TE WHARE WĀNANGA O TE ŪPOKO O TE IKA A MĀUI
VICTORIA
UNIVERSITY OF WELLINGTON

6 November 2007

Official Submission Before the Parliamentary Internet Caucus

Legal Foundations of the Internet; Technical Management and Coordination, and the Standards and Protocol Setting Processes; Identity and Description of Key Entities in Internet Governance

Philip J. Greene
InternetNZ Senior Research Fellow in Cyberlaw
Victoria University School of Law
November, 2007

This paper will discuss, generally, the technical management of the Internet, and identify and discuss the entities that play a role therein. Part I of this paper will address the history of the Internet which, by recounting how this particular network of networks was established, lay the groundwork for understanding how the current state of play in Internet Governance came to be. Included in this discussion will be an introduction to key concepts and processes, such as the Domain Name System, the Root Server System, and the Request for Comments process. Part II will follow in this explanation by identifying in greater detail the key entities that play a role in the technical management of the Internet, the development and implementation of its underlying protocols and standards, and the policy making and operational aspects of its overall structure.

Introduction - Technical Management of the Internet - Major Entities and Their Roles

"Internet Governance" is a term used to describe the manner in which the policies, protocols, parameters and standards of the Internet are defined and implemented. While it has a considerable amount of political significance, with respect to the management of the Domain Name System (DNS) root system, activities currently managed by the Internet Corporation for Assigned Names and Numbers (under a Cooperative Agreement with the U.S. Department of Commerce, as explained below), this paper will proceed under the more general notion of "Internet governance" as being related to the manner in which the major entities work together to ensure the Internet's continued operation.

FACULTY OF LAW Te Kauhanganui Tatai Ture
PO Box 600, Wellington, New Zealand
Phone +64-4-463 6366 Fax +64-4-463 6365
Website www.lawschool.vuw.ac.nz

The technical management, or "governance," of the Internet is conducted by a number of entities, working in concert. In describing their respective roles, and the manner in which they were established, it is helpful to provide a brief history of the origins of the Internet.

Indeed, to a greater extent, many of the entities currently responsible for Internet governance reflect a legacy of the entities that played a role in the Internet's birth. By reciting a brief history of the Internet, below, the key entities that have played and continue to play a role in Internet governance will be identified. Following this historical section, each key entity and its respective role will be examined in greater detail. The main entities that will be discussed in this paper include:

Internet Society (ISOC)
Internet Architecture Board (IAB)
Internet Engineering Task Force (IETF)
Internet Engineering Steering Group (IESG)
Internet Research Steering Group (IRSG)
Internet Research Task Force (IRTF)
Internet Corporation for Assigned Names and Numbers
Internet Assigned Numbers Authority (IANA)
VeriSign/Network Solutions, Inc. (NSI)
World Wide Web Consortium (W3C)

I. Brief Overview of Internet History, and Introduction to Internet Infrastructure

It has been said that war is the mother of invention, and in the case of the Internet, one might say that the Cold War was its impetus. Indeed, the basic Internet technology known as "packet-switching" came about as a result of United States Defense Department efforts designed to ensure continuity of communications in the event of nuclear attack or other emergency.[1] According to Charles Siegal:

> "In the late 1960s, the United States Department of Defense became concerned that an opponent in a nuclear war could easily demolish the lines of communication - the command and control system - that permitted the country's political and military leadership to communicate with its dispersed forces. To counter this possibility, the Air Force's newly-formed Advanced Research Projects Agency ("ARPA") funded the development of a network of computers and software based upon a group of distributed nodes. Rather than transmitting messages only from a central computer to remote computers, the messages were broken up into packets, each containing information about its source and destination, and sent into a web of nodes, to be reassembled at its destination. If one node was disrupted, the packets would be routed around it, maintaining the robustness of the system as a whole."[2]

This network became known as the ARPANET[3]. During the 1970s and early 1980s, ARPA (which was/is also known by the longer name DARPA, which stands for Defence Applied Research Projects Agency) experimented with various internetworking

protocols, led by researchers including Vinton Cerf and Robert Kahn. As more computers were added to the network, it became necessary for ARPA to develop, and then require the use of, a standard protocol. Siegal continues:

> "Interconnectivity is made possible by various programs (protocols) that permit the transmission of information - text, graphics, sound and video - between networks. The most important of these protocols are (1) the Transmission Control Protocol ("TCP"), which disassembles messages into packets at the source and reassembles them at the receiving end; and (2) the Internet Protocol ("IP"), which addresses the packets so they can be transmitted over various systems to their ultimate destination."[4]

Thus was installed the ubiquitous TCP/IP; on 1 January 1983, TCP/IP was officially recognized as ARPANET's host protocol.[5] As more and more computers, vendors and users were added to the Internet, communication, education and outreach were seen as vital to maintaining compatibility between the various computers, operating systems, and connected networks. As Siegal explains:

> " The Department of Defense required TCP/IP compatibility in many of its orders, but there was a lack of communication between users, inventors and vendors regarding how TCP/IP was supposed to work, and how it was to be used. A newly created DARPA Internet organization, the Internet Activities Board (IAB), held a conference to educate vendors through tutorials, design meetings and workshops. Vendors also began attending the meetings of the new DARPA working group, the Internet Engineering Task Force (IETF). In 1988, DARPA organizations organized the first Interop trade show. Vendors were very successful in the display of their products because they had worked to make their products interoperable with all the other products."[6]

As ARPANET grew in size and visibility, other researchers wanted to gain access to the Internet, and the focus began to shift from defense-based to research-based. During the 1980s, the National Science Foundation (an agency of the U.S. Government) established its own academic-centered network, NSFNet, in 1986, and also adopted TCP/IP as its network protocol.[7] By 1990, ARPANET was taken out of service, replaced by the NSFNet, which became the new "backbone" of the Internet.[8]

As NSFNet assumed responsibility for the fledgling "Internet," it also inherited oversight of and support for DARPA's Internet organizations, namely, the Internet Activities Board (IAB) and the Internet Engineering Task Force.[9] "Initially just one of the many task forces, the Internet Engineering Task Force (IETF) began to grow tremendously starting in 1985. Consequently, it gained prominence among the groups, and a substructure of working groups was created under it. The remaining DARPA working groups were combine into the Internet Research Task Force."[10]

Separately, in 1990 -1993, the World Wide Web was created and implemented by Tim Berners-Lee, a researcher at CERN, a particle physics laboratory in Geneva, Switzerland. Initially, the World Wide Web was purely text-, not graphics-based.

