'Rough Consensus and Running Code' and the Internet-OSI Standards War

Andrew L. Russell

The Johns Hopkins University

Internet historians recognize the technical achievements but often overlook the bureaucratic innovations of Internet pioneers. The phrase, "We reject: kings, presidents, and voting. We believe in: rough consensus and running code," was coined by David Clark in 1992. This article explains how the phrase captured the technical and political values of Internet engineers during a crucial phase in the Internet's growth.

Historians interested in conflict and consensus in technological systems stand much to gain from examining standardization—the processes of creating and implementing technical standards. Standardization requires sophisticated technical work as well as organizational and strategic coordination. In the late 1970s, it was reasonable to assume that formal standards-setting institutions, such as the International Telecommunication Union (ITU) and the International Organization for Standardization (ISO), would lead the coordination and standardization of information services, much as they did for international telecommunications. Beginning in 1977, ISO oversaw a large, ambitious attempt to define a network architecture called Open Systems Interconnection (OSI). Between the late 1970s and the early 1990s, the OSI seven-layer model became enshrined in computer science curricula and endorsed by governments around the world. Competing networks-including experimental TCP/IP networks-were expected to fade away once OSI protocols were standardized and implemented by users and manufacturers.

By the mid-1990s, however, ISO's slow standardization process had failed to keep up with alternative, informal mechanisms that were more effective at coordinating rapid technological change.¹ These informal mechanisms the focus of this article—provided vital institutional support for the eventual success and global deployment of the Internet architecture and Internet standards such as TCP/IP.

In most existing histories of the Internet, leaders such as Vinton Cerf and Robert Kahn receive plaudits for their technical work, especially for their roles in the creation of the core Internet standards TCP and IP.² What is often overlooked, however, is the extent to which the success of the Internet depended on organizational innovations directed by Cerf, Kahn, and other Internet pioneers such as Jon Postel and David Clark. Starting in the 1970s, researchers in groups such as the Internet Configuration Control Board (ICCB), the Internet Activities Board (IAB), and the Internet Engineering Task Force (IETF) coordinated Internet standards and architectural development. These institutions deserve special attention not only for their technical achievements, but also because they fostered a voluntary international consensus during a period of intense technical and institutional change in computing and telecommunications.

These organizational innovations emerged as responses to external competition from ISO as well as to internal strains in the processes of Internet standardization. As the Internet grew rapidly in the late 1980s and 1990s, groups such as the IAB and IETF started to become victims of their own success and struggled to preserve their founding principles. If we understand "politics" to mean relations of control between individuals and groups, the political values of Internet architects and engineers were especially evident as these people built political structures—standards bodies—from scratch.³

Standardization is technically and organizationally complex as well as deeply value-laden.⁴ Leaders and participants in the IAB and IETF articulated institutional rules, cultural traditions, and versions of their own history as they responded to challenges from within and without. A memorable phrase encapsulates the prevailing technical and organizational values of those who were involved in the Internet standards process from the mid-1970s to the present: "We reject: kings, presidents, and voting. We believe in: rough consensus and running code."⁵ This article explains how this phrase which became a motto for Internet standardization—articulated a common set of beliefs about the work culture and engineering style of Internet standardization.

In 1999, law professor Lawrence Lessig declared that "rough consensus and running code" had broad significance as "a manifesto that will define our generation."6 An examination of the origins of this credo-coined by David Clark in 1992-illustrates some technical, rhetorical, and philosophical differences between the Internet standards process and the competing ISO process. Most participants in the Internet standards process today consider "rough consensus and running code" to be a succinct and accurate description of the Internet standards process.⁷ For historians, "rough consensus and running code" stands as a revealing depiction of the international politics of computer internetworking in the last quarter of the 20th century, as well as a point of entry for investigating why contemporaries described the competition between the Internet and OSI as "the Internet religious war."8

Internet architecture: TCP/IP and the end-to-end argument

In 1996, the Internet Architecture Board (IAB) published RFC 1958, "Architectural Principles of the Internet," to record aspects of architectural approach practiced in the Internet community over the previous 25 years. They wrote: "in very general terms, the community believes that the goal is connectivity, the tool is the Internet Protocol, and the intelligence is end to end rather than hidden in the network."⁹ This section provides a brief overview of the distinctive aspects of Internet architecture by exploring these three principles: interconnection, the Internet Protocol, and the "end-to-end" argument.

After initial networking experiments with the Arpanet beginning in 1969, the Advanced Projects Research Agency (ARPA) continued to explore networking concepts for satellite and radio. The problem of how to enable communication between technically disparate systems also had implications for mobile military communications. In the early 1970s, ARPA developed a packet radio network (known as PRnet) based on the same packet-switching technology being tested in the Arpanet. Robert Kahn, as program manager for this project, identified the need to connect packet radio networks to large computers in the continental US via the Arpanet (PRnet was developed in Hawaii), but faced a problem of trying to connect what he later called "two radically different networks" with different network capacities, protocols, and routing systems.¹⁰ ARPA's concurrent development of satellite packet switching in SATnet further compounded this problem, leading Kahn to conclude that network interconnection could not be achieved on an ad hoc, network-to-network level.

To overcome this problem, Kahn in 1973 rekindled an effective partnership with Vint Cerf-the two had worked together on the first nodes of the Arpanet in 1969-and proposed a new way to transport data packets. This mechanism was a simple technical protocol (transmission control protocol, or TCP) and system of gateways (now known as routers) to transfer data packets between the dissimilar networks.¹¹ Kahn's ultimate goal was to make the network transparent, "a seamless whole," invisible to the user who would be more interested in accessing information over the network instead of the operation of the network itself.¹² This basic principle-interconnection via standard protocols-is the Internet's fundamental structure and defining feature.¹³ In 1978, Cerf and two other DARPA-funded network researchers, Danny Cohen and Jon Postel, split the functions of TCP into two protocols, TCP and the Internet Protocol (IP), that worked together in the now-familiar combination of TCP/IP.14

Network transparency and application autonomy are the two key elements of a design philosophy first articulated in 1981 by David Clark, Jerome Saltzer, and David Reed, three veterans of ARPA-sponsored networking experiments at the Massachusetts Institute of Technology's (MIT's) Laboratory for Computer Science. Their paper, "End-To-End Arguments in System Design," was written to explore a design principle that they claimed had "been used for many years with[out] explicit recognition."15 An outgrowth of substantial experience with TCP/IP networking, the end-to-end principle held that the Internet's complex functions should be performed at the endpoints, leaving only the (relatively) simple tasks of interconnection and data transport to the network. The authors commended the technical merits and simplicity of the end-toend model by concluding, "end-to-end arguments are a kind of 'Occam's razor' when it comes to choosing the functions to be provided in a communications subsystem."16

The end-to-end design principle thus calls for a simple standard for data transfer that allows new innovations to be added at the edges and on top of the "stupid network."¹⁷ From this distributed design follows decentralized control, both over the functionality of the network and over the content of the network traffic. Marjory Blumenthal and David Clark argued in 2001 that end-to-end principles form the heart of the

conventional understanding of the "Internet philosophy": freedom of action, user empowerment, end-user responsibility for actions undertaken, and lack of controls "in" the Net that limit or regulate what users can do.¹⁸

By building in minimum control and allowing the "broadest range of development" to end users, the technical standards of the Internet have, as Lessig wrote, "erected a free space of innovation."¹⁹

On 1 January 1983, TCP/IP replaced NCP (an older networking protocol) as the standard host protocol for the Arpanet.²⁰ "With the use of TCP/IP now mandatory on the Arpanet," Kahn recalled in 1994, "the rationale for connection of local area networks was evident and the growth of users and networks accelerated. It also led to a rethinking of the process that ARPA was using to manage the evolution of the network."²¹ The next section examines the evolution of this management process as the Internet community faced the problems and challenges that accompanied the growth of the network.

