sup> Or elsewhere: “the Internet Engineering Task Force (IETF) is an open global community of network designers, operators, vendors, and researchers producing technical specifications for the evolution of the Internet architecture and the smooth operation of the Internet.”<sup>28</sup>

---

25. See <http://www.isoc.org>.

26. For a detailed description of the IAB, see Brian Carpenter, “Charter of the Internet Architecture Board (IAB),” RFC 2850, BCP 39, May 2000.

27. Gary Malkin, “The Tao of IETF: A Guide for New Attendees of the Internet Engineering Task Force,” RFC 1718, FYI 17, October 1993.

28. Paul Hoffman and Scott Bradner, “Defining the IETF,” RFC 3233, BCP 58, February 2002.

The IETF is best defined in the following RFCs:

- “The Tao of IETF: A Guide for New Attendees of the Internet Engineering Task Force” (RFC 1718, FYI 17)
- “Defining the IETF” (RFC 3233, BCP 58)
- “IETF Guidelines for Conduct”<sup>29</sup> (RFC 3184, BCP 54)
- “The Internet Standards Process—Revision 3” (RFC 2026, BCP 9)
- “IAB and IESG Selection, Confirmation, and Recall Process: Operation of the Nominating and Recall Committees” (RFC 2727, BCP 10)
- “The Organizations Involved in the IETF Standards Process” (RFC 2028, BCP 11)

These documents describe both how the IETF creates standards and also how the entire community itself is set up and how it behaves.

The IETF is the least bureaucratic of all the organizations mentioned in this chapter. In fact it is not an organization at all, but rather an informal community. It does not have strict bylaws or formal officers. It is not a corporation (nonprofit or otherwise) and thus has no board of directors. It has no binding power as a standards creation body and is not ratified by any treaty or charter. It has no membership, and its meetings are open to anyone. “Membership” in the IETF is simply evaluated through an individual’s participation. If you participate via email, or attend meetings, you are a member of the IETF. All participants operate as unaffiliated individuals, not as representatives of other organizations or vendors.

The IETF is divided by topic into various Working Groups. Each Working Group<sup>30</sup> focuses on a particular issue or issues and drafts documents that

---

29. This RFC is an interesting one because of the social relations it endorses within the IETF. Liberal, democratic values are the norm. “Intimidation or ad hominem attack” is to be avoided in IETF debates. Instead IETFers are encouraged to “think globally” and treat their fellow colleagues “with respect as persons.” Somewhat ironically, this document also specifies that “English is the de facto language of the IETF.” See Susan Harris, “IETF Guidelines for Conduct,” RFC 3184, BCP 54, October 2001.

30. For more information on IETF Working Groups, see Scott Bradner, “IETF Working Group Guidelines and Procedures,” RFC 2418, BCP 25, September 1998.

are meant to capture the consensus of the group. Like protocols created by other standards bodies, IETF protocols are voluntary standards. There is no technical or legal requirement<sup>31</sup> that anyone actually adopt IETF protocols.

The process of establishing an Internet Standard is gradual, deliberate, and negotiated. Any protocol produced by the IETF goes through a series of stages, called the “standards track.” The standards track exposes the document to extensive peer review, allowing it to mature into an RFC memo and eventually an Internet Standard. “The process of creating an Internet Standard is straightforward,” they write. “A specification undergoes a period of development and several iterations of review by the Internet community and revision based upon experience, is adopted as a Standard by the appropriate body. . . , and is published.”<sup>32</sup>

Preliminary versions of specifications are solicited by the IETF as Internet-Draft documents. Anyone can submit an Internet-Draft. They are not standards in any way and should not be cited as such nor implemented by any vendors. They are works in progress and are subject to review and revision. If they are deemed uninteresting or unnecessary, they simply disappear after their expiration date of six months. They are not RFCs and receive no number.

If an Internet-Draft survives the necessary revisions and is deemed important, it is shown to the IESG and nominated for the standards track. If the IESG agrees (and the IAB approves), then the specification is handed off to the RFC editor and put in the queue for future publication. Cronyism is sometimes a danger at this point, as the old-boys network—the RFC editor, the IESG, and the IAB—have complete control over which Internet-Drafts are escalated and which aren’t.

---

31. That said, there are protocols that are given the status level of “required” for certain contexts. For example, the Internet Protocol is a required protocol for anyone wishing to connect to the Internet. Other protocols may be given status levels of “recommended” or “elective” depending on how necessary they are for implementing a specific technology. The “required” status level should not be confused however with mandatory standards. These have legal implications and are enforced by regulatory agencies.