However, the creation of the Mosaic web browser at the University of Illinois in 1993 brought the capability of viewing images and multimedia presentations.[11] World Wide Web standards and protocols are developed under the auspices of the World Wide Web Consortium, as discussed below.[12]

Birth of Domain Name System

During the 1970s, as the Internet continued to grow, it soon became apparent that a standardized system needed to be implemented in order to determine the identity of each computer, and the means by which one computer could connect to another. "Early in the history of the Internet, each host had both a name and address. Every host had its own copy of the table or database that correlated names and addresses. ... This early incarnation of the name system lacked domains, and as the Internet grew, the lack of coordination gave rise to problems. Craig Partridge is quoted as saying:

> When we started to get about two thousand hosts, that's when things really began to come apart. Instead of having one big mainframe with twenty thousand people on it, suddenly we were getting inundated with individual machines, and everyone wanted to be named Frodo."[13]

In short, a system of numbering and addressing had to be developed.[14] Under the TCP/IP system, each connected computer was assigned an Internet protocol (IP) address, a unique, 32-bit number. In 1977, a young computer scientist by the name of Jon Postel was given responsibility for the nascent "Domain Name System." The DNS had its humble origins as a single computer file, "hosts.txt," which was maintained by the Stanford Research Institute under a Defense Department contract.[15] Postel maintained the DNS at Stanford until 1988, at which time he, and the contract, moved to the University of Southern California's Information Sciences Institute (ISI).[16] In December 1988, this function, and the entity providing it, became known as the Internet Assigned Numbers Authority (IANA).[17] "Dr. Jon Postel of the University of Southern California's Information Sciences Institute (ISI) assumed the task of assigning blocks of IP addresses to computer networks. Because no two computers had the same IP address, it was possible to locate any computer on the Internet simply by knowing its IP address. TCP/IP made possible a system of routing that permitted a user to dispatch a message onto the Internet, knowing only the IP address of the computer he wished to reach, with confidence that the message would eventually reach its intended destination."[18]

Postel, with fellow colleague Paul Mockapetris, thus developed the "Domain Name System," or DNS, a system by which each IP address is "mapped" to a specific domain name. But Postel and Mockapetris went beyond simply assigning and recording names to IP numbers, they also established a hierarchical system, or "domain name space," for keeping track of them.[19]

> "That name space is divided into top-level domains, or TLDs. Each top-level domain is divided into second-level domains, and so on. Under the plan developed by Postel and Mockapetris, there were seven generic, three-letter top-

level domains: .com, .net, .org, .edu, .gov, .mil, and .int. In addition, there were two-letter country code top-level domains such as .jp, .us, and .fr. At the outset, it was thought that .com would be used by commercial entities, .net by entities involved with the Internet networking infrastructure, .org by nonprofit organizations, and .edu by educational institutions. Today, the restrictions on the first three of these have long since fallen away."[20]

As such, a typical domain name will be composed of a top-level domain, a second level domain, and perhaps a third-level, fourth-level and further sub-domains. They are located in either generic Top-Level Domains (.com, .net, .org, .edu., .mil, .gov, .biz, .info, and others, known as gTLDs) or in country-code domains, corresponding to the particular country or nation to which the domain is assigned (known as ccTLDs). For example, in the domain name Amazon.com, ".com" is the generic Top-Level Domain, while "Amazon" is the second-level domain. Similarly, in the domain name Internetnz.net.nz, "Internetnz" is the third-level domain, ".net" is the second-level, and .nz is the country code Top-Level Domain (ccTLD) for New Zealand. A complete list of gTLDs and ccTLDs, respectively, may be found at the IANA web site, located at:

http://www.iana.org/gtld/gtld.htm, and

http://www.iana.org/root-whois/index.html

Actually, ccTLDs do not necessarily represent countries or nations, per se. "To avoid political problems, Postel used the ISO 3166-1 country codes to define what entity would warrant a ccTLD. Because these codes were provided by the International Organization for Standardization, an international association of national standard-setting bodies, their objectivity successfully shielded IANA from the political pressure of deciding what was and was not a country."[21]

Under Postel's leadership, IANA's Network Information Center (NIC) oversaw both the implementation of the DNS,[22] including the establishment of gTLDs, and the delegation of ccTLDs.[23] "The first ccTLD, ".us," was created and delegated in March 1985. Two other delegations, ".uk" (United Kingdom) and ".il" (Israel), followed in the same year."[24]

Postel used a standards-making process known as the "Request for Comments" or RFC process, explained below, in proposing and implementing both the gTLD and ccTLD regimes. RFC 920, issued in October 1984, incorporated the above-referenced ISO-3166 list as ccTLDs.[25] RFC 1591, issued in March 1994, sets forth policy and procedure with respect to Domain Name System Structure and Delegation. This document "provides some information on the structure of the names in the Domain Name System (DNS), specifically the top-level domain names; and on the administration of domains."[26]

Root Server System

In addition to having a hierarchical system of generic and country code top-level

domains, the DNS system operates astride what is known as the Root Server System.[27] The Root Server System directs traffic amongst connected computers by use of each respective domain name and IP number, using information stored in the System.

When a domain name or email address is entered into a Web browser, your computer translates the domain name back into the IP address, which directs your computer as to the location of the desired site. There is no one, single database of all IP addresses on the Internet; it would be too vast. Rather, the domain name database is:

> "held in a distributed and hierarchical fashion among numerous computers across the Internet. At the top of the hierarchy is a computer that holds the root directory, which contains a list of authoritative computers for all of the top-level domains, both geographic (e.g., .uk, .de, .us, .jp) and generic (.com, .gov, .edu). These computers, in turn, contain lists of the computers that are authoritative for the second-level domains (e.g., mcdonalds.com, harvard.edu) within each top-level domain, and so on. The value of this distributed approach is that it allows address changes to be made at the most local level, and therefore obviates the need for any huge, centralized administrative body to keep track of such changes."[28]

The Root Server System is made up of a "root file," "zone files," containing the various domain names, and "name servers." "The Internet relies on an underlying centralized hierarchy built into the Domain Name System (DNS) to control the routing for the vast majority of Internet traffic. At its heart is a single data file, known as the 'root.'"[29] Actually, there are 13 "root servers," each identified by a letter of the alphabet, from A to M. The "A Root Server" is considered to be "authoritative," and is maintained by VeriSign, Inc., under a contract with the U.S. Department of Commerce. All other root servers accept and follow the A Root Server, and are continuously updated to mirror it.[30]

> "The DNS assigns domain names to IP addresses and informs querying computers which IP address belongs to a particular domain name - a process called name "resolving." Domain names are unique and have a hierarchical structure like a tree. The last part of each domain name, e.g. ".com", indicates the Top-Level-Domain (TLD) the domain name belongs to. There are fourteen generic TLDs like ".com" and around 244 country-code TLDs like ".uk." The second-to-last part of a domain name, e.g. berkeley in www.berkeley.edu, is called a second level domain and indicates the subdomain berkeley in the TLD .edu. The two decisive components of the DNS are the root file and name servers. The root file, also called "the dot," contains the authoritative information about TLDs. The root file lists all TLDs and, for each TLD, the name server, which stores a list of the (second-level) domain names of the TLD along with their corresponding IP addresses. By referring to a particular name server, the root file gives that name server control over the domain names in that TLD.
>
> The root file determines that all (second-level) domain names ending with

> .museum are stored with their IP addresses on the name server NIC.ICOM.ORG, which has the IP address 195.7.65.253. All domain names depend on the root file because it refers searching computers to particular name servers and, thus, gives these name servers the authority to assign domain names under the TLDs. The original root file is stored on the "A-root-server" of the company NSI-Verisign in Virginia. As a practical matter this root file is authoritative because other computers accept and follow it."[31]

Toward Privatization

As noted above, as DARPA relinquished oversight, circa 1990, the National Science Foundation (NSF), via the network NSFNet, was there to provide the "backbone," leadership and funding for the Internet. At this time, IANA continued to perform a number of services related to the DNS, namely, performing domain name registration services (registrar function), keeping track of the domains already registered (registry function), and other aspects of coordinating IP address allocation, and overseeing the root server system, from both a policy and operational standpoint. These functions came to be known as "the IANA function."[32] On December 31, 1992, NSF entered into a cooperative agreement with a Herndon, Virginia company by the name of Network Solutions, Inc. (NSI), to perform the operational aspects of the IANA function, while reserving to Postel and IANA policy oversight.[33]