Governance of Internet protocols, 1979–1992

Internet researchers created and presided over the governing bodies of Internet protocols—the ICCB, IAB, and IETF—to manage the growth of end-to-end networks based on TCP/IP. The history of these groups reveals the creation of increasingly formal structures by a core of Internet researchers who strove to preserve the design principles and work culture that had fostered the Internet's successful growth.

Cerf's involvement with ARPA started at Stanford University, where he worked from 1974 to 1976 to implement TCP/IP. Kahn, who joined ARPA's Information Processing Techniques Office as a program manager in 1972, hired Cerf to come to ARPA in 1976 to run a program to pursue and coordinate Internet research. Cerf later recalled that, upon moving to ARPA, "The single-minded goal was to get the Internet system up and running."²²

When he began at ARPA, Cerf worked with "only a few researchers" to "develop and test versions of the internet protocols."²³ Kahn, Cerf's manager at the time, recalled the deci-

sion formally to expand the involvement of the research community in making decisions about the network:

[W]hen we started the internet program in the mid 1970s, originally it was just me in the office running the program. And after Vint was hired, then it was just Vint running the program with me to kibitz. And he was so good at what he did that he basically had everything in his head. What I worried about was what would happen if he got hit by a truck? Number two, what would happen if he would ever have to leave? And number three, how was anybody else in the community ever going to be part of the thinking process. So he set up, after some discussions, a kind of kitchen cabinet, if you will, of knowledgeable people that he would convene periodically. These were mostly the workers in the field, the key people who were implementing protocols . . . When he left, that group stayed intact.²⁴

This group, Cerf's "kitchen cabinet," was the ICCB, created in 1979 and "chaired by David Clark from MIT to assist ARPA in the planning and execution of the evolution of the TCP/IP suite."²⁵ The ICCB expanded control over Internet development by bringing more of the network users—technical experts distributed in universities, firms, and government agencies—into Cerf's inner circle in a more formal way. For Kahn, the ICCB was important because it "brought a wider segment of the research community more directly into the decision-making aspects of the Internet project which, until then, had been almost solely undertaken by ARPA."²⁶

Cerf left his position as head of the Internet research program at ARPA in 1982, and was replaced in 1983 by Barry Leiner. Leiner and Clark, in response to the continued growth of the Internet community, disbanded the ICCB in September 1984 and created the Internet Advisory Board (IAB).²⁷ Clark continued his close involvement as the first chair of the IABa position that confers the title of "Internet Architect."28 Leiner created task forces within the IAB to focus on specific aspects of Internet technology (such as gateway algorithms, endto-end protocols, and security) in an attempt to keep discussions focused while accommodating growing numbers of participants.²⁹ The Board itself consisted of the chairs of the task forces. There were no elections in the IAB. Instead, as Vint Cerf explained in 1989, new members were appointed by the chairman of the IAB, with the advice and consent of the remaining members. Cerf wrote,

[m]embership changes with time to adjust to the current realities of the research interests of the participants, the needs of the Internet system and the concerns of the constituent members of the Internet.²³

Given this description, it is not difficult to see why the author and computer scientist Ed Krol described the IAB as a "council of elders."³⁰ The IAB cannot be characterized as a democracy, since nobody voted and the Board only let in the people they wanted. The very premise of the Internet-especially its protocol and network design-required that the Board make its decisions by consensus. Democracy, with its competing factions and its political compromises, was not an appropriate political model for the IAB or the Internet. Instead, the IAB operated with a leadership of experienced technicians and a rank and file organized by area of technical interest, more technocratic than democratic. The IAB increasingly served as the steward of TCP/IP and the Internet, but had no legal mandate or enforcement mechanisms. In other words, IAB-backed protocols were de facto standards, whose status as standards depended on broad consensus and widespread implementation.

One of the IAB's task forces, the Internet Engineering Task Force, first met in 1986. Due to the tremendous growth in the Internet's practical and engineering side, there was soon "an explosion in the attendance at IETF meetings" that compelled IETF chair Phill Gross to create a substructure for the group. In 1987, the IETF formed separate working groups to oversee specific topics; in 1989 the IETF organized the working groups into areas and designated area directors, who formed the Internet Engineering Steering Group (IESG). The IETF's effectiveness was underscored that same year, when the IAB "recognized the importance of the IETF, and restructured the standards process explicitly to recognize the IESG as the major review body for standards."31 Having assigned responsibility to the IETF for the short-term engineering of the Internet, the IAB streamlined the rest of its task forces into the Internet Research Task Force with small research groups, such as the End-to-End Research Group and Internet Architecture Task Force, dedicated to long-term issues in the evolution of the Internet.³²

By the early 1990s, participation at IETF meetings—which were held three times a year and open to anyone interested—continued to increase at an explosive pace, reflecting growing interest from the research community as well as from the commercial community. The

List of Acronyms	
ARPA	Advanced Research Projects Agency
CLNP	ConnectionLess Network Protocol
DARPA	Defense Advanced Research Projects Agency
DoD	Department of Defense
GOSIP	Government Open Systems Interconnection Profile
IAB	Internet Advisory Board (1984–1986); Internet Activities
	Board (1986–1992); Internet Architecture Board (July
	1992–present)
ICCB	Internet Configuration Control Board
IESG	Internet Engineering Steering Group
IETF	Internet Engineering Task Force
IPv4	Internet Protocol version 4
ISO	International Organization for Standardization
ISOC	Internet Society
ITU	International Telecommunication Union
NCP	Network Control Program
NIST	National Institute of Standards and Technology
OSI	Open Systems Interconnection
PRnet	Packet Radio Network
RFC	Request for Comments
TCP/IP	Transmission Control Protocol/Internet Protocol
TP-4	Transport Protocol Class 4

increasing size, value, and internationalization of the IETF and Internet standards process brought significant legal and practical problems to the fore, such as antitrust liability, copyright protection, and the needs to detach the process from the US government and to accommodate international participation.³³ In 1990, Cerf began to formulate the idea of a private board of overseers that would act as a legal and organizational umbrella for the IAB and IETF and, at the same time, serve as a global coordinating mechanism for promoting the Internet. This umbrella, the Internet Society, was launched as a private, professional membership organization in January 1992, with Cerf as its president. In July 1992, the IAB changed its name from the Internet Activities Board to the Internet Architecture Board and became part of the Internet Society.³⁴

A "Brief History of the Internet"—coauthored by a number of Internet pioneers (including Cerf, Clark, Kahn, and Leiner) describes this formation of community governance as a "steady evolution of organizational structures designed to support and facilitate an ever-increasing community working collaboratively on Internet issues."³⁵ These structures combined two models of governance. The first model, the structure led by Cerf that coordinated the development of the Internet at ARPA, was a self-selected, experienced group—a

