

TOWARD AN ECONOMICS OF THE DOMAIN NAME SYSTEM

MILTON MUELLER

Syracuse University School of Information Studies

1. Introduction

In August 2000, an arbitration panelist of the World Intellectual Property Organization (WIPO) ordered the transfer of the domain name <barcelona.com> to the Barcelona City Council. The domain name had been registered by a New York-based start-up company, Barcelona.com, Inc., back in 1996. The company planned to develop the domain into an Internet portal for Barcelona residents, tourists and enthusiasts. Barcelona.com was in the process of negotiating with potential investors when the City filed a complaint under ICANN's Uniform Domain Name Dispute Resolution Policy.

The Barcelona City Council based its complaint on the city's many registered trademarks that included the word 'Barcelona,' many of which had just been registered in 1996 or 1995. Although Spanish trademark law prohibits the registration of place names, the City claimed that its trademarks, such as '*Barcelona Excelentísimo Ayuntamiento de Barcelona*,' could be confused with the domain name <barcelona.com>. The City Council also claimed that the entrepreneurs were cybersquatters who had registered the name in bad faith only to sell it to the city. WIPO panelist Marino Porzio agreed, and ordered the Internet entrepreneurs to hand over their name. Its prospective portal business was destroyed.

Years of court challenges and appeals over the right to the name followed. Finally, on June 2nd, 2003, a U.S. Court of Appeals denied all of the City Council's arguments regarding trademark rights, bad faith use of the domain and cybersquatting. A Spanish municipality was told by a U.S. court that it had no

rights to the name under Spanish or U.S. law. It was declared a reverse domain name hijacker, and ordered to restore the name to the original registrant. On August 16th, 2003, the domain was transferred back to Barcelona.com, Inc.

The Barcelona.com conflict was only one of thousands caused by the rise of the domain name industry after the rise of the World Wide Web in 1993. The rise of a domain name market has produced conflicts over trademark rights, personal names, misspellings of famous names, legal jurisdiction, and the creation of new generic top-level domains such as *.web*. In essence, the Internet domain name space created a new kind of value, and law and institutions had to make rapid and major adjustments to human efforts to capture that value. This chapter focuses on the economics of the domain name system. The Internet has commoditized identifiers on an unprecedented scale, creating global identifier-related services markets with their own unique economic characteristics. In the market for names, technical functions interact with semantic meaning in complex and sometimes explosive ways.

The term “identifier” refers to a string of characters or symbols that can be used to give a name to users or objects on a network, such as servers, routers, printers, documents, or multimedia content. Unique identifiers serve a critical role in most networks or online environments. (Berners-Lee 1996; Berners-Lee 1998; Paskin 1999) By uniquely identifying an object on a network, names make them retrievable or accessible. (Saltzer 1993) A telephone number is an identifier, and so is an ISBN number (the international code for books used by publishers). An Internet Protocol (IP) address is an identifier used by routers to determine how to get packets from their origins to their destination. This chapter, however, concentrates exclusively on Internet domain names. What makes domain names of particular interest is that they (unlike IP addresses or ISBN codes) have evolved into things that can be bought and sold in a market. The supply of domain names to Internet users generates an estimated \$2.5 billion annually in revenues. Despite the economic significance of this market and its centrality in provoking policy debates about Internet governance, there is surprisingly little literature about the economics of domain names. (FTC 1998; Rood 2000; Mueller 2002a)

Domain names are composed of alphanumeric characters in a hierarchical structure. It was the rise of the World Wide Web after 1994 that created most of the demand for domain name registrations.

(Mueller 2002a, Chapter 6) Three-level domain names became synonymous with Web site locators (URLs). Domain names began to appear in television ads, on the sides of buses, on business cards. For a short time, guessing the domain name associated with an organization in order to find its Web site became a common way of navigating the Internet. In late 1995, responding to the rush of web-inspired domain name registrations, the National Science Foundation authorized the US-based domain name registry Network Solutions, Inc. (NSI) to charge annual fees for registering and supporting domain names. The market for registrations doubled in size annually during the Internet boom of the second half of the 1990s.¹ During its frenzied peak in the first two quarters of 2000, revenues grew at a 150% annual growth rate. Responding to artificial scarcity in the supply of names, robust secondary markets for domain names began to evolve. Speculators registered hundreds of names and advertised them for sale to the highest bidder.

(Table 1)

In line with the general depression in Internet-related markets, the number of domain name registrations worldwide declined by about 8 percent from October 2001 to June 2002, as large numbers of speculative and promotional registrations in the generic top-level domains <.com>, <.net>, and <.org> were not renewed. Domain name registrations in most country-code TLDs, however, still grew at double-digit rates, and so did the new top-level domains <.info> and <.biz>. The domain name market as a whole returned to positive if modest growth in the last quarter of 2002. About 46.5 million domain names are registered worldwide as of the third quarter of 2002.