> "This meant that Postel (through his employer, USC) and Network Solutions would share authority over the naming system in various ways. Network Solutions became the sole registrar for the main nonmilitary domains (dot-com, dot-net, dot-org, and dot-edu. But Postel retained "policy" authority: the power to decide, for example, the number and content of the top-level domains. In some ways, this new arrangement fortified Postel's position. He still made big-picture decisions, and Network Solutions just executed the details. But in fact the transfer of partial authority to Network Solutions was a crucial turning point in Internet history. For the first time, administration of part of the Internet's naming system would be in the hands of a for-profit company. And also for the first time, a private firm with interests at odds with the founding vision would be exercising real authority over their creation. Perhaps even more significantly, Network Solutions became the custodian of the physical root server. As a result, the text files that make up the root were transferred to the Network Solutions offices in Herndon, Virginia, one hour west of Washington, D.C."[34]

In truth, NSI was but one of three commercial entities involved in conducting the "Internet Network Information Center" element of the IANA function. NSF entered into cooperative agreements with Atlantic Telephone and Telegraph (AT&T), to handle the directory and database services component, and with General Atomics, to handle the information services component. This collaborative suite of services became known as "the InterNIC."[35]

In 1997, U.S. Government oversight of the Internet changed hands once again, and "the Commerce Department was chosen to replace the National Science Foundation

as the lead agency," specifically, within the National Telecommunications and Information Administration (NTIA).[36] "On July 1, 1997, a Presidential Executive Order known as the "Presidential Directive on Electronic Commerce" authorized the Secretary of Commerce to "support efforts to make the governance of the domain name system private and competitive and to create a contractually based self-regulatory regime that deals with potential conflicts between domain name usage and trademark laws on a global basis."[37]

The next day, NTIA then sought comments on DNS policy issues, issuing a "Request for Comments" of its own. Issued 2 July 1997, "Request for Comments in the Matter of Registration and Administration of Internet Domain Names" sought "comments on the current and future system(s) for the registration of Internet domain names.[38] The RFC offered the following as background:

> "The rapid growth in the use of the Internet has led to increasing public concern about the current Internet domain name registration systems. .... The enormous growth and commercialization of the Internet has raised numerous questions about current domain name registration systems. In addition, the present system will likely undergo modification when the National Science Foundation's cooperative agreement (NSF agreement) with Network Solutions Inc. to register and administer second-level domains for three top-level domains expires in 1998....
>
> The United States Government played a central role in the initial development, deployment, and operation of domain name registration systems, and through the NSF agreement as well as Defense Advanced Research Projects Agency (DARPA) agreement(s) continues to play a role. In recent years, however, Internet expansion has been driven primarily by the private sector. The Internet has operated by consensus rather than by government regulation. Many believe that the Internet's decentralized structure accounts at least in part for its rapid growth.
>
> The Government has supported the privatization and commercialization of the Internet through actions such as the transition from the NSFNET backbone to commercial backbones. The Government supports continued private sector leadership for the Internet and believes that the transition to private sector control should continue. The stability of the Internet depends on a fully interconnected and interoperable domain name system that must be preserved during any transition.
>
> Various private sector groups have proposed systems for allocating and managing generic top level domains (gTLDs). The Government is studying the proposals and the underlying issues to determine what role, if any, it should play. The Government has not endorsed any plan at this time but believes that it is very important to reach consensus on these policy issues as soon as possible. The United States Government seeks the views of the public regarding these proposals and broader policy issues as well."[39]

As noted above, the U.S. Government was keenly aware that the Cooperative Agreement with NSI was set to expire on September 30, 1998.[40] NTIA released two documents in 1998 to continue the policy-making process with respect to the DNS debate. The first, released in January 1998, was titled "A Proposal to Improve Technical Management of Internet Names and Addresses," and went by the nickname, "The Green Paper."[41] "The Green Paper proposed the creation of a new not-for-profit corporation, 'operating as a private entity for the benefit of the Internet as a whole,' to manage domain names, IP addresses, and the root server network. Current IANA staff would move to the new organization. The U.S. government would transfer existing IANA functions, the root system, and the appropriate databases to the new corporation; the government would 'participate in policy oversight to assure stability' for no more than two years. The Green Paper emphasized that the new organization and its board 'must derive legitimacy from the participation of key stakeholders.'"[42]

The Green Paper, released in late January, 1998, received a considerable amount of public comment.[43] Over the course of the next several months, the U.S. Government, led by the administration of President William Clinton (the "Clinton Administration") made several revisions to the proposals set forth in the Green Paper, and on June 3, 1998, NTIA issued a document titled "Management of Internet Names and Addresses." This document became known as "The White Paper."[44]

The White Paper echoed the Green Paper's call for a private entity to take over policy and operational oversight for the DNS, and referred to this entity as "NewCo." Noting that "overall policy guidance and control of the TLDs and the Internet root server system should be vested in a single organization that is representative of Internet users," the White Paper proposed:

> "the creation of a private, not-for-profit corporation (the new corporation) to manage the coordinated functions in a stable and open institutional framework. The new corporation should operate as a private entity for the benefit of the Internet as a whole. The new corporation would have the following authority:
>
> 1. to set policy for and direct the allocation of number blocks to regional number registries for the assignment of Internet addresses;
>
> 2. to oversee the operation of an authoritative root server system;
>
> 3. to oversee policy for determining, based on objective criteria clearly established in the new organization's charter, the circumstances under which new top-level domains are added to the root system; and
>
> 4. to coordinate the development of other technical protocol parameters as needed to maintain universal connectivity on the Internet."[45]

The White Paper did not envision a Government role, per se, in the creation of this NewCo, rather, it contemplated the NewCo to be created by outside entities, and

then which the U.S. Government would then "recognize."[46]

> "... the U.S. Government is prepared to recognize, by entering into agreement with, and to seek international support for, a new, not-for-profit corporation formed by private sector Internet stakeholders to administer policy for the Internet name and address system."

....