"council of elders." Historians of the Internet unequivocally praise this group as a source of its astounding growth. Frequently described as a meritocracy, this close-knit network of people worked together since the early Arpanet days (many as graduate students at MIT or the University of California, Los Angeles, or as engineers at Bolt Beranek and Newman, the consulting firm that had designed aspects of the Arpanet) and provided the bulk of the technical and bureaucratic leadership necessary to keep the Internet up and running.³⁶ (For more on BBN, see the special issues of Annals: vol. 27, no. 2 and vol. 28, no. 1.) "Relatively few, competent, highly motivated people were involved," recalled Larry Press, "and they had considerable autonomy."37

Kahn's observation about the ICCB's role as a sort of kitchen cabinet provides an insight into the second model of governance, which has received less attention from historians: the function of the ICCB, IAB, and IETF as mechanisms for engaging and directing the efforts of a distributed group of Internet researchers. The Internet user community was small enough in 1979 that the ICCB functioned simultaneously as both a "council of elders" and a "grassroots mechanism" for Internet standards.³⁸ By 1992, the IAB and Internet Society maintained the character of the council of elders. responsible for architectural and bureaucratic oversight, while the IETF, as its name implies, assumed responsibility for the distributed, hands-on tasks involved in the engineering and implementation of protocols and provided a forum for interested newcomers.

The transfer of responsibility for technical standards from the ICCB to the IAB and then to the IETF demonstrated a strong desire on the part of the council of elders to engage and empower the broader community that wanted to contribute to the further development of the Internet. Under this system, the Internet's architectural oversight remained with the reconstituted IAB, while the efforts of participants in the IETF were channeled toward the technical aspects of protocol development and implementation.

OSI in Europe and GOSIP in the US

This development of TCP/IP networks within the ICCB, IAB, and IETF occurred while other groups developed competing protocol architectures for computer internetworking. Vendors such as IBM and DEC (Digital Equipment Corporation) offered proprietary networking solutions that were difficult if not impossible to connect to other networks. Both the TCP/IP Internet and the open systems interconnection (OSI) model were designed as nonproprietary networks that would allow users more flexibility. Although they shared some common architectural features, these two models—the Internet and OSI—emerged within different organizational contexts under somewhat different motivations. This section describes the ideas behind the development of OSI, and its adoption by the US Department of Defense (DoD).³⁹

OSI is an example of a broader movement toward "open systems" that "encouraged compatibility between systems from different manufacturers."⁴⁰ The emergence of many types of open systems in the 1970s and 1980s (such as for stereo components and microcomputers) helped smaller producers and consumers who did not want to be locked in to proprietary products from a single manufacturer.⁴¹ For computer networks, open systems such as OSI and the TCP/IP Internet emerged as alternatives that could challenge the dominance of IBM and its System Network Architecture. These various standards coexisted throughout the late 1970s and 1980s, as networking executives and engineers developed strategies and organizational structures to facilitate network interconnection.⁴² In short, open systems articulated a network interconnection strategy that at the same time facilitated a more decentralized industry structure.

OSI dates from 1977, when, according to one of its protagonists, the ISO "recognized the special and urgent need for standards for heterogeneous informatic networks" and formed a subcommittee to create a Reference Model of Open Systems Interconnection.⁴³ The goal of this seven-layer OSI Reference Model was not to define the internal operations of networks, but only to standardize the external interfaces between networks: in other words, to set the ground rules for network interconnection.⁴⁴

During the 1980s and early 1990s, OSI enjoyed widespread support from national governments, particularly in Western Europe, North America, and the Far East.⁴⁵ OSI enjoyed this level of support due in part to the strategic position of its sponsor, ISO. ISO was an "official" international standards body, meaning that it was populated by representatives from national governments who, in most cases, acted on behalf of the interests of their national telecommunications and computer firms.⁴⁶ ISO's organizational culture—concerned with defining and controlling the future of information and telecommunication services on behalf of its representatives from national governments—resembled contemporary democratic bodies insofar as it featured voting, partisan compromises, and rule-making behavior designed to protect financial interests. Such processes stand in stark contrast to the research and military orientation of the people and institutions that developed Internet protocols.

The international support for OSI influenced the officials in the DoD who were responsible for procuring equipment for their own computer networks.47 Even though they had sponsored the development of Internet standards, DoD managers believed in the early 1980s that OSI networks were likely to emerge as de facto and de jure global standards. To better understand the competing standards, the DoD asked the National Research Council in 1983 to evaluate TCP and TP-4, its functional counterpart in the OSI Reference Model. The final 1985 report presented three options: keep the two as costandards; adopt TP-4 as soon as it was shown to be ready for military networks; or keep TCP and defer indefinitely a decision on TP-4.48 The DoD supported the second option and planned to "move ultimately toward exclusive use of TP-4."49

OSI continued to gain momentum in August 1988, when the National Bureau of Standards published a version of OSI for US federal agencies called GOSIP (for "Government Open Systems Interconnection Profile") version 1, which was built around TP-4 and many OSI specifications. By August 1990, federal agencies were required to procure GOSIP-compliant products.⁵⁰ Through this procurement requirement, the government intended to stimulate the market for OSI products.⁵¹ However, many network administrators resisted the GOSIP procurement policy and continued to operate TCP/IP networks, noting that the federal mandate, by specifying only procurement, did not prohibit the use of products built around the more familiar and more readily available TCP/IP.52

'OSI Bigots' and 'IP Bigots': Cultural dimensions of a standards war

A spirited rivalry between respective advocates of OSI and TCP/IP networks emerged as they fought for jurisdiction over standards for computer internetworks. Richard des Jardins an early contributor to the ISO Reference Model and president of the GOSIP Institute captured the intensity of this rivalry when, in a 1992 article, he compared the "OSI Bigots" and the "IP Bigots" to people who objected to "the convergence of cultures and races in the world at large."⁵³ OSI continued to gain momentum in August 1988, when the National Bureau of Standards published a version of OSI for US federal agencies called GOSIP.

Tensions between the OSI and Internet communities became apparent in the early 1980s, well before GOSIP came into being. For example, in their 1983 paper describing the similarities between the ARPA and ISO protocol architectures, Danny Cohen and Jon Postel painted the ISO model as an abstraction, far too rigid in its reliance on seven interrelated levels, and inappropriate to be used "as a model for all seasons."54 In an unusually colorful paper, Cohen and Postel-both of whom were instrumental in the early history of TCP/IP-mockingly speculated that "mystical" traditions such as Early Zoroastrianism, New Testament celestial beings, and the Christian seven deadly sins might have "shaped the choice of Seven."55 Another Internet advocate, in his 1991 "technical travelogue" of networking in 21 countries across the world, suggested that trying to implement OSI over slow, low-quality lines was "akin to looking for a hippopotamus capable of doing the limbo."56

The resentment of Cohen, Postel, and their Internet colleagues stemmed from their frustration with the technical aspects of OSI as well as with ISO as a bureaucratic entity. Where TCP/IP was developed through continual experimentation in a fluid organizational setting, Internet engineers viewed OSI committees as overly bureaucratic and out of touch with existing networks and computers. OSI's political and formal process did not endear the TCP/IP Internet community-who were accustomed to a decentralized division of labor throughout the standards process-to the ISO Reference Model. In a scathing 1985 critique of OSI and its advocates, one veteran of the Arpanet and Internet community, Mike A. Padlipsky, characterized the ARPA Internet Reference Model as "Descriptive" and ISO Reference Model as "Prescriptive." Another networking pioneer, David Mills, agreed in a 2004 interview: "Internet standards tended to be those written