The significance of domain names goes beyond their relatively modest economic niche in the market for Internet-related services. The DNS sits at the heart of the Internet infrastructure. Its technical functions are critical to the ability of users to establish connections to each other, just as consistent use of telephone numbers is critical to interoperability of the telephone system. Because control of the top level of the domain name hierarchy gives its administrator considerable power over the domain name industry and its customers, commercialization of domain name services stimulated major institutional changes in international governance arrangements. (Mueller 2002a) The US government and key corporate actors in the Internet business desired to place domain name regulation into a global jurisdictional framework and to avoid control of the DNS by traditional intergovernmental organizations. (NTIA 1998) This led to the

¹ See NSI Annual Report, 2000.

“privatization” of the regulatory responsibilities, policymaking and coordinating functions associated with domain names. Much of this happened because the semantic properties of domain names led to legal and political conflicts over trademark rights (and other types of claimed rights to names), resulting in new international rules and rulemaking mechanisms for arbitrating rights to names.

Aside from that, how things are named on the Internet – a problem that involves many actual or proposed standards and services in addition to DNS – profoundly affects the Net’s usability and efficiency. Many of the economic and policy issues initially raised by DNS are being or might be repeated by the growth of other naming systems. (Bechtold 2002) Thus, experience with DNS may be taken as an example (for better or worse) of what economic issues to anticipate with different naming designs.² The DNS protocol also may play a role in new information and telecommunication services. The ENUM protocol, for example, uses the DNS to map telephone numbers to IP addresses, facilitating efficient interconnection of IP telephony systems and traditional telephone networks using E.164 numbering as identifiers. (Huston 2002; Hwang and Mueller 2002) Already, controversies regarding DNS economics and policy have spilled over into the implementation of ENUM. (Cannon 2001; McTaggart 2001; Rutkowski 2001)

The market for domain names and related services will grow as the Internet diffuses. The institutional changes provoked by the growth and commercialization of DNS will have a lasting impact on communication industries. It is important, therefore, to understand how the DNS functions, the economic properties of domain names as resources, the supply of domain name services, and the industrial organization characteristics of the domain name supply industry and its regulator, the Internet Corporation for Assignment Numbers (ICANN).

The exposition is divided into the following parts. First, a basic technical description of the assignment and resolution of domain names is provided. The next section describes the demand characteristics of domain names. Next, the characteristics of the supply industry are examined. Finally, the chapter looks at the policy issues associated with the industry. Most of the focus is on the economics of demand and supply, which are the least-analyzed aspects of the domain name problem. A great deal of

² For example, the “Handle” system of digital object identifiers also has a centralized root structure like DNS, which may lead to governance issues if it becomes widely adopted. On the other hand, the handle system is designed to permit unlimited expansion of the top level. Kahn, R. and R. Wilensky (1995). A Framework for Distributed Digital Object Services, Corporation for National Research Initiatives. **2002**.

policy has been made and debated with little application of economic theory and no empirical research into the underlying economics. One purpose of this chapter is to identify and call attention to those areas where research is needed.

2. Technical description of DNS.

Domain names most likely have gone unnoticed by economists because before one can understand the economics one must make sense of a complicated Internet protocol known as DNS (Domain Name System) and its implementation by Internet service providers. (Mockapetris 1987; Albitz and Liu 1992) What follows is a brief overview of the technical structure of DNS. The DNS protocol specifies a name space, permitted structures and characters for names, formats for storing data associated with names, and a mechanism for retrieving that data in client-server interactions. The two basic aspects of domain name are assignment (the process of delegating exclusive control of a name within the DNS name space to a responsible party) and resolution (the process of discovering what IP address is associated with a given domain name).

2.1 The Name Space and Name Assignment

Domain names are alphanumeric strings used to name computers on the Internet. Domain names are most often seen as parts of email addresses and Web URLs. In the email address vogelsang@bu.edu, <bu.edu> is a domain name registered by Boston University and <vogelsang> is a user ID assigned by BU. The name <www.nokiagirls.com> is a three-level domain name that functions as a web URL. Domain names are organized hierarchically, with each level representing the delegation of responsibility from one organization or entity to another. The starting point of the hierarchy is known as “the root,” which is unnamed. For example, the operator of the <com> domain was delegated the right to register names under <com> from the authority over the DNS root; the owner of the <nokiagirls> domain was delegated that authority by the <.com> domain, and the web server <www> was so named by the owner of the <nokiagirls> domain. In the most commonly used notation, dots separate the different levels of the hierarchy. The label farthest to the right is termed the “top-level domain” (TLD), the next level to the left is known as the second-level domain, and so on down the hierarchy. The DNS was designed to permit deep

hierarchies but in actual practice it is rare to see more than four levels used, and most of the real work is done by name servers at the top and second levels.