> "We anticipate that the new corporation's organizers will include representatives of regional Internet number registries, Internet engineers and computer scientists, domain name registries, domain name registrars, commercial and noncommercial users, Internet service providers, international trademark holders and Internet experts highly respected throughout the international Internet community. These incorporators should include substantial representation from around the world."[47]

In October 1998, a proposal was submitted to the Commerce Department, that a newly-incorporated California non-profit known as the Internet Corporation for Assigned Names and Numbers (ICANN), would be the NewCo envisioned in the White Paper. Jon Postel, representing IANA, was a key figure in this proposal.[48] On November 25, 1998, the Commerce Department entered into a Memorandum of Understanding with ICANN, and within a few months ICANN agreed with ISI to take over the IANA function. On 26 February 1999, the Commerce Department officially recognized ICANN as the "NewCo" identified in the White Paper,[49] the entity empowered with "overall policy guidance and control of the TLDs and the Internet root server system."[50]

The Green Paper and the White Paper also addressed the "competitive" aspects of the DNS, and saw the need for reform, stating that "[t]he system for registering second-level domain names and the management of the TLD registries should become competitive and market-driven."[51] Indeed, dating from 1993, Network Solutions, Inc. was the sole entity carrying out the IANA function, under agreement with the U.S. Government (NSF and Commerce, respectively).[52] The two "competitive"aspects of the IANA function were the "registry" function and the "registrar" function, in that they were seen as being capable of being performed by more than one entity, and not by the monopoly that NSI was seen as being.[53] The Green Paper defined a "registry" as being "responsible for maintaining a TLD's zone files, which contain the name of each SLD in that TLD and each SLD's corresponding IP number. Under the current structure of the Internet, a given TLD can have no more than one registry."[54] In other words, for each given TLD, such as .com or .org, the registry was the entity responsible for keeping track of all domain names registered in that TLD. A "registrar," on the other hand, is an entity that actually issues domain registrations to the registrants; as defined by the Green Paper, "[a] 'registrar' acts as an interface between domain-name holders and the registry, providing registration and value-added services. It submits to the registry zone file information and other data (including contact information) for each of its customers in a single TLD."[55]

What evolved from this process came to be known as the Shared Registry System.[56] On October 6, 1998, NSI and the Commerce Department amended their Cooperative Agreement to allow for the development of a shared registry system, allowing other registrars to provide domain name registration services in .com, .net and .org. Known as Amendment 11, this document paved the way for the shared registration system that is in place today, with different entities serving as registries for the respective TLDs, and hundreds of entities serving as registrars.[57] For example, NSI (through its parent company, VeriSign, as explained below) remains as the sole registry for .com and .net, while a company by the name of NeuLevel serves as registry for .biz.[58]

The non-competitive aspects of the "IANA function"remained with ICANN, namely, providing coordination, security and stability for the DNS,[59] as is explained in more detail, below.

The Request For Comments (RFC) Process

From its early days, the Internet standards-setting community has relied on a process known as the "Request for Comments" or "RFC Process." During the summer of 1968, representatives from the four ARPA computer science contractors (University of California at Los Angeles (UCLA), Stanford Research Institute (SRI), University of California at Santa Barbara (UCSB), and University of Utah) gathered to discuss the concept of a packet-switched network. This group became known as the Network Working Group, and they began to meet regularly.[60] "The Network Working Group issued its first documentation in April 1969. Steven Crocker, the author of the first standards document, entitled it a "Request for Comments" (RFC) because, as Vint Cerf relates, "we were just graduate students at the time and so had no authority. So we had to find a way to document what we were doing without acting like we were imposing anything on anyone."[61] Crocker later recalled:

> "I remember having great fear that we would offend whomever the official protocol designers were, and I spent a sleepless night composing humble words for our notes. The basic ground rules were that anyone could say anything and that nothing was official. And to emphasize the point, I labeled the notes "Request for Comments."[62]

The name stuck, as other Network Working Group members also labeled their standards documents "RFCs." An RFC could range from a working document to technical reports to meeting notes. "If a protocol seemed interesting, someone implemented it and if the implementation was useful, it was copied to similar systems on the net."[63]

To this day, the RFC Process is used in standard and protocol setting. The process is supervised by the RFC Editor, working in conjunction with the Internet Engineering Task Force's leadership, as explained below.[64] The RFC Process "now constitutes the official publication channel for Internet standards and other statements by the Internet research and engineering community (including the occasional poem and

humourous spoof. .... Officially adopted RFCs are numbered and available free of charge to anyone via the Internet. A number of sites on the Internet contain complete collections of the RFC series, including http://www.rfc-editor.org and http://community.roxen.com/developers/idocs/rfc/."[65]

The Internet Standards Process, Rough Overview

In the pages that follow, the various entities and their respective roles will be described in greater detail. This thumbnail sketch, however, offers some illumination as to the process, as well as its open, pragmatic, and consensus-based nature:

> "The Internet's technical standards are set through a process that bears a surprising resemblance to the low-tech ideal of the New England town meeting. The body with de facto responsibility for setting Internet standards is the Internet Engineering Task Force ("IETF"), which is composed primarily of computer scientists, programmers, and engineers. The IETF holds open meetings three times per year, at which various parties are permitted to comment on numerous issues surrounding Internet technical standards. After these meetings, working groups within the IETF consider the various comments and come to consensus around a particular technical standard. These standards are then approved by the Internet Architectural Board ("IAB") and then published on the Internet in the form of a Request for Comment ("RFC").
>
> The standards are then adopted and implemented voluntarily by the wider Internet community. The central and salient fact about the Internet coordination process is that no central body has the de jure authority to mandate adoption of the standards published in the RFCs. The Internet is a network with distributed intelligence. Because no single computer controls the Internet, the adoption of a given standard cannot be made at a single locus but, instead, must be adopted in a distributed fashion by all of the computers on the Internet. The miraculous part is that this occurs without any formal mandate or legal obligation With a surprising degree of noncentralized coordination, the standards are voluntarily adopted by thousands of system operators all throughout the Internet."[66]

II. Technical Management of the Internet - Major Entities

Inasmuch as the Internet is a "network of networks,"[67] computers and networks connected to one another, it is necessary that all such computers and networks connecting to the Internet share the same network protocols. "Every network functions by having all computers attached to it recognize common network protocols. A network protocol consists of a set of rules that two or more computers must follow to exchange messages. Protocols serve as coordination tools describing both the format of messages that can be sent as well as the way a computer should respond to each message."[68]

> "Almost all Internet technological standards are developed and set by the group consisting of the Internet Society (ISOC) and the units operating under the auspices of ISOC: the Internet Architecture Board (IAB), the Internet

> Engineering Steering Group (IESG), the Internet Engineering Task Force (IETF), the Internet Research Steering Group (IRSG), the Internet Research Task Force (IRTF), and the RFC Editor. It is important to note that, while these units are responsible to ISOC, ISOC allows them a large degree of independence in their technical work. Internet domain names and IP addresses are the province of the Internet Corporation for Assigned Names and Numbers (ICANN) and its Internet Assigned Numbers Authority (IANA). World Wide Web standards are developed by the World Wide Web Consortium (W3C)."[69]

While it is impossible to list all entities which play a role in the overall operation of the Internet, the following organizations are generally considered to be primarily responsible for decision-making, the setting of Internet standards and protocols, the management of the DNS, and for other aspects of Internet governance.