Apart from the external OSI threat, the Internet community faced many internal problems throughout the Internet's rapid growth during the 1980s and early 1990s.

for implementers. International standards were written as documents to be obeyed."⁵⁷ "Another way of putting it," Padlipsky wrote, "is that whereas the Descriptive approach is suitable for technology, the Prescriptive approach is suitable for theology."⁵⁸

Apart from the external OSI threat, the Internet community faced many internal problems throughout the Internet's rapid growth during the 1980s and early 1990s. As Internet advocates battled against OSI, they also continued to struggle with the organizational problems of their own standardization process. Some critics felt that the IAB "failed at times to provide a solid agenda and timetables of engineering problems" for the IETF to address.⁵⁹ The informal character of the IAB's oversight of the IETF had created problems in the past, especially when IETF engineers perceived that IAB decisions favored the commercial interests of vendors over the technical consensus of the IETF.⁶⁰

Tensions between the IAB and IETF exacerbated a disagreement about a major technical obstacle to Internet growth. One of the chief concerns of Internet architects in the late 1980s revolved around addressing and routing problems built into the current version of TCP/IP (IPv4): the finite amount of address space in IPv4 was projected to be running out quickly. If the exponential growth of Internet users continued, the bottleneck would prevent new connections to the network.⁶¹ The IAB perceived that a solution might be reached through the OSI functional counterpart to IP called ConnectionLess Network Protocol, or CLNP.⁶²

Although it was aware of strong opposition to OSI within the Internet community, the IAB felt that working with CLNP could help the Internet overcome the address space problem. From this perspective, the technical rationale for incorporating CLNP into the architecture of the Internet supported the community's broader interests and the IAB's mandate— keeping the network open for anyone who wanted to connect.⁶³

Internet's constitutional crisis

During its June 1992 meeting in Kobe, Japan, the IAB developed a draft discussion proposal to use CLNP as the basis for a larger address space. To the IAB, this seemed a responsible path to take, given the limits of the IPv4 address space and the desire for the Internet to accommodate as many users as possible. As IAB member Christian Huitema recalled in 1996,

The IAB discussed [the draft proposal to incorporate CLNP] extensively. In less than two weeks, it went through eight successive revisions. We thought that our wording was very careful, and we were prepared to discuss it and try to convince the Internet community. Then, everything accelerated. Some journalists got the news, an announcement was hastily written, and many members of the community felt betrayed. They perceived that we were selling the Internet to the ISO and that headquarters was simply giving the field to an enemy that they had fought for many years and eventually vanquished. The IAB had no right to make such a decision alone.⁶⁴

To the "general membership of the IETF," Carl Cargill commented, "this was rank heresy."65 Scott Bradner preferred a political metaphor to the military and religious metaphors used by Huitema and Cargill: he later referred to the events that followed as a "constitutional crisis."66 The technical discussion proposed by the IAB inspired a movement within the IETF that challenged the organizational hierarchies of Internet standards. The fact that a mere proposal—it is important to note that the IAB's CLNP discussion was in no way a final decisionwould provoke such outrage from hundreds of engineers and computer scientists reflects the passion and commitment of engineers in the pitched battle of a standards war.

Many IETF participants, while aware of the limitations of IPv4, assumed that TCP/IP and OSI "will coexist for a long time," and certainly did not anticipate that another group would attempt to change the Internet's fundamental protocol without first consulting the IETF.⁶⁷ Tradition, supported by IAB and IETF documentation, gave the IETF the right to standardize protocols. Additionally, the IETF, consisting mostly of academic and government researchers, resented that OSI was a complex and costly system, driven by the political con-

cerns of ISO—the "standards elephant."⁶⁸ This cultural conflict—which was, by 1992, over a decade old—made CLNP especially unpalatable to the IETF as a replacement for their favored Internet Protocol.

At the July 1992 IETF meeting in Cambridge, Massachusetts, irate IETF participants protested to the Internet Society about what they perceived as a unilateral decision by the IAB. Approximately 700 IETF participants demanded that the newborn Internet Society intervene and ensure that the IETF would remain in control of the standards process.⁶⁹ In the face of a massive "palace revolt," the IAB, with Vint Cerf prominent among them, relented. As Cerf addressed the IETF, he slowly removed the layers of his signature three-piece suit, performing a striptease that revealed a Tshirt: "IP on Everything." The T-shirt, according to Cerf, was to reiterate a goal of the IAB: to run IP on every underlying transmission medium.⁷⁰ Like Cerf, David Clark turned his plenary presentation into a memorable occasion, one that would "rally the troops" and reaffirm the values of the Internet community.71

'We reject: kings, presidents, and voting'

Clark framed his talk in terms of architectural choices the IETF would have to make. After spending several minutes urging the audience to focus on network security and the basic assumptions of the protocol architecture, Clark considered how the IETF should "manage the process of change and growth."5 As he reminded the IETF audience of the vital importance of the values of the process by which they made standards, Clark punctuated his discussion with his summary of the IETF approach: "We reject: kings, presidents, and voting. We believe in: rough consensus and running code."5 The IETF community responded with overwhelming approval. "Rough consensus and running code" was so popular that Marshall Rose, a vocal participant in the IETF "palace revolt," created the ultimate form of computer geek approval: Tshirts with the phrase emblazoned across the front. (On the occasion of the IETF's 20th anniversary in 2006, Clark delivered an encore of this presentation. The video of his talk is available from http://ietf20.isoc.org/videos/.)

"Rough consensus and running code" generated and sustained this level of enthusiasm because of the way it framed the successful aspects of the IETF process in opposition to the ISO process. The rough-consensus component of this motto refers to the decision-making process within IETF working groups. Since its inception, the IETF never had members, only participants, and hence it could not have a formal voting structure. In the tradition of Cerf's ICCB discussions, IETF leaders encouraged newcomers to contribute their expertise, and approved proposals that enjoyed broad support within the group. IETF veterans place an acceptable level of agreement at around 80 to 90 percent: a level high enough to demonstrate strong support, but flexible enough to work in the absence of unanimity. In short, rough consensus was an apt description of this informal process in which a proposal must answer to criticisms, but need not be held up if supported by a vast majority of the group.⁷² To IETF participants, this process was vastly superior to the bureaucratic and political approach of ISO.73

As a complement to rough consensus, running code means "multiple actual and interoperable implementations of a proposed standard must exist and be demonstrated before the proposal can be advanced along the standards track."74 Since most standards begin with a proposal from an individual or group within a working group-and not from the IAB or IETF leadership-the party behind the proposal must provide multiple working versions of the proposal. This burden of proof on the proposed standard facilitates the adoption of new IETF standards across the Internet's diverse computing platforms. Running code also evokes a major difference between the IETF and ISO approaches: where the IETF protocols represented "the result of intense implementation discussion and testing," ISO committees developed a theoretical model that was difficult to alter or implement fully.75 According to Lyman Chapin, a participant in both ISO and Internet standardization,

it didn't take long to recognize the basic irony of OSI standards development: there we were, solemnly anointing international standards for networking, and every time we needed to send electronic mail or exchange files, we were using the TCP/IP-based Internet!⁷⁶

Or, as Internet pioneer Einar Stefferud was fond of saying, "OSI is a beautiful dream, and TCP/IP is living it."⁷⁷

Beyond serving as a concise description of the IETF's organizational and technical approach, "rough consensus and running code" also served as a means of self-identification and a positive summary of the IETF's model for standards development. The internal divisions exacerbated by the controversy over CLNP prompted a good deal of reflection among those who were committed to defending the traditions of the IETF. Although most of his presentation was devoted to the pressing technical and organizational problems within the Internet standards community, Clark's memorable phrase was an attempt to unite the fractured community by contrasting it to their OSI rivals. After their successful campaigns against the IAB and the CLNP proposal, one can imagine IETF engineers leaving the July 1992 meeting with a certain sense of optimism about the future of the Internet.