In purely technical terms, the supply of unique domain names is inexhaustible. Within any single top-level domain, second-level domain names can use 63 character spaces and 37 available characters, meaning that there are 37^{63} possible combinations of characters available. This is an unimaginably large number, and that array is available in one level of the hierarchy alone. The DNS protocol permits the hierarchy to go down more than 100 levels. Moreover, there could be thousands more top-level domains. Of course, huge parts of the available pool of names are undesirable from a human point of view, being either meaningless or too long to be usable, or both.

To function properly as a global identifier, each domain name must be unique and exclusively assigned. Early attempts to have one central naming authority coordinate all name assignments broke down under the weight of constant increases in the number of host computers on the network. (Vixie 1994) The DNS solved this problem by creating a *hierarchical* name space and delegating the authority to assign domain names down the levels of the hierarchy. In this arrangement, a central authority is responsible for assigning unique and exclusive names at the top of the hierarchy. Each assignee of a top-level name then takes responsibility for coordinating the assignment of unique and exclusive second-level names. The owner of a second-level name can in turn assign unique third-level names, and so on.

2.2 Resolution, Name Servers and BIND software

As unique identifiers, domain names are sometimes referred to as Internet “addresses.” But the *real* addresses that guide the routing of data packets are Internet Protocol (IP) addresses, 32-bit binary numbers uniquely assigned to a host. The domain name serves as a higher-level, meaningful name. In order to route information across the Internet a domain name must be mapped to its corresponding IP address, a process known as “resolving a domain name,” or as “resolution.” That layer of indirection serves two purposes. First, semantically meaningful names are easier to key in and remember than strings of numbers. Second, higher-level names can provide a more stable, persistent identifier than an IP address. The IP addresses assigned to groups of computers are determined by the topology and routing practices of the Internet, which tend to change relatively frequently as networks grow, shrink, or are reorganized. By

separating a computer's public identifier from its IP address, and relying on some mapping function to translate between the two, a network manager can give computers stable names which can be used consistently even when their IP addresses change.

The DNS can be characterized as a highly distributed global database, wherein the responsibility for entering, updating, and serving up data on request is handled by hundreds of thousands of independently operated computers known as *name servers*. At the top of the DNS hierarchy, the root-level authority registers top-level domain names (TLDs), and the root servers handle requests to resolve TLDs. The 250 or so top-level domain name registries operate name servers that contain lists of all second-level domain names registered under that top-level domain, and pointers to two or more name servers capable of operating that domain. Thus an economist might view the DNS as a distributed network of compatible name servers, the function of which is to enable assignment and resolution of globally unique names.

Whenever end users use a domain name, their email or browser client (or some other software) must send a query out to resolve the name into an IP address. Typically, one resolves a name by sending a query containing the domain name to the nearest name server operated by an organization or its Internet Service Provider. That name server first checks to see if it has already stored the answer to the query in its cache after responding to a previous query for the same domain name. Thus, the relevant records for the most popular domain names, such as <google.com> and <yahoo.com>, will already be stored and there will be no need to waste time and cycles searching for it. Caching is critical to the scalability of DNS. Most DNS queries never need to go out into the Internet; if they did, the reliability and scalability of the Internet would be impaired. Only if the local name server does not know the answer, or if an earlier record has expired, must it initiate the process of resolving the name.

The resolution of a domain name follows the same hierarchical structure as the assignment of names. A query first goes to the root, which tells the resolver which top-level domain name server to query; the resolver then queries the appropriate top-level name server, which tells it which second-level domain name server stores the records for that domain, and so on.

In principle, anyone can write software that implements the DNS protocol. In reality, an implementation software known as BIND became established in the 1990s and dominates the industry. BIND is open-source software but its revision and release are controlled by a group of IETF veterans

organized as the Internet Software Consortium. BIND's dominance has important economic consequences, particularly regarding the potential for competing DNS roots. The software release contains a list of the IP addresses of the official root servers, which are used as default values by the name server unless those values are manually altered by the name server administrator.