The Internet Society (ISOC)

The Internet Society (ISOC), a non-profit corporation based in Reston Virginia (suburb of Washington, DC), "is a professional membership society with more than 100 organization and over 20,000 individual members in over 180 countries. It provides leadership in addressing issues that confront the future of the Internet, and is the organizational home for the groups responsible for Internet infrastructure standards, including the Internet Engineering Task Force (IETF) and the Internet Architecture Board (IAB)."[70]

As noted above, the founders and architects of the Internet, the TCP/IP protocols, the DNS system, and related standards and technologies, were all of the ARPANET/NSFNet hierarchy. Seeking to install some form of legal structure to their efforts and network (both in a literal and figurative sense), the ISOC was born. "Beginning in 1991 the ARPA cadre did attempt to place an organizational and legal capstone around their efforts. They founded a private, nonprofit organization called the Internet Society...."[71]

Established in 1992, ISOC is something of an oversight and coordinating body for the organizations that play a role in protocol and standard setting, namely, the Internet Architecture Board (IAB), Internet Engineering Task Force (IETF), Internet Engineering Steering Group (IESG), Internet Research Steering Group (IRSG), the Internet Research Task Force (IRTF), and the RFC Editor, which are considered to be its "component bodies."[72] The relations of each respective entity will be explained in more detail, below. ISOC's role is to foster global cooperation and coordination on Internet issues, and its membership is composed of companies, government agencies, foundations and individuals active in the open development, testing, and innovation of Internet standards, protocols, and related infrastructure technologies. "Since ISOC's formation, other Internet organizations have accepted it as the over-arching Internet authority."[73]

Internet Architecture Board (IAB)

Originally chartered as the Internet Activities Board in 1983, the IAB "replaced a standing advisory committee for DARPA's Internet program that had been around since 1979,"[74] namely, the Internet Configuration Control Board (ICCB).[75] Vint Cerf established the ICCB "as an informal committee to advise DARPA and to guide the technical evolution of the protocol suite.[76] The IAB "develops guidelines for research into problems with and improvements in the Internet's architecture. The IAB members are the Internet Engineering Task Force ("IETF") Chair plus twelve other trustees selected by the IETF nominating committee and approved by ISOC."[77]

The IAB oversees "big picture" aspects of the Internet, and oversees the IETF's long-term planning and coordination, and also serves as a liaison or clearinghouse for conveying information to relevant persons or organizations.[78] According to its Web site, "[t]he IAB is chartered both as a committee of the Internet Engineering Task Force (IETF) and as an advisory body of the Internet Society (ISOC). Its responsibilities include architectural oversight of IETF activities, Internet Standards Process oversight and appeal, and the appointment of the RFC Editor. The IAB is also responsible for the management of the IETF protocol parameter registries."[79] The IAB formally came under the ISOC umbrella in June of 1992 as one of its components.[80] ISOC mandates the IAB to oversee the architecture of the Internet, including its protocols and other standards.[81]

While the ISOC has oversight and jurisdiction over the IAB, the IAB exercises a considerable amount of independence and control over its activities. "With respect to technology, the IAB is considered to be a committee of the IETF."[82] Indeed, with respect to its composition, voting members of the IAB are nominated by the IETF, and are then appointed by the ISOC Board of Trustees. "The IETF chair, who is chair of IESG as well, is also a voting member. The voting members select one of themselves to serve as chair of IAB. Non-voting members, mainly from associated bodies, also exist. Members serve as individuals, and not as representatives of companies or organizations."[83]

The IAB also plays a role with respect to the actions of the Internet Engineering Steering Group (IESG); "IESG decisions may be appealed to the IAB. IAB rulings are final, with the exception that claims that the IAB proceeded unreasonably may be appealed to the ISOC Board of Trustees. The appointment of an organization as RFC Editor is subject to IAB approval."[84]

As a final comment on the IAB, it is important to note that, in chartering and empowering the IETF, the IAB ceded a considerable amount of its responsibilities and power to the IETF. "The public's role in Internet standard setting became formalized in 1986, when the IAB created the Internet Engineering Task Force (IETF). Although the IAB also created several other task forces at the same time, the IETF was special: the IAB asked it to take over 'the general responsibility for making the Internet work and for the resolution of all short-and mid-range protocol and architectural issues required to make the Internet function effectively.' In so doing, the IAB divested itself of the main

part of the standard-creation work and relegated itself to making the final decisions in a supervisory, appellate, and managerial role. The IETF became the main forum in which the technical standards were proposed, tested, and debated. The IAB remained primarily a reviewing body, with increasingly little direct participation in the standards drafting process.[85]

Internet Engineering Task Force (IETF)

The IETF is perhaps best summed up by RFC 4677, also known as "The Tao of IETF: A Novice's Guide to the Internet Engineering Task Force," as follows:

> "The Internet Engineering Task Force is a loosely self-organized group of people who contribute to the engineering and evolution of Internet technologies. It is the principal body engaged in the development of new Internet standard specifications. The IETF is unusual in that it exists as a collection of happenings, but is not a corporation and has no board of directors, no members, and no dues. Its mission includes the following:
>
> * Identifying, and proposing solutions to, pressing operational and technical problems in the Internet;
>
> * Specifying the development or usage of protocols and the near-term architecture to solve such technical problems for the Internet;
>
> * Making recommendations to the Internet Engineering Steering Group (IESG) regarding the standardization of protocols and protocol usage in the Internet;
>
> * Facilitating technology transfer from the Internet Research Task Force (IRTF) to the wider Internet community; and,
>
> * Providing a forum for the exchange of information within the Internet community between vendors, users, researchers, agency contractors, and network managers."[86]

Quite succinctly, the IETF's Mission Statement, which is found at RFC 3935, states that "[t]he goal of the IETF is to make the Internet work better. The mission of the IETF is to produce high quality, relevant technical and engineering documents that influence the way people design, use, and manage the Internet in such a way as to make the Internet work better."[87] The Mission Statement continues:

> "These documents include protocol standards, best current practices, and informational documents of various kinds. The IETF will pursue this mission in adherence to the following cardinal principles:
>
> * Open process - any interested person can participate in the work, know what is being decided, and make his or her voice heard on the issue. Part

of this principle is our commitment to making our documents, our WG mailing lists, our attendance lists, and our meeting minutes publicly available on the Internet.

* Technical competence - the issues on which the IETF produces its documents are issues where the IETF has the competence needed to speak to them, and that the IETF is willing to listen to technically competent input from any source. Technical competence also means that we expect IETF output to be designed to sound network engineering principles - this is also often referred to as "engineering quality".

* Volunteer Core - our participants and our leadership are people who come to the IETF because they want to do work that furthers the IETF's mission of "making the Internet work better".

* Rough consensus and running code - We make standards based on the combined engineering judgement of our participants and our real- world experience in implementing and deploying our specifications.

* Protocol ownership - when the IETF takes ownership of a protocol or function, it accepts the responsibility for all aspects of the protocol, even though some aspects may rarely or never be seen on the Internet. Conversely, when the IETF is not responsible for a protocol or function, it does not attempt to exert control over it, even though it may at times touch or affect the Internet. "[88]

The notion of "rough consensus and running code" has become the IETF's mantra, its trademark. "Unlike standards communities grounded in coalitions of vendors or carriers, the early IETF considered interoperability and empowerment of the end user to be basic norms. The standards themselves were nonproprietary. All documentation was open, noncopyrighted, and freely available online. The community 'believes that the value of technical ideas should not be decided by vote but by empirical proof of feasability or, in the language of the engineers, by running code.' The community's political modus operandi was reflected in its famous credo, coined by David Clark in 1992: 'We reject presidents, kings and voting; we believe in rough consensus and running code.'"[89]

How is "rough consensus" measured? Certainly not by formal voting, as is explained in the Tao of IETF:

"Another aspect of Working Groups that confounds many people is the fact that there is no formal voting. The general rule on disputed topics is that the Working Group has to come to 'rough consensus,' meaning that a very large majority of those who care must agree. The exact method of determining rough consensus varies from Working Group to Working Group. Sometimes consensus is determined by 'humming' -- if you agree with a proposal, you hum when prompted by the chair; if you disagree, you keep your silence. Newcomers

find it quite peculiar, but it works. It is up to the chair to decide when the Working Group has reached rough consensus