Conclusion

By the time National Institute of Standards and Technology (NIST) abandoned GOSIP in favor of TCP/IP in 1994, the grand future planned for OSI was on the rapid decline.⁷⁸ The market for network protocols had tipped in favor of TCP/IP, epitomized by the popularity of a new application—the World Wide Web—that was designed to take advantage of the Internet's end-to-end architecture.⁷⁹ The veterans of the Internet-OSI standards war were no doubt wiser from the experience, but the millions of users who got on the Internet in the mid-1990s were oblivious to the fact that their new toy was the product of a protracted international struggle. For those who read the last issue of Telecommunications Policy in 1993, William Drake provided an insightful summary:

The debate is not merely about the comparative efficacy of two sets of standards, but it is rather between two competing visions of how international standardization processes and network development should be organized and controlled.⁸

The religious, political, and military metaphors that participants used to describe the competition between the Internet and OSI confirm that this was no mere technical dispute. At the height of the "religious war" between TCP/IP and OSI, tensions within the IETF and IAB over architecture and organizational power created a "constitutional crisis." This crisis, a divisive force in a community that had always prided itself on its attention to due process and consensus, forced engineers in the IETF and IAB to examine their core procedural beliefs. In other words, strains in the technical architecture-the address space-prompted strains in the organizational architecture. Forged in the face of this crisis, the credo "rough consensus and running code" articulated a political philosophy, a style of network architecture and of engineering. While it is now common to see participants in the

Internet standards community refer to the "rough consensus and running code" ideal, Clark's rejection of alternative forms of decision making—kings, presidents, and voting—reminds us of the close links between network standards and international politics.

Throughout the 1970s, 1980s, and 1990s, the open and decentralized technical architecture of the Internet embodied the technical and organizational values of its design group. Internet architecture and organizations were created at the same time, by the same people, as part of an effort "to unite the community behind a single objective—to focus the effort and guarantee the continued growth of the Internet."⁸⁰ As the evidence in this article demonstrates, defining a single goal and then uniting the community behind it was no small feat.

These organizational and procedural problems intensified throughout the 1990s, and continue to haunt the IETF and IAB—as well as the broader reaches of information technology standards-today.81 The structure and process of Internet standards set precedents and influenced the development of subsequent efforts to set standards for digital networks in institutional imitators such as the World Wide Web Consortium, the Global Grid Forum, and, most recently, the Voting Systems Performance Rating.⁸² The Internet standards community not only introduced technological innovations; it also pioneered organizational innovations such as the use of electronic mailing lists to build consensus around technical work, a process open to all interested stakeholders, and the free and unrestricted distribution of standards.

As it embodied a new style of standardization, the Internet standards community constantly dealt with problems that stemmed from the tension between centralized authority and grassroots initiatives, as well as the rising influence of commercial values. Since 1992, IETF participants in multiple working groups and mailing lists have spent tremendous amounts of energy defining formal procedures for the IETF, to the point that many in the IETF feel that their technical work is suffering from a lack of attention.⁸³ In light of its history, it seems certain that the IETF's ongoing efforts to refine and reform its structure and process will dictate the future success of its standards. Organizational and political conflicts within standards bodies define the terrain within which effective collaborations can take place. Standards wars such as the Internet-OSI conflict provide ample research opportunities for historians who want to understand computers, networks, and the people who designed and used them.

Acknowledgments

I gratefully acknowledge the close reading and thoughtful comments of Scott Bradner, Vinton G. Cerf, Stuart W. Leslie, Mark Pittenger, and Philip J. Weiser.

References and notes

- For descriptions of the features (and problems) of formal standards setting, see P.A. David and M. Shurmer, "Formal Standards-Setting for Global Telecommunications and Information Services," *Telecommunications Policy*, Oct. 1996, pp. 789-815; and W.J. Drake, "The Transformation of International Telecommunications Standardization: European and Global Dimensions," *Telecommunications in Transition: Policies, Services, and Technologies in the European Community,* C. Steinfield, J.M. Bauer and L. Caby, eds., Sage Publications, 1994, pp. 71-96.
- The best single account of Internet history is J. Abbate, *Inventing the Internet*, MIT Press, 1999. Cerf and Kahn were awarded the 2004 Turing Award for their pioneering efforts.
- 3. According to one history, the Internet's architectural principles "embody some value judgments and reflect the fundamental political and ethical beliefs of the scientists and engineers who designed the Internet"; National Research Council, *The Internet's Coming of Age*, National Academy Press, 2000, p. 35.
- See H. Nissenbaum, "How Computer Systems Embody Values," *Computer*, Mar. 2001, pp. 118-120.
- D.D. Clark, "A Cloudy Crystal Ball: Visions of the Future," plenary presentation at 24th meeting of the Internet Engineering Task Force, Cambridge, Mass., 13-17 July 1992. Slides from this presentation are available from http://ietf20.isoc.org/ videos/future ietf 92.pdf.
- L. Lessig, Code and Other Laws of Cyberspace, Basic Books, 1999, p. 4.
- On "rough consensus and running code" as IETF credo, see E. Huizer, *IETF-ISOC Relationship*, IETF RFC 2031, Oct. 1996; http://www.ietf.org/rfc/ rfc2031.txt; and S. Bradner, "The Internet Engineering Task Force," *Open Sources: Voices from the Open Source Revolution*, C. DiBona, S. Ockman, and M. Stone, eds., O'Reilly, 1999, p. 50.
- 8. W. Drake, "The Internet Religious War," *Telecommunications Policy*, Dec. 1993, p. 643.
- B. Carpenter, ed., Architectural Principles of the Internet, IETF RFC 1958, June 1996; http://www. ietf.org/rfc/rfc1958.txt.
- R.E. Kahn, oral history interview by J.E. O'Neill, 24 Apr. 1990, OH 192, Charles Babbage Inst. (CBI), Univ. of Minnesota, Minneapolis. Kahn continued, noting that "all the details were different. I don't mean conceptually they were differ-

ent. They were sort of the same genre. Just like, say, Chinese and Americans are of the same genre except one speaks Chinese and one speaks English, one lives on one side of the world, they go to sleep during your daytime, etc."