32. Scott Bradner, “The Internet Standards Process—Revision 3,” RFC 2026, BCP 9, October 1996.

The actual stages in the standards track are:

1. *Proposed Standard*. The formal entry point for all specifications is here as a Proposed Standard. This is the beginning of the RFC process. The IESG has authority via the RFC editor to elevate an Internet-Draft to this level. While no prior real-world implementation is required of a Proposed Standard, these specifications are generally expected to be fully formulated and implementable.

2. *Draft Standard*. After specifications have been implemented in at least two “independent and interoperable” real-world applications, they can be elevated to the level of a Draft Standard. A specification at the Draft Standard level must be relatively stable and easy to understand. While subtle revisions are normal for Draft Standards, no substantive changes are expected after this level.

3. *Standard*. Robust specifications with wide implementation and a proven track record are elevated to the level of Standard. They are considered to be official Internet Standards and are given a new number in the “STD” subseries of the RFCs (but also retain their RFC number). The total number of Standards is relatively small.

Not all RFCs are standards. Many RFCs are informational, experimental, historic, or even humorous<sup>33</sup> in nature. Furthermore, not all RFCs are full-fledged Standards; they may not be that far along yet.

---

33. Most RFCs published on April 1 are suspect. Take, for example, RFC 1149, “A Standard for the Transmission of IP Datagrams on Avian Carriers” (David Waitzman, April 1990), which describes how to send IP datagrams via carrier pigeon, lauding their “intrinsic collision avoidance system.” Thanks to Jonah Brucker-Cohen for first bringing this RFC to my attention. Brucker-Cohen himself has devised a new protocol called “H2O/IP” for the transmission of IP datagrams using modulated streams of water. Consider also “The Infinite Monkey Protocol Suite (IMPS)” described in RFC 2795 (SteQven [sic] Christey, April 2000), which describes “a protocol suite which supports an infinite number of monkeys that sit at an infinite number of typewriters in order to determine when they have either produced the entire works of William Shakespeare or a good television show.” Shakespeare would probably appreciate “SONET to Sonnet Translation” (April 1994, RFC 1605), which uses a fourteen-line decasyllabic verse to optimize data transmission over Synchronous Optical Network (SONET). There is also the self-explanatory “Hyper Text Coffee Pot Control Protocol (HTCPCP/1.0)” (Larry Masinter, RFC 2324, April 1998), clearly required reading for any sleep-deprived webmaster.

In addition to the STD subseries for Internet Standards, there are two other RFC subseries that warrant special attention: the Best Current Practice (BCP) documents and informational documents known as FYI.

Each new protocol specification is drafted in accordance with RFC 1111, “Request for Comments on Request for Comments: Instructions to RFC Authors,” which specifies guidelines, text formatting and otherwise, for drafting all RFCs. Likewise, FYI 1 (RFC 1150) titled “F.Y.I. on F.Y.I.: Introduction to the F.Y.I. Notes” outlines general formatting issues for the FYI series. Other such memos guide the composition of Internet-Drafts, as well as STDs and other documents. Useful information on drafting Internet standards is also found in RFCs 2223 and 2360.<sup>34</sup>

The standards track allows for a high level of due process. Openness, transparency, and fairness are all virtues of the standards track. Extensive public discussion is par for the course.

Some of the RFCs are extremely important. RFCs 1122 and 1123 outline all the standards that must be followed by any computer that wishes to be connected to the Internet. Representing “the consensus of a large body of technical experience and wisdom,”<sup>35</sup> these two documents outline everything from email and transferring files to the basic protocols like IP that actually move data from one place to another.

Other RFCs go into greater technical detail on a single technology. Released in September 1981, RFC 791 and RFC 793 are the two crucial documents in the creation of the Internet protocol suite TCP/IP as it exists today. In the early 1970s Robert Kahn of DARPA and Vinton Cerf of Stanford Uni-

---

Other examples of ridiculous technical standards include Eryk Salvaggio’s “Slowest Modem,” which uses the U.S. Postal Service to send data via diskette at a data transfer rate of only 0.002438095238095238095238 kb/s. He specifies that “[a]ll html links on the diskette must be set up as a href=‘mailing address’ (where ‘mailing address’ is, in fact, a mailing address).” See Eryk Salvaggio, “Free Art Games #5, 6 and 7,” *Rbizome*, September 26, 2000. See also Cory Arcangel’s “Total Asshole” file compression system that, in fact, enlarges a file exponentially in size when it is compressed.