The lack of formal voting has caused some very long delays for some proposals, but most IETF participants who have witnessed rough consensus after acrimonious debates feel that the delays often result in better protocols. (And, if you think about it, how could you have "voting" in a group that anyone can join, and when it's impossible to count the participants?) Rough consensus has been defined in many ways; a simple version is that it means that strongly held objections must be debated until most people are satisfied that these objections are wrong."[90]

Internally, most IETF work is done in its working groups, which are organized by topic into several areas (e.g., routing, transport, security, etc.), and conducted for the most part via electronic discussion groups of "mailing lists." The IETF holds meetings several times a year. Its working groups are organized by "areas," led by Area Directors, or ADs, whom are also members of the Internet Engineering Steering Group (IESG).[91]

As noted above, the IAM provides architectural oversight for the IETF. "The IAB also adjudicates appeals when someone complains that the IESG has failed. The IAB and IESG are chartered by the Internet Society (ISOC) for these purposes. The General Area Director also serves as the chair of the IESG and of the IETF, and is an ex-officio member of the IAB."[92]

Internet Engineering Steering Group (IESG)

Formed in 1989, the Internet Engineering Steering Group (IESG) is the IETF's internal management arm, within the ISOC umbrella. At the helm of the IESG is the IETF/IESG chair, and IESG decisions are subject to appeal to the IAB. IESG's membership includes the IETF Area Directors, the above-referenced IETF/IESG chair, and a small number of ex-officio and liaison members. In addition to creating IETF working groups, the IESG handles the vetting and approval of all IETF standards, and generally manages the standards process according to guidelines set forth by the ISOC Trustees.[93]

The IESG's role within the Internet standards process is described in great detail in Section 6 of RFC 2026. "The mechanics of the Internet Standards Process involve decisions of the IESG concerning the elevation of a specification onto the standards track or the movement of a standards-track specification from one maturity level to another. Although a number of reasonably objective criteria ... are available to guide the IESG in making a decision to move a specification onto, along, or off the standards track, there is no algorithmic guarantee of elevation to or progression along the standards track for any specification. The experienced collective judgment of the IESG concerning the technical quality of a specification proposed for elevation to or advancement in the standards track is an essential component of the decision-making process."[94]

Internet Research Task Force (IRTF)

As its name implies, the primary function of the IRTF is research, primarily into more longer-ranging, abstract areas of Internet research, including Internet protocols, applications, architecture and technology.[95] Indeed, "[w]hile the IETF focuses on engineering and standards, the IRTF focuses on research. The IRTF investigates Internet topics that are too uncertain or too advanced to be standardized at the moment. When IRTF produces a specification that is suitable for standardization, it is processed via IETF."[96]

As another of the IAB organizations, the IRTF was established in 1989, and is divided into various Research Groups (RGs), each having a Chair. The IRTF Chair may remove an RG Chairs, however this decision may be appealed to the IAB. The IRTF Chair is appointed by the IAB, and reports thereto. Consistent with its longer-view, Big Picture approach, IRTF Research Groups are generally of a longer duration. [97]

IRTF Research Groups generally are of a longer duration, in terms of membership continuity, in order to promote the development of research collaboration and teamwork in exploring research issues. Members also tend more to be individual contributors, rather than representatives of companies or organizations. The IRTF is managed by the IRTF Chair (appointed by the IAB) in consultation with the Internet Research Steering Group (IRSG). The IRSG membership includes the IRTF Chair, the chairs of the various Research Group and possibly other individuals ("members at large") from the research community.[98]

At the time of this writing, there were 14 active Research Groups engaged in research activities, as follows:

Anti-Spam Research Group (ASRG)
Crypto Forum Research Group
Delay-Tolerant Networking Research Group (DTNRG)
End-to-End Research Group Charter
Internet Measurement Research Group
IP Mobility Optimizations (Mob Opts) Research Group
Network Management Research Group Charter (NMRG)
Peer-to-Peer Research Group
Routing Research Group
Transport Modeling Research Group
Internet Congestion Control Research Group
Scalable Adaptive Multicast Research Group
End Middle End Research Group
Host Identity Protocol (HIP) Research Group[99]

Internet Research Steering Group (IRSG)

Formed in 1989, the Internet Research Steering Group (IRSG) is the management and oversight arm of the IRTF. Led by the IRTF Chair (who reports to the

IAB, and manages in consultation), the IRSG is made up of the respective chairs of the various IRTF Research Groups (such as those named, above), as well as at-large members from the research community, as appropriate.[100]

RFC Editor

The RFC Editor is the entity in charge of editing, managing, publishing and maintaining the authoritative archive for the Request For Comments (RFC) series of documents, which are the Internet's documents of record. As noted above, the RFC process began in 1969, initiated by graduate student Steve Crocker of UCLA, and was originally housed within the SRI Network Information Center. It later moved to the Information Sciences Institute (ISI) at the University of Southern California (USC). There, Jon Postel ran the RFC Editor for decades until his passing in October, 1998. It remains at ISI, operating under the auspices of ISOC, which appoints an organization as RFC Editor on the recommendation of IAB. The IAB vets the general policy followed by the RFC Editor.[101]

Internet Corporation for Assigned Names and Numbers (ICANN)

As noted above, ICANN retains the non-competitive aspects of the so-called "IANA function." In particular, and pursuant to its Bylaws, ICANN:

1. Coordinates the allocation and assignment of the three sets of unique identifiers for the Internet, which are

    a. Domain names (forming a system referred to as "DNS");

    b. Internet protocol ("IP") addresses and autonomous system ("AS") numbers; and

    c. Protocol port and parameter numbers.

2. Coordinates the operation and evolution of the DNS root name server system.

3. Coordinates policy development reasonably and appropriately related to these technical functions.

ICANN Structure

ICANN is led by its Chairman and Board of Directors, which are guided by a number of advisory committees, Supporting Organizations, and other offices. The following entities serve as part of the governing structure of ICANN, all in support of and/or in conjunction with the ICANN Board of Directors:

President/CEO, which is supported by ICANN staff;
Governmental Advisory Committee (GAC)

At-Large Advisory Committee (ALAC)
Security and Stability Advisory Committee (SSAC)
Root Server System Advisory Committee (RSSAC)
Technical Liaison Group (TLG)
Country Code Names Supporting Organization (ccNSO)
Generic Names Supporting Organization (GNSO)
Address Supporting Organization (ASO)
Ombudsman

A complete list, organizational chart, and further information on ICANN's structure may be found on the ICANN Web site, at http://www.icann.org/structure/. These entities provide information and feedback to the Board on the relevant issues of the day, offering international perspective from the technical, commercial, governmental, academic and user-oriented communities.[102] In addition to the above, ICANN empanels numerous Working Groups, Task Forces, and other groups charged with responsibility for researching and discussing relevant issues of the day, and then making report and recommendations to the respective ICANN entity (such as the ccNSO, GNSO, etc.). For example, at the time of this writing, ICANN working groups, task forces, and other components are discussing the impact posed by, and potential policy responses to, issues such as the privacy implications of domain name registration information made available under the WHOIS service; domain "tasting," implementation of so-called International Domain Names (IDNs, which rely on non-ASCII text, such as Chinese characters), and implementation of Internet Protocol Version 6 (ITV6). For a complete list of current issues being considered before ICANN, see the Current Topics list, found at http://www.icann.org/topics/.[103] ICANN holds meetings three times a year at locations around the globe.[104]

