1

Introduction: The Problem of the Root

For two days in July 1998, one hundred and fifty people gathered in a windowless hotel convention room in Reston, Virginia. The crowd comprised techies in T-shirts, trademark lawyers in suits, academic and business people, and a small but significant number of Europeans, Latin Americans, and Asians. The meeting had an ambitious goal: to "prepare a model, a set of common principles, a structure and general charter provisions" for the formation of a global governance body for an Internet naming and addressing authority.[1] The meeting was compared to an Internet "constitutional convention" by some. But the delegates to this convention were not diplomats or legislators, and its participants held no formal credentials. There had been some attempts to encourage preregistration, but for all practical purposes attendance was completely open—anyone who walked in could participate. A call had been issued by a self-appointed, hastily assembled, and loosely defined steering committee, whose membership remained fluid and controversial for weeks afterwards. Aside from a few basic agenda and scheduling decisions, the process was made up on the spot. There were no formal committee chairs; facilitators either volunteered or were appointed. There were not even arrangements for breakout rooms for subgroups to work in, so the committees had to huddle in corners of the same noisy room and sometimes shout to make themselves heard.

The Reston meeting was the first in what turned out to be a series of four such conferences known as the International Forum on the White Paper (IFWP). Reston, Virginia, was an appropriate location for the inaugural meeting; it was ground zero of the commercial Internet explosion of the

mid-1990s. The region was home to Network Solutions, Inc. (NSI), the government contractor that had turned domain name registration into a multimillion dollar business and that was the site of the critical A root server, the central source of data for coordinating the world's Internet names. Reston itself was the headquarters of the Internet Society. The Pentagon and the National Science Foundation, whose sponsorship of the Internet had pushed it to the brink of global critical mass, were only a few miles away. So was the Corporation for National Research Initiatives (CNRI), which hosted the secretariat of the Internet Engineering Task Force (IETF) and once served as the organizational home of Robert Kahn and Vint Cerf, the joint inventors of the Internet protocol. Commercial firms that had risen to prominence with the Internet, such as MCI, PSINet and America Online, located their headquarters nearby.

For several years it had been clear that the Internet was no longer a subsidized tool of education and research but a vibrant new global medium. The Internet was growing at exponential rates, and its importance to the economy was becoming increasingly evident. But key technical functions such as name and address management were still performed under contracts with the U.S. military and the National Science Foundation. Foreign governments were becoming increasingly restive about unilateral U.S. control of such an important part of the global communication infrastructure. Network Solutions' unplanned-for and increasingly lucrative monopoly over domain name registration was also a point of growing contention.

The transition process, everyone knew, would be risky and controversial. Domain names and IP (Internet Protocol) addresses stood at the core of the Internet's operation. If they were handled poorly, the Internet could break. As the stakes grew higher, however, the Internet community had fallen into rancorous battles over policy and control. The years of escalating tension became known as the domain name wars. Finally, in July 1997, the U.S. Department of Commerce initiated a formal proceeding to privatize the domain name system (NTIA 1997). The result was a policy document officially titled "Management of Internet Names and Addresses" but universally known in Internet circles as simply "the White Paper" (NTIA 1998b).

With the release of the White Paper on June 3, 1998, the U.S. government took an unusual approach to the transition. Instead of using its rule-

making powers to settle issues, instead of creating an organization and specifying the rules it would follow, it threw the responsibility back to the warring parties, back to what it called private sector stakeholders. The government's announced intention was to "recognize . . . and seek international support for a new, not-for-profit corporation formed by private sector Internet stakeholders" (NTIA 1998b, 31749). That new corporation, not the U.S. government, would make the difficult policy decisions. It was up to the Internet community itself to form this organization and come to the U.S. government with a single proposal that commanded the unified support of the global Internet community. This had to be done in only four months.