34. See Jon Postel and Joyce Reynolds, “Instructions to RFC Authors,” RFC 2223, October 1997, and Gregor Scott, “Guide for Internet Standards Writers,” RFC 2360, BCP 22, June 1998.

35. Robert Braden, “Requirements for Internet Hosts—Communication Layers,” RFC 1122, STD 3, October 1989.

versity teamed up to create a new protocol for the intercommunication of different computer networks. In September 1973 they presented their ideas at the University of Sussex in Brighton and soon afterward finished writing the paper “A Protocol for Packet Network Intercommunication,” which was published in 1974 by the IEEE. In that same year Vint Cerf, Yogen Dalal, and Carl Sunshine published “Specification of Internet Transmission Control Program” (RFC 675), which documented details of TCP for the first time. RFC editor Jon Postel and others assisted in the final protocol design.<sup>36</sup> Eventually this new protocol was split in 1978 into a two-part system consisting of TCP and IP. (As mentioned in earlier chapters, TCP is a reliable protocol that is in charge of establishing connections and making sure packets are delivered, while IP is a connectionless protocol that is only interested in moving packets from one place to another.)

One final technology worth mentioning in the context of protocol creation is the World Wide Web. The Web emerged largely from the efforts of one man, the British computer scientist Tim Berners-Lee. During the process of developing the Web, Berners-Lee wrote both HTTP and HTML, which form the core suite of protocols used broadly today by servers and browsers to transmit and display Web pages. He also created the Web address, called a Universal Resource Identifier (URI), of which today’s “URL” is a variant: a simple, direct way for locating any resource on the Web.

As Berners-Lee describes it:

The art was to define the few basic, common rules of “protocol” that would allow one computer to talk to another, in such a way that when all computers everywhere did it, the system would thrive, not break down. For the Web, those elements were, in decreasing order of importance, universal resource identifiers (URIs), the Hypertext Transfer Protocol (HTTP), and the Hypertext Markup Language (HTML).<sup>37</sup>

So, like other protocol designers, Berners-Lee’s philosophy was to create a standard language for interoperation. By adopting his language, the computers would be able to exchange files. He continues:

---

36. Mueller, *Ruling the Root*, p. 76.

37. Tim Berners-Lee, *Weaving the Web* (New York: HarperCollins, 1999), p. 36.

What was often difficult for people to understand about the design was that there was nothing else beyond URIs, HTTP, and HTML. There was no central computer “controlling” the Web, no single network on which these protocols worked, not even an organization anywhere that “ran” the Web. The Web was not a physical “thing” that existed in a certain “place.” It was a “space” in which information could exist.<sup>38</sup>

This is also in line with other protocol scientists’ intentions—that an infoscape exists on the Net with no centralized administration or control. (But as I have pointed out, it should not be inferred that a lack of centralized control means a lack of control as such.)

Berners-Lee eventually took his ideas to the IETF and published “Universal Resource Identifiers in WWW” (RFC 1630) in 1994. This memo describes the correct technique for creating and decoding URIs for use on the Web. But, Berners-Lee admitted, “the IETF route didn’t seem to be working.”<sup>39</sup>

Instead he established a separate standards group in October 1994 called the World Wide Web Consortium (W3C). “I wanted the consortium to run on an open process like the IETF’s,” Berners-Lee remembers, “but one that was quicker and more efficient. . . . Like the IETF, W3C would develop open technical specifications. Unlike the IETF, W3C would have a small full-time staff to help design and develop the code where necessary. Like industry consortia, W3C would represent the power and authority of millions of developers, researchers, and users. And like its member research institutions, it would leverage the most recent advances in information technology.”<sup>40</sup>

The W3C creates the specifications for Web technologies and releases “recommendations” and other technical reports. The design philosophies driving the W3C are similar to those at the IETF and other standards bodies. They promote a distributed (their word is “decentralized”) architecture, they promote interoperability in and among different protocols and different end systems, and so on.

In many ways the core protocols of the Internet had their development heyday in the 1980s. But Web protocols are experiencing explosive growth

---

38. Berners-Lee, *Weaving the Web*, p. 36.

39. Berners-Lee, *Weaving the Web*, p. 71.

40. Berners-Lee, *Weaving the Web*, pp. 92, 94.

today. Current growth is due to an evolution of the concept of the Web into what Berners-Lee calls the Semantic Web. In the Semantic Web, information is not simply interconnected on the Internet using links and graphical markup—what he calls “a space in which information could permanently exist and be referred to”<sup>41</sup>—but it is enriched using descriptive protocols that say what the information actually is.