ICANN Function, continued

Pursuant to Section 3 of ICANN's Articles of Incorporation (as revised 21 November 1998), ICANN is authorized to:

> "pursue the charitable and public purposes of lessening the burdens of government and promoting the global public interest in the operational stability of the Internet by (i) coordinating the assignment of Internet technical parameters as needed to maintain universal connectivity on the Internet; (ii) performing and overseeing functions related to the coordination of the Internet Protocol ("IP") address space; (iii) performing and overseeing functions related to the coordination of the Internet domain name system ("DNS"), including the development of policies for determining the circumstances under which new top-level domains are added to the DNS root system; (iv) overseeing operation of the authoritative Internet DNS root server system; and (v) engaging in any other related lawful activity in furtherance of items (i) through (iv)."[105]

In doing so, Section 4 of the Articles states that ICANN "shall operate for the benefit of the Internet community as a whole, carrying out its activities in conformity with relevant principles of international law and applicable international conventions

and local law and, to the extent appropriate and consistent with these Articles and its Bylaws, through open and transparent processes that enable competition and open entry in Internet-related markets. To this effect, the Corporation shall cooperate as appropriate with relevant international organizations."

Internet Assigned Numbers Authority (IANA)

The "IANA function," as described above, has evolved from a one-man operation, in the name of Jon Postel (both at Stanford and at USC), to a shared-responsibility between IANA (based in Southern California) and NSI (in Herndon, Virginia), to the current state, which entails oversight of the root server system, and the policy and coordination responsibilities currently performed by ICANN. In particular, "IANA oversees IP address allocation, manages the DNS (this includes root server system oversight and top-level domain delegation), and coordinates protocol parameter assignment. All Internet domain names and IP addresses are allocated from IANA, either directly or, much more likely, indirectly through IANA's delegation of authority via a worldwide system of Internet registries and registrars."[106]

From the IANA Web site:

> One of the IANA functions that ICANN performs is to keep the root-zone up to date and maintain the authoritative Whois database for generic top-level domains (gTLDs) and country code top-level domains (ccTLDs). The following information is provided with the goal of assisting the Internet community to understand the process that is followed in addressing requests to add/change information in the root-zone and/or in the authoritative Whois database. Collectively, the end result of the processing of these requests is referred to as "Root Management." These requests include the following:
>
> * Primary nameserver changes;
> * Secondary nameserver changes;
> * Change of administrative and technical contact;
> * Change of data for the sponsoring organization, administrative Contact, and technical Contact;
> * Change of data for the URL for registration services and port 80 whois server;
> * Redelegations of oversight and management authority for TLDs

The present-day IANA function is performed by ICANN, pursuant to contract between the Commerce Department and ICANN, entered into on 21 March 2001.[107]

VeriSign/Network Solutions, Inc. (NSI)

Just as IANA's role and function has evolved over the years, so, too, has that of NSI. Beginning in 1993, NSI carried out the registry and registrar function at the gTLD level, as described above. As the Shared Registry System has been implemented, NSI's responsibilities have been scaled back. NSI, a subsidiary of VeriSign, Inc., currently

serves as the sole registry for the .com and .net gTLDs, and also operates as a registrar for a number of gTLDs. NSI also offers a suite of value-added proprietary services, such as web hosting, web site design, WHOIS services, and related services.[108]

World Wide Web Consortium (W3C)

The World Wide Web Consortium (W3C) develops common protocols for the evolution of the World Wide Web (WWW), to ensure its interoperability. The W3C develops, tests and implements new standards for use on the WWW. While similar to the IETF in scope and operation, its focus is limited to that of the WWW and related technologies. The W3C is similar to the IETF in that it develops technological standards, but its focus is more tightly directed, at the Web and associated technologies. The W3C was founded by the creator of the World Wide Web, Tim Berners-Lee in, October, 1994 at the Massachusetts Institute of Technology (MIT), in the U.S., in conjunction with the European Organization for Nuclear Research (CERN), located in Geneva, Switzerland.[109]

1. Marcus Maher, "An Analysis of Internet Standardization" (1998), 3 Va. J.L. & Tech. 5 (Spring, 1998)

2. Charles Siegal, "Rule Formation in Non-Hierarchical Systems" (1998), 16 Temp. Envtl. L. & Tech. J. 173, 179.

3. "Rule Formation," above, 179.

4. "Rule Formation," above, 180.

5. "An Analysis of Internet Standardization," above, at 9.

6. "An Analysis of Internet Standardization," above, at 10.

7. "Rule Formation," above, at 180.

8. "An Analysis of Internet Standardization," above, at 10.

9. "An Analysis of Internet Standardization," above, at 10.

10. "An Analysis of Internet Standardization," above, at 10.

11. "An Analysis of Internet Standardization," above, at 7-8. See also, http://www.w3.org/History.html (Current as of 25 July 2007).

12. Ben Segal, "A Short History of Internet Protocols at CERN" (April, 1995), found at http://ben.home.cern.ch/ben/TCPHIST.html (Current as of 25 July 2007).

13. Karl. M. Manheim & Lawrence B. Solum, "An Economic Analysis of Domain Name Policy" (Spring/Summer, 2003), 25 Hastings Comm. & Ent. L.J. 359, at 373.

14. Jonathan Weinberg, "ICANN and the Problem of Legitimacy" (October, 2000), 50 Duke L.J. 187, 194.

15. Jack Goldsmith and Tim Wu, "Who Controls the Internet? Illusions of a Borderless World," Oxford University Press (2006), at 33.

16. "Who Controls the Internet?" above, at 33-34.

17. Milton Mueller, "Ruling the Root: Internet Governance and the Taming of Cyberspace," MIT Press, Cambridge, Mass. (2002, 2004), at 93

18. "ICANN," above, at 194.

19. "ICANN," above, at 195.

20. "ICANN," above, at 196.

21. Peter K. Yu, "The Origins of ccTLD Policymaking,"(Fall, 2004), 12 Cardozo J. Int'l & Comp. L. 387, 390-391.

22. Karl. M. Manheim & Lawrence B. Solum, "An Economic Analysis of Domain Name Policy" (Spring/Summer, 2003), 25 Hastings Comm. & Ent. L.J. 359, at 372.

23. "Ruling the Root," above, at 79. See also, "The Origins of ccTLD Policymaking," above, at 390.

24. "The Origins of ccTLD Policymaking," above, at 390.

25. "Ruling the Root," above, at 79.

26. RFC 1591, found at http://tools.ietf.org/html/rfc1591

27. "An Economic Analysis," above, at 379.

28. Joseph P. Liu, "Legitimacy and Authority in Internet Coordination: A Domain Name Case Study" (Spring, 1999), 74 Ind. L.J. 587, at 590, 591.

29. A. Michael Froomkin, "Wrong Turn in Cyberspace: Using ICANN to Route Around the APA and the Constitution" (2000), 50 Duke L.J. 17, at 17.

30. List of root servers, http://www.root-servers.org/ (current as of 25 July 2007); Markus Mullern, "Who Owns the Internet? Ownership As a Legal Basis for American Control of the Internet" (Spring, 2005), 15 Fordham Intell. Prop. & Ent. L.J. 709, at 713-714.

31. "Who Owns the Internet?," above, at 714-715.