1.1 A Constitutional Moment

Hence, the unusual gathering in Virginia. The IFWP was the response of those who took literally the U.S. government's call for private sector leadership. It was conceived as an open, neutral arena that would bring the key parties involved in the domain name wars together in face-to-face meetings. Tamar Frankel, a Boston University law professor who was expert in corporate governance structures but largely innocent of the Internet and its controversies, agreed to preside over the meetings. Many participants in the Reston meeting reveled in the government's willingness to keep its hands off and allow the "Internet community" to resolve the problems on its own. The words *consensus* and *self-governance* were on everyone's lips. Ira Magaziner, the Clinton administration policy adviser who had supervised the White Paper proceeding, gave the Reston gathering a kind of official blessing with an opening speech and then left to allow "the private sector" to do its work. Jon Postel, the respected Internet technologist who had managed the number space and domain name delegations for many years, sent a letter from California expressing his hopes that the forum would succeed. The Reston meeting was followed by quickly organized counterparts in Geneva, Singapore, and Buenos Aires. The ultimate result, for better or worse, was the Internet Corporation for Assigned Names and Numbers (ICANN).

The IFWP seemed to initiate a unique form of international organization. Normally, policy for global resources such as Internet names and

numbers would be coordinated through established institutions, such as national governments, trade associations, standards bodies, international treaties, or formal international organizations. The Internet was different, however. It seemed to call forth an entirely new spirit for collective action. It had created a perplexing set of issues that eluded resolution by any one government or organization. There was no suitable legal or organizational framework in place. Various organizations—the Internet Society, the International Telecommunication Union (ITU), alternative "root server confederations"—had tried and failed to create one.

The type of problem that the White Paper set out to solve was not entirely unprecedented. The telegraph and postal systems, radio, satellites, air travel, and maritime transport all had raised similar issues in the past. These problems had been handled by collective action among nation-states through formal treaties or intergovernmental organizations such as the ITU. Something different was happening here. The intellectual, commercial, and political climate surrounding the Internet militated against the involvement of states and state-derived international organizations. True, the U.S. government had set the stage for the process by holding a formal proceeding and issuing a policy statement. It still held substantial power over who would be selected to administer the authority. But the method it was using deviated sharply from traditional ones. Indeed, at the initial IFWP meeting, Magaziner presented the White Paper as an epochal change in the nature of international organization. Drawing on a distinction between "industrial society" and "information society" that was popular at the time, Magaziner suggested that the White Paper's methods were more appropriate to the information age. "We believe that the Internet as it develops needs to have a different type of coordination structure than has been typical for international institutions in the industrial age. [G]overnmental processes and intergovernmental processes by definition work too slowly and somewhat too bureaucratically for the pace and flexibility of this new information age."[2] The Harvard professor Lawrence Lessig, on the other hand, a critic of the administration's private sector approach, complained that "we are creating the most significant jurisdiction since the Louisiana purchase, and we are building it outside the review of the Constitution."[3]

A scene from the International Forum on the White Paper is thus a fitting way to open this book. Although it was only one of many episodes in

the process, it was perhaps the purest exemplar of what David Post (1998) has called "cyberspace's constitutional moment." The Internet's growth created a need for a new kind of social contract. Its crucial central coordinating functions needed governing arrangements that were both technically robust and capable of winning the support and cooperation of global, diverse, constantly expanding, and often conflicting groups of interested parties. The Internet's structure was so distributed, and the organizations that built it were so diverse and so informal, however, that no single group, not even the U.S. government, possessed the legitimacy and authority to pull it all together on its own. If the IFWP process seemed ramshackle and ad hoc, it was because it had the task of bootstrapping authority on a global scale in an absurdly compressed time span. There was, for precisely this reason, something exhilarating about the IFWP's brief moment. Like the first meetings of the Long Parliament in the English revolution of 1640,[4] the apparent power vacuum produced a heady feeling of self-determination. It encouraged idealistic pronouncements based on first principles. It fostered the illusion that the needed governance arrangements could be designed from scratch. And the IFWP, like the Long Parliament, was ultimately bypassed and superseded by more powerful forces impatient with the transaction costs of an open, democratic process. Yet, by creating expectations of open public participation and private sector consensus the IFWP had a lasting impact on the process.