- V.G. Cerf and R.E. Kahn, "A Protocol for Packet Network Intercommunication," *IEEE Trans. Comm.*, vol. COM-22, no. 5, 1974, pp. 637-648. For an account of the decisive international contributions to TCP/IP, especially from French computer scientists such as Louis Pouzin who had developed the Cyclades network, see J. Abbate, *Inventing the Internet*, pp. 123-133.
- R. Kahn, "Testimony before the Subcommittee on Basic Research of the Committee on Science on the Subject of Internet Domain Names," 31 Mar. 1998; available from http://www.cnri. reston.va.us/testimony.html.
- 13. D.D. Clark, "The Design Philosophy of the DARPA Internet Protocols," *Proc. SIGCOMM 88, Computer Communication Rev.*, vol. 18, no. 4, 1988, pp. 106-107.
- 14. Cerf, Cohen, and Postel split the protocol into TCP and IP to reduce the requirements on network gateways and leave complex tasks such as tracking reliable packet delivery to the computers at the ends of the network. See J. Abbate, *Inventing the Internet*, pp. 129-130; V.G. Cerf, "Protocols for Interconnected Packet Networks," *Computer Communication Rev.*, vol. 18, no. 4, 1980, pp. 10-11; and J. Postel, ed., "DOD Standard Internet Protocol," IETF RFC 760, Jan. 1980; http://www.ietf.org/rfc/rfc760.txt.
- 15. J.H. Saltzer, D.P. Reed, and D.D. Clark, "End-To-End Arguments in System Design," *ACM Trans. Computer Systems*, vol. 2, no. 4, 1984, p. 277.
- 16. J.H. Saltzer, D.P. Reed, and D.D. Clark, "End-to-End Arguments," p. 287.
- 17. D.S. Isenberg, "The Rise of the Stupid Network," Computer Telephony, Aug. 1997, pp. 16-26.
- M.S. Blumenthal and D.D. Clark, "Rethinking the Design of the Internet: The End to End Arguments vs. the Brave New World," ACM Trans. Internet Technology, vol. 1, no. 1, 2001, pp. 70-109.
- 19. Lessig notes that end-to-end concepts are not unique to the Internet, and are distinguished for their ability to provide a broad range of services. Two examples are the electricity grid and the highway systems: they are simple architectures with minimal requirements for participation (standard plugs or motor vehicles) that do not constrain the behavior of the participants. L. Lessig, *The Future of Ideas: The Fate of the Commons in a Connected World*, Random House, 2001, pp. 26, 39.
- 20. For a discussion of the strategy behind this transition, see J. Postel, "NCP/TCP Transition Plan,"

IETF RFC 801, Nov. 1981; http://www.ietf.org/ rfc/rfc801.txt.

- 21. R.E. Kahn, "The Role of the Government in the Evolution of the Internet," *Comm. ACM*, vol. 37, no. 8, 1994, p. 16. Kahn was director of ARPA's Information Processing Techniques Office from August 1979 until September 1985. Because Cerf left ARPA in October 1982 and Barry Leiner did not replace him until August 1983, Kahn personally managed the transition to TCP/IP. R.E. Kahn, oral history interview, OH 192, CBI.
- 22. V.G. Cerf, oral history interview by J.E. O'Neill, 24 Apr. 1990, OH 191, CBI.
- V. Cerf, "The Internet Activities Board," IETF RFC 1160, May 1990; http://www.ietf.org/rfc/ rfc1160.txt.
- 24. R.E. Kahn oral history interview, OH 192, CBI. The term "kitchen cabinet" dates from the administration of President Andrew Jackson, who preferred to consult with an informal group of advisors—allegedly in the White House kitchen instead of his formal "Parlor" cabinet. See R.V. Remini, Andrew Jackson: The Course of American Freedom, 1822–1832, Johns Hopkins Univ. Press, 1998, pp. 315-330.
- 25. V. Cerf, as told to B. Aboba, "How the Internet Came to Be," *The Online User's Encyclopedia: Bulletin Boards and Beyond,* B. Aboba, Addison-Wesley, 1993. Kahn and Cerf note that the ICCB, unlike the IAB and IETF, functioned in a time where there was very little commercial interest in the Internet, no personal computers, and a nascent networking industry. R.E. Kahn and V.G. Cerf, "What Is The Internet (And What Makes It Work)," Dec. 1999; http://www.cnri.reston.va. us/what is internet.html.
- 26. R.E. Kahn, "The Role of the Government," p. 16.
- 27. The acronym "IAB" remained consistent since 1984, but the "A"—and the meanings behind it—have changed. From 1984 to 1986, the IAB was the Internet Advisory Board; in 1986 its name changed to the Internet Activities Board; in 1992 it changed once again, this time to the Internet Architecture Board. See "A Brief History of the Internet Advisory/Activities/Architecture Board"; http://www.iab.org/about/history.html.
- Clark was Internet Architect from 1983 to 1989; Cerf served from 1989 to 1992, and was followed by Lyman Chapin (through March 1993).
- 29. See B. Leiner et al., "A Brief History of the Internet"; http://www.isoc.org/internet/history/brief. shtml, and R.E. Kahn, oral history interview, OH 192, CBI. Task forces or working groups have been a consistent feature in the history of technical standards bodies. See C. Cargill, Open Systems Standardization: A Business Approach, Paladin Consulting, 1997.
- 30. E.Krol, "FYI on 'What is the Internet?" IETF RFC

1462, May 1993; http://www.ietf.org/rfc/ rfc1462.txt.

- See B. Leiner et al., "A Brief History of the Internet"; http://www.isoc.org/internet/history/brief. shtml. Since 1989, the number of IETF areas has shifted between seven and ten.
- B. Leiner et al., "A Brief History of the Internet"; and R. Kahn, "The Role of the Government," p. 17. For Internet Research Task Force mission and activities, see http://www.irtf.org.
- 33. See S. Bradner, "The Internet Engineering Task Force," *OnTheInternet*, vol. 7, no. 1, 2001, p. 24; and J. Abbate, *Inventing the Internet*, pp. 207-208.
- 34. V. Cerf, "IETF and ISOC," 18 July 1995; http://www.isoc.org/internet/history/ietfhis.shtm I. Kahn also worked to provide institutional support for the growth of the Internet when he founded the Corporation for National Research Initiatives in 1986. See http://cnri.reston.va.us/ about_cnri.html.
- 35. B. Leiner et al., "A Brief History of the Internet." In a 1993 article, Dave Crocker concurred: "In general, the IETF is applying its own technical design philosophy to its own operation." D. Crocker, "Making Standards the IETF Way," *StandardView*, vol. 1, no. 1, 1993, p. 54. The 1968 musings of Melvin Conway are also strikingly relevant: "There is a very close relationship between the structure of a system and the structure of the organization which designed it." See M.E. Conway, "How Do Committees Invent?" *Datamation*, vol. 14, no. 4, 1968, p. 30; and "Conway's Law," at http://www.catb.org/~esr/jargon/html/C/ Conways-Law.html.
- A.L. Norberg and J.E. O'Neill, *Transforming Computer Technology*, Johns Hopkins Univ. Press, 1996; and T.P. Hughes, *Rescuing Prometheus*, Pantheon Books, 1998, pp. 15-68.
- L. Press, "Seeding Networks: The Federal Role," Comm. ACM, vol. 39, no. 10, 1996, pp. 11-18.
- The description of the ICCB as a "grass-roots mechanism" comes from Kahn, "The Role of the Government," p. 18.
- 39. Various aspects of OSI adoption—too numerous and complex to describe in full in this article permeate Internet history. Although we lack a comprehensive historical analysis of the Internet, OSI, and other computer internetworking architectures of the period, three excellent comparisons may be found in J. Abbate, *Inventing the Internet*, pp. 147-169; National Research Council, *Global Networks and Local Values: A Comparative Look at Germany and the United States*, National Academy Press, 2001, pp. 23-45; and T. Egyedi, "'Tension between Standardisation and Flexibility' Revisited: A Critique," *Standardisation and Innovation in Information Technology: Conf.*

Proc. 1st IEEE Conf. Standardisation and Innovation in Information Technology, K. Jakobs and R. Williams, eds., IEEE Press, 1999, pp. 65-74.