For example, the word “Galloway” is meaningless to a machine. It is just a piece of information that says nothing about what it is or what it means. But wrapped inside a descriptive protocol it can be effectively parsed: “<surname>Galloway</surname>.” Now the machine knows that Galloway is a surname. The word has been enriched with semantic value. By making the descriptive protocols more complex, one is able to say more complex things about information, namely, that Galloway is *my* surname, and my given name is Alexander, and so on. The Semantic Web is simply the process of adding extra metalayers on top of information so that it can be parsed according to its semantic value.

Why is this significant? Before this, protocol had very little to do with meaningful information. Protocol does not interface with content, with semantic value. It is, as I have said, against interpretation. But with Berners-Lee comes a new strain of protocol: protocol that cares about meaning. This is what he means by a Semantic Web. It is, as he says, “machine-understandable information.”

Does the Semantic Web, then, contradict my earlier principle that protocol is against interpretation? I’m not so sure. Protocols can certainly *say* things about their contents. A checksum does this. A file-size variable does this. But do they actually know the meaning of their contents? So it is a matter of debate as to whether descriptive protocols actually add intelligence to information, or whether they are simply subjective descriptions (originally written by a human) that computers mimic but understand little about. Berners-Lee himself stresses that the Semantic Web is not an artificial intelligence machine.<sup>42</sup> He calls it “well-defined” data, not interpreted data—and

---

41. Berners-Lee, *Weaving the Web*, p. 18.

42. Tim Berners-Lee, “What the Semantic Web Can Represent,” available online at <http://www.w3.org/DesignIssues/RDFnot.html>.

in reality those are two very different things. I promised in the introduction to skip all epistemological questions, and so I leave this one to be debated by others.

As this survey of protocological institutionalization shows, the primary source materials for any protocological analysis of Internet standards are the RFC memos. They began circulation in 1969 with Steve Crocker's RFC "Host Software" and have documented all developments in protocol since.<sup>43</sup> "It was a modest and entirely forgettable memo," Crocker remembers, "but it has significance because it was part of a broad initiative whose impact is still with us today."<sup>44</sup>

While generally opposed to the center-periphery model of communication—what some call the "downstream paradigm"<sup>45</sup>—Internet protocols describe all manner of computer-mediated communication over networks. There are RFCs for transporting messages from one place to another, and others for making sure it gets there in one piece. There are RFCs for email, for webpages, for news wires, and for graphic design.

Some advertise distributed architectures (like IP routing), others hierarchical (like DNS). Yet they all create the conditions for technological innovation based on a goal of standardization and organization. It is a peculiar type of anti-federalism through universalism—strange as it sounds—whereby

---

43. One should not tie Crocker's memo to the beginning of protocol per se. That honor should probably go to Paul Baran's 1964 Rand publication "On Distributed Communications." In many ways it served as the origin text for the RFCs that followed. Although it came before the RFCs and was not connected to it in any way, Baran's memo essentially fulfilled the same function, that is, to outline for Baran's peers a broad technological standard for digital communication over networks.

Other RFC-like documents have also been important in the technical development of networking. The Internet Experiment Notes (IENs), published from 1977 to 1982 and edited by RFC editor Jon Postel, addressed issues connected to the then-fledgling Internet before merging with the RFC series. Vint Cerf also cites the ARPA Satellite System Notes and the PRNET Notes on packet radio (see RFC 2555). There exists also the MIL-STD series maintained by the Department of Defense. Some of the MIL-STDs overlap with Internet Standards covered in the RFC series.

44. Steve Crocker, "30 Years of RFCs," RFC 2555, April 7, 1999.

45. See Minar and Hedlund, "A Network of Peers," p. 10.

universal techniques are levied in such a way as ultimately to revert much decision making back to the local level.

But during this process many local differences are elided in favor of universal consistencies. For example, protocols like HTML were specifically designed to allow for radical deviation in screen resolution, browser type, and so on. And HTML (along with protocol as a whole) acts as a strict standardizing mechanism that homogenizes these deviations under the umbrella of a unilateral standard.

Ironically, then, the Internet protocols that help engender a distributed system of organization are themselves underpinned by a distributed, bureaucratic institutions—be they entities like ICANN or technologies like DNS.