1.2 The Root

What problem precipitated this constitutional moment? What great issue animated these global negotiations? The object of the controversy was control of a seemingly obscure set of technical functions related to naming and addressing computers on the Internet. Data communication on the Internet takes place by breaking messages into smaller units called packets and routing them from network to network. In order to know where to go, each packet must carry a numerical address, known as an Internet Protocol (IP) address. Every computer connected to the Internet must have a unique IP address. To supplement these numerical addresses, the computers, routers, and other resources connected to the network can be given user-friendly names like *www.yahoo.com,* known as domain names.

Many vital activities on the Internet, such as email or the World Wide Web, use domain names rather than IP numbers as addresses. But for packets to flow across the network, the user-friendly names must be translated into IP addresses. Both kinds of addresses—domain names and IP numbers—are valuable resources, a kind of virtual real estate that can be bought and sold.

It was name and address management that created the controversies that led to the IFWP. The specific set of functions at issue can be summarized as

- The authority to set policy for and to manage the allocation and assignment of Internet Protocol addresses
- The authority to add new names to the top level of the Internet domain name hierarchy
- The responsibility for operating root servers that distribute authoritative information about the content of the top level of the domain name space

These functions are defined more precisely and discussed in greater detail in chapters 2 and 3. Although they may sound uninteresting, they are the technical underpinnings of what the Internet is all about. We tend to speak of the Internet as if it were a *thing,* but in reality the Internet is entirely virtual; it consists of nothing but a software protocol suite known as TCP/IP.[5] The software enables any computer in the world to exchange information with any other computer, regardless of the particular physical networks to which they are attached or the hardware they use. It does this largely by giving computers addresses and names, and providing instructions about how to use them. Consistent and scalable naming and addressing protocols are at the core of TCP/IP's design. The functions enumerated previously are needed to ensure that the names and addresses will be unique. Throughout this book, I refer to that cluster of functions as "the root."[6]

The root is the point of centralization in the Internet's otherwise thoroughly decentralized architecture. The root stands at the top of the hierarchical distribution of responsibility that makes the Internet work. It is the beginning point in a long chain of contracts and cooperation governing how Internet service providers and end users acquire and utilize the addresses and names that make it possible for data packets to find their destinations.

Addresses and names must be globally unique. Ensuring uniqueness in an open, rapidly growing network with millions of users is a coordination problem of some magnitude. The root is the Internet's answer to the problem of coordinating unique identifiers.

The security and stability of the root server system is critical to the viability of any service or function that relies on the Internet. No one disputes the operational significance of the root, and hence no one disputes the need for the formation of permanent, stable organizational arrangements to control—to govern—those functions. But the word *governance* has wider implications.

1.3 Governance

During the debates over the formation of ICANN, an interesting dialogue evolved over the use of the term "Internet governance." To some, "governance" meant the legal and organizational arrangements for management of the root functions. This narrow construction of the term was analogous to the way we use "corporate governance" to refer to the articles and bylaws of an organization, how board members are elected, and so on.

To many others, however, "Internet governance" raised troubling questions. Aside from being a single point of failure, the domain name system (DNS) root is also, potentially, a single point for the surveillance of users and the control of access to cyberspace. The strategic lever of the root, many believed, could be used to enforce public policy and to regulate or control Internet users. "Internet governance" sounded a lot like "a government of the Internet." As David Post (1998) observed,

> If the person or entity controlling the root servers determines that a $1,000 fee (or a certificate of good standing from the California Secretary of State, or a pledge to abide by the laws of Uzbekistan, or a promise not to transmit encrypted messages, or . . .) is required to register a name-number combination and place it in these publicly accessible databases, those who cannot or will not pay the fee, obtain the certificate, or make the required promises, are effectively banished from the global system.

Indeed, the original creators of ICANN always attempted to distance themselves from the term "governance." They preferred to say "technical management." As Esther Dyson put it, ICANN "governs the plumbing,

not the people. It has a very limited mandate to administer certain (largely technical) aspects of the Internet infrastructure in general and the Domain Name System in particular."[7]

The White Paper itself utilized *governance* in both senses, referring at one point to the "bottom-up governance that has characterized the development of the Internet to date" (NTIA 1998b, 31749) and claiming at another that "the U.S. government policy applies only to management of Internet names and addresses and does not set out a system of Internet 'governance'" (31743).