- J. Abbate, "Government, Business, and the Making of the Internet," *Business History Rev.*, vol. 75, no. 1, Spring 2001, p. 167. See also M. Libicki, *Information Technology Standards: Quest for the Common Byte*, Digital Press, 1995, pp. 75-129; J.S. Quarterman and S. Wilhelm, *UNIX, POSIX, and Open Systems: The Open Standards Puzzle*, Addison-Wesley, 1993; C. Cargill, "Evolution and Revolution in Open Systems," *StandardView*, vol. 2, no. 1, 1994, pp. 3-13; and S. Schindler, "Open Systems, Today and Tomorrow—A Personal Perspective," *Computer Networks*, vol. 5, no. 3, 1981, pp. 167-176.
- R.N. Langlois, "Networks and Innovation in a Modular System: Lessons from the Microcomputer and Stereo Component Industries," *Research Policy*, vol. 21, no. 4, 1992, pp. 297-313. See also R. Sanchez and J.T. Mahoney, "Modularity, Flexibility, and Knowledge Management in Product and Organization Design," *Strategic Management J.*, vol. 17, Winter special issue, 1996, pp. 63-76.
- 42. See P.E. Green Jr., ed., Network Interconnection and Protocol Conversion, IEEE Press, 1988; and R.J. Cypser, Communicating for Cooperating Systems: OSI, SNA, and TCP/IP, Addison-Wesley, 1991. IBM executives and engineers were keenly aware of the importance of network migration and integration. See, for example, L.M. Branscomb, "Computer Communications in the Eighties-Time to Put It All Together," Computer Networks, vol. 5, no. 1, 1981, pp. 3-8; P. Janson, R. Molva, and S. Zatti, "Architectural Directions for Opening IBM Networks: The Case of OSI," IBM Systems J., vol. 31, no. 2, 1992, pp. 313-335; and T.J. Routt, "Integration Strategies for APPN and TCP/IP," Business Communications Rev., Mar. 1995, pp. 43-49.
- 43. H. Zimmerman, "OSI Reference Model—The ISO Model of Open Architecture for Open Systems Interconnection," *IEEE Trans. Comm.*, vol. COM-28, no. 4, 1980, p. 425. ISO, created in 1947, is a worldwide federation of national standards bodies.
- R. des Jardins, "Overview and Status of the ISO Reference Model of Open Systems Interconnection," *Computer Networks*, vol. 5, no. 2, 1981, pp. 77-80.
- 45. See, for example, T. Whitty, "OSI: the UK approach," Comm., vol. 7, no. 2, 1990, pp. 20-24; L. Caffrey, "EPHOS: Towards a European GOSIP," Computer Networks and ISDN Systems, vol. 19, no. 3-5, 1990, pp. 265-269; and R. Cowan, ed., Information Technology Standards: The Economic Dimension, Organisation for Economic Co-operation and Development, 1991, pp. 23-30.

- 46. For an analysis of ISO's position within the system of European standardization bodies, see S.K. Schmidt and R. Werle, Coordinating Technology: Studies in the International Standardization of Telecommunications, MIT Press, 1998, pp. 39-57.
- M. Witt, "Moving from DoD to OSI Protocols: A First Step," Computer Communication Rev., vol. 16, no. 2, 1986, pp. 2-7.
- 48. National Research Council, Transport Protocols for Department of Defense Data Networks: Report to the Department of Defense and the National Bureau of Standards Committee on Computer-Computer Communication Protocols, Board on Telecommunications and Computer Applications Commission on Eng. and Technical Systems, National Academy Press, 1985.
- J. Postel, "A DoD Statement on the NRC Report," IETF RFC 945, May 1985; http://www.ietf.org/ rfc/rfc945.txt.
- 50. "U.S. Government Open Systems Interconnection Profile," US Federal Information Processing Standards Publication 146, version 1, Aug. 1988, cited in V. Cerf and K. Mills, "Explaining the Role of GOSIP," IETF RFC 1169, Aug. 1990; http://www.ietf.org/rfc/rfc1169.txt. See also P. Janson et al., "Architectural Directions for Opening IBM Networks," p. 314 ("Many government agencies around the world, including the U.S. Department of Defense, require OSI on all systems they purchase").
- 51. Throughout the latter half of the 20th century, DoD procurement policies were effective means for driving market growth in semiconductors, computers, and software. See D.M. Hart, "Corporate Technological Capabilities and the State: A Dynamic Historical Interaction," Constructing Corporate America: Historical Perspectives on Big Business, Society, and Politics, K. Lipartito and D.B. Sicilia, eds., Oxford Univ. Press, 2004, pp. 168-187.
- 52. TCP/IP and OSI, despite the rhetoric of many of their proponents, could be engineered (with substantial effort) to work together. Since 1987, according to Cerf and NIST's Kevin Mills, there had been efforts "within the Internet community to enable integration of OSI, coexistence of OSI with TCP/IP, and interoperability between OSI and TCP/IP"; V. Cerf and K. Mills, IETF RFC 1169. See also E. Huizer, "The IETF Integrates OSI Related Work," ConneXions, vol. 7, no. 6, 1993, pp. 26-28.
- 53. R. des Jardins, "OSI is (Still) a Good Idea," ConneXions, vol. 6, no. 6, 1992, p. 33. Des Jardins added, "Let's continue to get the people of good will from both communities to work together to find the best solutions, whether they are two-letter words or three-letter words, and let's just line up the bigots against a wall and shoot them."

- D. Cohen and J. Postel, "The ISO Reference Model and Other Protocol Architectures," *Information Processing 83: Proc. IFIP 9th World Computer Congress, R.E.A. Mason, ed., North-Holland,* 1983, p. 34.
- 55. D. Cohen and J. Postel, "The ISO Reference Model," p. 30.
- 56. C. Malamud, *Exploring the Internet: A Technical Travelogue*, Prentice Hall PTR, 1992, p. 191.
- 57. D.A. Mills, personal interview with author, 26 Feb. 2004.
- 58. M.A. Padlipsky, The Elements of Networking Style, and Other Animadversions on the Art of Intercomputer Networking, iUniverse, 2000, 1st ed., Prentice Hall, 1985, p. 11. In his characteristically witty prose, Padlipsky recommended that the ISO Reference Model, or ISORM, be pronounced "Eyesore-mmm."
- 59. C. Malamud, *Exploring the Internet: A Technical Travelogue*, p. 196.
- 60. C. Malamud, *Exploring the Internet: A Technical Travelogue*, p. 151.
- 61. Published in December 1991, IETF RFC 1287 noted "increasing strains on the fundamental architecture, mostly stemming from continued Internet growth." D. Clark et al., "Towards the Future Internet Architecture," IETF RFC 1287, Dec. 1991; http://www.ietf.org/rfc/rfc1287.txt. Lyman Chapin, IAB chairman in 1992, said in May 1992 that the shortage of Internet addresses was "definitely the most significant engineering problem on the Internet now." E. Messmer, "Internet Architect Gives Long-Term View," *Network World*, vol. 9, 18 May 1992, p. 37.
- 62. For one proposal to incorporate CLNP, see R. Callon, "TCP and UDP with Bigger Addresses (TUBA), A Simple Proposal for Internet Addressing and Routing," IETF RFC 1347, June 1992; http://www.ietf.org/rfc/rfc1347.txt.
- In addition to this technical rationale, Clark suggested that ISO and the American National Standards Institute (ANSI) were pressuring the IAB to implement ISO as the Internet host protocol.
 D.D. Clark, personal communication, 27 Oct. 2001.
- 64. C. Huitema, *IPv6: The New Internet Protocol*, Prentice Hall PTR, 1998, p. 2. Minutes of the IAB discussion are available from http://www.iab.org/documents/iabmins/IABmins.1992–06–18.html. According to "The Internet Standards Process," written by L. Chapin and published as IETF RFC 1310 in March 1992, the IAB (still at this time the Internet Activities Board) delegated to the IETF "the primary responsibility for the development and review of potential Internet Standards from all sources." Internet Activities Board, L. Chapin, chair, "The Internet Standards Process," IETF RFC 1310, Mar. 1992;

http://www.ietf.org/rfc/rfc1310.txt.