Thus it is an oversight for theorists like Lawrence Lessig (despite his strengths) to suggest that the origin of Internet communication was one of total freedom and lack of control.<sup>46</sup> Instead, it is clear to me that the exact

---

46. In his first book, *Code and Other Laws of Cyberspace* (New York: Basic Books, 1999), Lessig sets up a before/after scenario for cyberspace. The “before” refers to what he calls the “promise of freedom” (p. 6). The “after” is more ominous. Although as yet unfixed, this future is threatened by “an architecture that perfects control” (6). He continues this before/after narrative in *The Future of Ideas: The Fate of the Commons in a Connected World* (New York: Random House, 2001), where he assumes that the network, in its nascent form, was what he calls free—that is, characterized by “an inability to control” (p. 147). Yet “[t]his architecture is now changing” (p. 239), Lessig claims. The world is about to “embrace an architecture of control” (p. 268) put in place by new commercial and legal concerns.

Lessig’s discourse is always about a process of becoming, not of always having been. It is certainly correct for him to note that new capitalistic and juridical mandates are sculpting network communications in ugly new ways. But what is lacking in Lessig’s work, then, is the recognition that control is endemic to all distributed networks that are governed by protocol. Control was there from day one. It was not imported later by the corporations and courts. In fact distributed networks *must* establish a system of control, which I call protocol, in order to function properly. In this sense, computer networks are and always have been the exact opposite of Lessig’s “inability to control.”

While Lessig and I clearly come to very different conclusions, I attribute this largely to the fact that we have different objects of study. His are largely issues of governance and commerce while mine are technical and formal issues. My criticism of Lessig is less to deride his contribution, which is inspiring, than to point out our different approaches.

opposite of freedom—that is, control—has been the outcome of the last forty years of developments in networked communications. The founding principle of the Net is control, not freedom. *Control has existed from the beginning.*

Perhaps it is a different type of control than we are used to seeing. It is a type of control based on openness, inclusion, universalism, and flexibility. It is control borne from high degrees of technical organization (protocol), not this or that limitation on individual freedom or decision making (fascism).

Thus it is with complete sincerity that Berners-Lee writes: “I had (and still have) a dream that the web could be less of a television channel and more of an interactive sea of shared knowledge. I imagine it immersing us as a warm, friendly environment made of the things we and our friends have seen, heard, believe or have figured out.”<sup>47</sup> The irony is, of course, that in order to achieve this social utopia computer scientists like Berners-Lee had to develop the most highly controlled and extensive mass media yet known. Protocol gives us the ability to build a “warm, friendly” technological space. But it becomes warm and friendly through technical standardization, agreement, organized implementation, broad (sometimes universal) adoption, and directed participation.

I stated in the introduction that protocol is based on a *contradiction* between two opposing machines, one machine that radically distributes control into autonomous locales, and another that focuses control into rigidly defined hierarchies. This chapter illustrates this reality in full detail. The generative contradiction that lies at the very heart of protocol is that *in order to be politically progressive, protocol must be partially reactionary.*

To put it another way, in order for protocol to enable radically distributed communications between autonomous entities, it must employ a strategy of universalization, and of homogeneity. It must be anti-diversity. It must promote standardization in order to enable openness. It must organize peer groups into bureaucracies like the IETF in order to create free technologies.

To be sure, the two partners in this delicate two-step often exist in separate arenas. As protocol pioneer Bob Braden puts it, “There are several vital

---

47. Cited in Jeremie Miller, “Jabber,” in *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, ed. Andy Oram (Sebastapol: O’Reilly, 2001), p. 81.

kinds of heterogeneity.”<sup>48</sup> That is to say, one sector can be standardized while another is heterogeneous. The core Internet protocols can be highly controlled while the actual administration of the Net can be highly uncontrolled. Or, DNS can be arranged in a strict hierarchy while users’ actual experience of the Net can be highly distributed.

In short, control in distributed networks is not monolithic. It proceeds in multiple, parallel, contradictory, and often unpredictable ways. It is a complex of interrelated currents and counter-currents.

Perhaps I can term the institutional frameworks mentioned in this chapter a type of *tactical standardization*, in which certain short-term goals are necessary in order to realize one’s longer-term goals. Standardization is the politically reactionary tactic that enables radical openness. Or to give an example of this analogy in technical terms: DNS, with its hierarchical architecture and bureaucratic governance, is the politically reactionary tactic that enables the truly distributed and open architecture of the Internet Protocol. It is, as Barthes put it, our “Operation Margarine.” And this is the generative contradiction that fuels the Net.

---

48. Bob Braden, personal correspondence, December 25, 2002.