The two meanings define the fundamental dilemma of Internet governance: the intersection of technical management and regulatory control. Where does one end and the other begin?

It is clear that technical management decisions have direct and immediate economic consequences. Decisions made by those who control the root profoundly affect the structure of the rapidly growing market for domain name registration. At the beginning of 2001, that market was valued at about US$1.5 billion, and it had doubled annually for the preceding five years. As discussed in later chapters, it was the conflict over who would be assigned the right to register names under new top-level domains that catalyzed much of the global governance debate. It is a question a root manager cannot avoid making decisions about. The economic value of IP addresses is harder to estimate, because end users are not allowed to trade them in a market. (That, of course, is itself a policy decision of some magnitude that straddles the economic and the technical.) But most experts would view addresses as even more valuable assets than domain names. IP addresses are essential inputs into networked services. Their cost and availability will have a major impact on the business plans of telecommunication service providers and equipment manufacturers in the burgeoning digital economy.

The importance of root governance goes well beyond the dollar value of any real or imagined market for names and addresses, however. As unique identifiers, IP addresses can be used to identify and track users. Similarly, domain name registration records directly reveal to the world the name, email address, and physical address of the registrant. The domain name system establishes a mechanism for the identification and surveillance of the denizens of cyberspace. Consider, then, the security and privacy im-

plications of the policies adopted by an Internet naming and addressing authority. Contradicting privacy concerns are demands by some government agencies to use domain name registration data to facilitate identifying and sanctioning Internet users who break the law. A domain name record can, in fact, function very much like an Internet "driver's license." Here is another policy tug of war that cannot be sidestepped by whoever administers the root.

A similar tension hangs over domain name–trademark conflicts. The domain name system allows almost anyone to think of a name, register it (if it is not already taken), and publish it globally. The brand equity of trademark holders often conflicts with the ability of individuals and small businesses to express ideas and achieve visibility in cyberspace. As discussed in later chapters, major intellectual property holders succeeded in linking domain name registration to the adjudication of trademark–domain name disputes. Indeed, they are trying to leverage the root's ability to monitor and police intellectual property in even more ambitious ways.

The assignment of domain names also intersects with content regulation, or what Americans call free-speech or First Amendment questions. Several interest groups and politicians have called for the creation of a *.xxx* top-level domain in order to clearly identify and segregate sexually explicit material. By the same logic, many businesses and consumers have called for a *.kids* domain that would only contain "child-appropriate" content. But if a domain name authority assigns a *.xxx* domain, is it encouraging governments to use their powers to force all sexual material into that domain? If so, who decides what is X-rated on a global basis? If the root administrator gives someone the *.kids* domain, is it taking responsibility that the sites under that label really are suitable for children? More broadly, should a domain name administrator be concerned with the authenticity of the content associated with a specific domain name?

The tendency for policy demands to be placed on the administration of the root cannot be dismissed. And that does not even begin to touch upon the geopolitical questions. For if one concedes that control of the root is economically, technically, and politically important, then one cannot avoid the issue of how that power is distributed among the world's nations, geographic regions, and cultures. Would Americans feel comfortable if the root of the domain name system were located in China? If not, how

can they expect the Chinese to be happy about its location in the United States?

The uncomfortable fact is that the two meanings of "Internet governance" are inseparably linked. Centralization of control at the root does create levers for the intrusion of politics, policy, and regulation. If these powers are not to be expanded or abused, the governance structure (in the narrow organizational sense) must be designed to prevent this from happening. There is no way to institutionalize control of the root without confronting the larger governance issues. Investigating the nature of these issues forms the central theme of this book.

1.4 Institutionalization

The tools I use are drawn from institutional economics. Institutional economics looks at the interaction of law, economics, and politics; it examines how societies solve collective action problems by defining property rights and establishing governance arrangements. It is interested in technology insofar as it creates new resources that must be incorporated into legal and institutional regimes, or causes changes in transaction costs or relative prices that lead to a breakdown in a preexisting order.