- 65. C. Cargill, *Open Systems Standardization*, p. 257; S. Bradner, "The Internet Engineering Task Force," *OnTheInternet*, p. 24; V. Cerf, "IETF and ISOC"; http://www.isoc.org/internet/history/ ietfhis.shtml.
- 66. S. Bradner, "IETF," OnTheInternet, p. 24.
- 67. Chapin remarked in May 1992, "The most comprehensive solution [to the shortage of Internet addresses] is to replace the Internet Protocol in the Internet with the Open Systems Interconnection Connectionless Network Protocol. That idea is already almost universally accepted." Clearly Chapin was mistaken in his assessment of community support for CLNP. E. Messmer, "Internet Architect," p. 46.
- ISO's designation as a "standards elephant" comes from David Clark's presentation to the July 1992 IETF plenary, "A Cloudy Crystal Ball: Visions of the Future" (see ref. 5).
- 69. C. Cargill, Open Systems Standardization, p. 257.
- 70. V.G. Cerf, personal communication, 27 Jan. 2002. A photograph of Cerf wearing his shirt is available at http://ietf20.isoc.org/videos/ip_on_ everything.jpg. Beginning in 1994, the IETF developed a new version of the Internet Protocol, known as IPv6, to succeed IPv4; but for a variety of reasons beyond the scope of this article, a universal transition to IPv6 has not yet occurred. Interested readers might visit http://playground. sun.com/pub/ipng/html/.
- 71. S. Bradner, "IETF," On The Internet, p. 24.
- 72. S. Bradner, "IETF," Open Sources, pp. 47-53. See also D. Crocker, "Making Standards the IETF Way," pp. 48-54; and S. Harris, "The Tao of the IETF—A Novice's Guide to the Internet Engineering Task Force," IETF RFC 3160, Aug. 2001; http://www.ietf.org/rfc/rfc3160.txt.
- 73. The informal name of the document series in which Internet Standards are published— "Requests for Comments"-underlines the intention for Internet standards to be descriptive documents, not final and unchangeable prescriptions. IETF standards are published as RFCs, but not all RFCs are standards. See RFC Editor et al., "30 Years of RFCs," IETF RFC 2555, 7 Apr. 1999; http://www.ietf.org/rfc/rfc2555.txt ("Hiding in the history of the RFCs is the history of human institutions for achieving cooperative work"). Among the humorous and offbeat RFCs, two classics are R. Callon, ed., "The Twelve Networking Truths," IETF RFC 1925, 1 Apr. 1996; http://www.ietf.org/rfc/rfc1925.txt; and D. Waitzman, "A Standard for the Transmission of IP Datagrams on Avian Carriers," IETF RFC 1149, 1 Apr. 1990; http://www.ietf.org/rfc/rfc1149.txt.
- 74. S. Bradner, "IETF," *OnTheInternet*, p. 26. Bradner's summary is telling: "In brief, the IETF

operates in a bottom-up task creation mode and believes in 'fly before you buy.'" Bradner, "IETF," *Open Sources*, p. 51. Bradner is a longtime participant of the IETF who served as an area director and IESG member between 1993 and 2003.

- 75. M.A. Padlipsky, *Elements of Networking Style*, p. 104.
- 76. Lyman Chapin, quoted in G. Malkin, "Who's Who in the Internet: Biographies of IAB, IESG, and IRSG Members," IETF RFC 1336, May 1992; http://www.ietf.org/rfc/rfc1336.txt.
- Einar Stefferud, quoted in M.T. Rose, "Comments on 'Opinion: OSI Is (Still) a Good Idea," ConneXions, vol. 6, no. 8, 1992, pp. 20-21.
- 78. See M. Libicki, Information Technology Standards, pp. 108-119; D.C. Wood, "Federal Networking Standards: Policy Issues," StandardView, vol. 2, no. 2, 1994, pp. 218-223; J.S. Quarterman, "The Demise of GOSIP," Matrix News, vol. 4, no. 10, 1994, p. 6; and R. Hunt, "The Future of TCP/IP and OSI/GOSIP—Migration or Coexistence?" Proc. Networks 93 Conference: Held in Birmingham, June-July 1993, Blenheim Online, 1993, pp. 423-437.
- T. Berners-Lee, Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by its Inventor, HarperSanFrancisco, 1999, p. 16.
- C. Huitema, *IPv6: The New Internet Protocol*, p. 1. Huitema was on the IAB from 1991 to 1996 and was Internet Architect from March 1993 to July 1995.
- 81. For a lengthy and enthusiastic appraisal of the IETF as "an international phenomenon that conforms well to the requirements of [Jürgen] Habermas's discourse ethics," see A.M. Froomkin, "Habermas@Discourse.Net: Toward a Critical Theory of Cyberspace," *Harvard Law Rev.*, vol. 16, no. 3, 2003, pp. 749-873. For a more critical appraisal of the problems of informal standardization bodies (including the IETF), see P.J. Weiser, "The Internet, Innovation, and Intellectual Property Policy," *Columbia Law Rev.*, vol. 103, no. 3, 2003, pp. 573-576.
- 82. "Rough consensus and running code" provided a model for standards development for the World Wide Web. See T. Berners-Lee, *Weaving the Web*, p. 98. See also "Global Grid Forum Overview: Structure and Process"; http://www.gridforum. org/L_About/Struc_Proc.htm; and "Voting System Performance Rating Charter"; http://vspr. org/charter ("The underlying approach and structure of VSPR is modeled after the IETF").
- 83. As one participant stated recently, the "processprocess" is "sucking resources out of IETF's technical/standards work." S. Dawkins, "Two New Proposals/Drafts—Details," 10 Feb. 2004; available from http://eikenes.alvestrand.no/pipermail/

solutions/2004-February/000969.html. For a summary of the IETF's recent problems, see E. Davies, ed., "IETF Problem Statement," IETF RFC 3774, May 2004; http://www.ietf.org/rfc/ rfc3774.txt.



Andrew L. Russell is a PhD candidate in the Department of the History of Science and Technology at Johns Hopkins University. He holds a BA in history from Vassar College and an MA in history from the University of Colorado at Boulder. He also

worked at the Harvard Information Infrastructure Project (1997–1999).

Readers may contact Andrew Russell about this article at Dept. History of Science and Technology, The Johns Hopkins University, 3505 N. Charles St., Baltimore, MD 21218; arussell@jhu.edu.

For further information on this or any other computing topic, please visit our Digital Library at http://www.computer.org/publications/dlib.