32. "ICANN and the Problem of Legitimacy," above, at 198.

33. "ICANN and the Problem of Legitimacy," above, at 198. See also, "Management of Internet Names and Addresses," also known as "The White Paper," 3 June 1998, found at

http://www.ntia.doc.gov/ntiahome/domainname/6_5_98dns.htm, current as of 28 July 2007.

34. "ICANN and the Problem of Legitimacy," above, at 198-199, also, "Who Controls the Internet," above, at 35.

35. "Ruling the Root," above, at 102.

36. "Ruling the Root," above, at 157.

37. "Presidential Directive on Electronic Commerce," July 1, 1997, at http://www.technology.gov/digeconomy/presiden.htm, current as of 27 July 2007.

38. "Request for Comments on the Registration and Administration of Internet Domain Names," http://www.ntia.doc.gov/Ntiahome/Domainname/DN5NOTIC.htm, current as of 27 July 2007.

39. "Request for Comments on the Registration and Administration of Internet Domain Names," Background section, above.

40. NSF Cooperative Agreement No. NCR-9218742, art. 7, also, Christopher P. Rains, "Comment: A Domain By Any Other Name: Forging International Solutions for the Governance of Internet Domain Names" (Spring, 2000), 14 Emory Int'l L. Rev. 355.

41. "ICANN and the Problem of Legitimacy," above, at 205, also, see http://www.ntia.doc.gov/ntiahome/domainname/dnsdrft.htm, current as of 27 July 2007.

42. "ICANN and the Problem of Legitimacy," above, at 205-206, also, see http://www.ntia.doc.gov/ntiahome/domainname/dnsdrft.htm, current as of 27 July 2007.

43. "Ruling the Root," above, at 64. See also, "Management of Internet Names and Addresses," also known as "The White Paper," issued 3 June 1998, found at http://www.ntia.doc.gov/ntiahome/domainname/6_5_98dns.htm

44. "Ruling the Root," above, at 64. See also, "Management of Internet Names and Addresses," also known as "The White Paper," issued 3 June 1998, found at http://www.ntia.doc.gov/ntiahome/domainname/6_5_98dns.htm.

45. "The White Paper," above.

46. "The White Paper," above.

47. "The White Paper," above.

48. "Ruling the Root," above, at 179. Also, "ICANN and the Problem of Legitimacy," above, at 209-120.

49. "Ruling the Root," above, at 184.

50. "The White Paper," above.

51. The Green Paper, above, see "The Competitive Functions." Also, "The White Paper," above, see "The Transition."

52. The "White Paper," above, see "The Transition."

53. "Who Controls the Internet," above, at 35-36. Also, "Ruling the Root," above, 180-182.

54. "The Green Paper," above, see "The Competitive Functions."

55. "The Green Paper," above, see "The Competitive Functions."

56. "Ruling the Root," above, at 183.

57. See The InterNIC's "Accredited Registrar Directory," at http://www.internic.net/regist.html, current as of 27 July 2007. See also, "ICANN Accredited Registrars," at http://www.icann.org/registrars/accredited-list.html, current as of 27 July 2007. See also, ICANN Registry Listing, at http://www.icann.org/registries/listing.html, current as of 27 July 2007.

58. See "InterNIC nAccredited Registrar Directory," at http://www.internic.net/regist.html, current as of 27 July 2007. See also, "ICANN Accredited Registrars," at http://www.icann.org/registrars/accredited-list.html, current as of 27 July 2007. See also, ICANN Registry Listing, at http://www.icann.org/registries/listing.html, current as of 27 July 2007.

59. ICANN Bylaws, at http://www.icann.org/general/bylaws.htm#I, current as of 27 July 2007.

60. A. Michael Froomkin, "Habermas@Discourse.Net: Toward a Critical Theory of Cyberspace"(January, 2003), 116 Harv. L. Rev. 749, at 783-784.

61. "Habermas," above, at 784.

62. "Habermas," above, at 784. Also, see RFC 1000, supra note 134, at 3 (quoting Stephen D. Crocker, first RFC editor).

63. "Habermas," above, at 784.

64. "Ruling the Root," above, at 32.

65. "Ruling the Root," above, at 32.

66. "Legitimacy and Authority in Internet Coordination," above, at 595-596.

67. "ICANN," above, at 193.

68. Josh A. Goldfoot, Antitrust Implications of Internet Administration, 84 Va. L. Rev. 909 (August 1998), at 910.

69. Alex Simonelis, "A Concise Guide to the Major Internet Bodies," *Ubiquity* magazine, Volume 6, Issue 5, (February 2005), at http://www.acm.org/ubiquity/views/v6i5_simoneli.html (current as of 27 July 2007).

70. ISOC Web site, at http://www.isoc.org/isoc/, current as of 27 July 2007.

71. "Ruling the Root," above, at 94.

72. Alexander Gigante, "Blackhole in Cyberspace: The Legal Void in the Internet" (Spring, 1997), 15 J. Marshall J. Computer & Info. L. 413, at 416-418; also, "A Concise Guide to the Major Internet Bodies," above.

73. "Blackhole in Cyberspace," above, at 418.

74. "Ruling the Root," above, at 90.

75. "A Concise Guide," above.

76. "Habermas," above, at 785.

77. "Blackhole in Cyberspace," above, at 417.

78. "Tao of IETF, A Novice's Guide to the Internet Engineering Task Force," RFC 4677, September 2006.

79. Internet Architecture Board Web site, at http://www.iab.org/ (Current as of 27 July 2007).

80. "Habermas," above, at 788.

81. "A Concise Guide," above.

82. "A Concise Guide," above.

83. "A Concise Guide," above.

84. "A Concise Guide," above.

85. "Habermas," above, at 787.

86. "Tao of IETF," above.

87. RFC 3935 (October, 2004), at http://www.faqs.org/rfcs/rfc3935.html, current as of 27 July 2007.

88. RFC 3935, above.

89. "Ruling the Root," above, at 91.

90. Tao, above.

91. "Overview of the IETF," http://www.ietf.org/overview.html, current as of 17 July 2007

92. "Overview of the IETF," above.

93. "A Concise Guide," above.

94. RFC 2026, "The Internet Standards Process – Revision 3," October, 1996.

95. "Blackhole," above, at 417. Also, see IRTF Web site, at www.irtf.org, current as of 27 July 2007.

96. "A Concise Guide," above.

97. "A Concise Guide," above.

98. IRTF Web site, found at www.irtf.org, current as of 27 July 2007.

99. IRTF Web site, at http://www.irtf.org/groups, current as of 27 July 2007.

100. "A Concise Guide," above.

101. "A Concise Guide," above.

102. "A Concise Guide," above.

103. ICANN Web site, "Current Topics," at http://www.icann.org/topics/, current as of 27 July 2007.

104. ICANN Web site, "Meetings," see http://www.icann.org/meetings/, current as of 27 July 2007.

105. ICANN Articles of Incorporation, as revised 21 November 1998, found at http://www.icann.org/general/articles.htm, current as of 27 July 2007.

106. "A Concise Guide," above.

107. "Contract Between ICANN and the United States Government for Performance of the IANA Function," 21 March 2001, http://www.icann.org/general/iana-contract-21mar01.htm, current as of 27 July 2007.

108. Network Solutions web site, at http://www.nsi.com, current as of 27 July 2007.

109. "A Concise Guide," above.