The root—not specific people or organizations—is the protagonist of this story. The development of internetworking endowed the name and address spaces with enormous social value. The Internet's origins in informal, noncommercial, and relatively nonpolitical research and education organizations, however, placed these valuable resources outside the control of existing institutions. The root was essentially unowned, and its inherently global nature made it difficult for nation-states and traditional international organizations to respond. Consequently, as the Net became public and commercial, it fostered an international struggle over the definition of property rights and governance arrangements. The governance problem could only be solved through the development of new institutional arrangements. This is therefore a case study in institutional innovation, all the more interesting and complex because it happened on an international scale.

Admittedly, *institutionalization* is an ugly and seemingly unexciting word. How much more interesting to talk about the vast amounts of

money that can be made from e-commerce or the exciting new capabilities of information technology. But *institutionalization* is the only word that gets to the essence of what happened (and continues to happen) to the Internet from 1996 to 2001.

When we ask who controls the Internet, the response typically takes one of two extremes. The first, favored by many technologists, is to say that no one controls it. The Internet is inherently uncontrollable. Technology is more powerful than governments, traditions, cultures; the Internet "routes around" censorship, and so on. The other extreme is to search for the names of a clique of people or corporations who are said to have overwhelming power to issue authoritative commands. The Internet is run by MCI, or AOL, or the U.S. government. Both responses, I think, miss the point. For any complex sociotechnical system, especially one that touches as many people as the Internet, control takes the form of *institutions,* not commands. Contending parties work out rules and procedures that make their interactions less costly, more stable and predictable. They supplement these rules with organizations that monitor compliance and sanction those who break the rules. In such a process, control is never perfect and no one gets exactly what he wants. But it is false and misleading to say that there is no control, no social constraint. Some parties have more bargaining power than others. Rules are never perfectly fair or neutral; they are always formulated and implemented in ways that favor some types of interests over others. Not everyone has the same amount of resources to devote to monitoring and enforcing their rights. Some people break the rules and get away with it. In short, there are winners and losers in any institutionalization process. And there is always continuing pressure for the modification of the rules in ways that reflect the special interests of various parties. The value of the institutional perspective is precisely that it provides a framework for understanding these kinds of interactions.

1.5 Goals and Plan of the Book

I have three related objectives in writing this book. One is to tell the story of Internet governance objectively and comprehensively, and in the process apply what we know about property rights economics and institutional analysis to the story. Another is to synthesize a technological understand-

ing of DNS and IP addressing with the economic and institutional analysis. This is necessary if we are to understand how technical systems are shaped by political and institutional constraints, and vice versa, and how the development of technical systems can be frozen or diverted into unexpected paths by legal and political pressures. Finally, I want to assess what is really at stake in this matter, to discuss and evaluate contrasting claims about the significance of ICANN and its new regime.

The book is organized into three parts. Part I is framework and background: it analyzes name and number spaces in technical and economic terms, and then elaborates the theories of property rights and institutional change that can be applied to the issue. This part draws on the work of Gary Libecap (1989) on the initial formation of property rights, Elinor Ostrom (1990; 1994) on collective action to resolve common pool problems, and John Richards (1999) on international regimes.

Part II is historical. It traces the growth of the root, the development of property rights conflicts, and the emergence of a new institutional framework to resolve those conflicts. It shows how organized interest groups, particularly intellectual property holders, deliberately reached for control of the root, the centralized point of coordination and control, to impose an order upon the Internet more to their liking. They were joined by an entrenched technical hierarchy that wanted to solidify its role in the management of the Internet and lacked the vision to understand what they were giving up to get it.

Part III explores the stakes and the longer-term policy and social issues posed by the institutionalization of the Internet under ICANN. It characterizes ICANN as a new international regime, one that is likely to become more politicized and to attract more direct and formal participation by governments as it matures. The new regime is analogous to radio broadcasting regulation, in that it uses its exclusive control of a resource to regulate the economic structure of an industry and to sanction various forms of user behavior. Unlike broadcast regulation, however, this is an explicitly global regime and has been placed outside the normal institutional constraints of national governments. The book also explores the World Intellectual Property Organization's attempt to use the ICANN regime to create a new system of global property rights in names.