2.3 *The Root Servers*

To make sure that all domain names on the Internet are unique and can be globally resolved, there must be a coordinated source of information about what top level domains exist and at what IP addresses their name servers can be found. That function is performed by the root administrator(s). The root administrator(s) determines what names exist at the top of the DNS hierarchy. The list of recognized TLDs and associated records of their IP addresses is known as the *root zone file*. As we shall see in section 4.1 below, this technical coordination function provides a point of leverage that can also provide control over market entry and the conduct of suppliers and consumers. The root is the one point of centralization on what is otherwise a highly distributed system.

There are 13 official root servers. The existence of multiple root servers provides redundancy in case one root server crashes; they are also geographically distributed in order to reduce query time and expense for users in different parts of the world.³ The number of root servers is technically limited, however, to a maximum of thirteen.⁴ The limit on the number of root servers has important economic consequences. Until some new way of coordinating a larger number of root servers is invented, Internet growth increases the load on existing root servers rather than providing an incentive to supply additional root servers.

3. **The Demand for Domain Names**

Domain names combine technical and human functions. On the technical side, they provide a unique symbol, similar to a telephone number, which facilitates the identification of objects on the network

³ In a recent coordinated denial of service attack on the root servers, 7 of the 13 root servers were disabled, yet most Internet users noticed nothing, or a slight increase in waiting time. David McGuire and Brian Krebs, "Attack on Internet Called Largest Ever," *Washingtonpost.com*, October 22, 2002.

⁴ The limit comes from the ability of a single UDP packet to carry the IP addresses of all the servers.

and helps guide the movement of information from its source to its destination. As words, symbols or phrases that are *meaningful* as well as unique, they also perform psychological and communicative functions, such as making an address easier to remember or more likely to attract attention. The semantic aspects of demand introduce radically different economic properties, and therefore must be distinguished from the technical aspects and analyzed separately.

3.1 Demand for Technical Functions

In order to have an Internet identity end users need an assignment of a unique alphanumeric string devoted to their exclusive use, as well as continuous maintenance of public database records reserving that string and storing the data needed to resolve the name. This aspect of the service may be called the *registration service*. Users of Internet domain names also need a service that, when queried at any time by other users of the public Internet, responds with the information needed to resolve the name. This is called *name service* or *name resolution service*. Registration and name resolution services are needed regardless of whether the domain name is meaningful or not. As noted earlier in the technical section, the layer of indirection provided by a domain name is valuable quite apart from its semantic functions. In this regard, it is worth noting that for hundreds of millions of non English-reading Internet users around the world, the ASCII characters used by the standard DNS protocol are basically meaningless, being based on an alphabet that is mostly illegible to them.

For small-scale consumers, the supply of registration and name resolution services can be efficiently bundled with the provision of basic Internet access and web hosting services. (However, this may increase switching costs, as the ISP has operational control of the name and any attempt to change ISPs might affect the name's functionality. There is a longer discussion of switching costs in 3.1.1 below.) Larger consumers often self-provide their own DNS servers or contract with separate specialized vendors for the various components of the service (registration, name service, Internet access, and hosting).

3.1.1 Portability and Sunk Costs

The costs of the technical services associated with domain names are quite low relative to the total costs of having an Internet presence. Yet the value of an established name can be enormous, when it has

been used for a long time and numerous directories, links or publicity materials direct traffic to the name. Consumers of domain name registration services incur sunk costs when they invest in publicizing their unique identifier (e.g., web site or email address). This restricts their ability to easily change names, which may affect their ability to change service suppliers. As a name (e.g., <firm.com>) becomes well known by external parties or embedded in their stationery, directory listings, and so on, changing to a new domain name (e.g., <firm.biz>) can involve substantial switching costs. Additional investments must be made to inform consumers of the new name. The change may result in foregone sales to consumers who fail to become aware of the new name in time, another switching cost. While it is technologically possible, even easy, to automatically redirect traffic from the old domain name to a new one, this requires retaining a subscription to the old domain for some time. It is theoretically possible, therefore, that a supplier could raise the price to locked-in customers.

In a (non-empirical) study of the significance of sunk costs on competition in the registry market, the U.S. Federal Trade Commission's Competition and Economics Bureau wrote:

The economic analysis of markets with switching costs has identified a number of factors that, in appropriate circumstances, can diminish the ability and the incentive of a supplier to act opportunistically with respect to its locked-in customers: ... (1) the extent to which prospective customers are aware of the possibility of supplier opportunism; (2) the extent to which customers have effective means (e.g., enforceable long-term contracts) to protect themselves against opportunism; (3) the intensity of competition among suppliers; and (4) the importance of reputation and repeat business to suppliers' current and future profits. (FTC 1998)

All four of these factors apply in some degree to the domain name market, although their effectiveness has been limited by artificial constraints on competition imposed by ICANN and the U.S. Department of Commerce. Most customers, particularly those most dependent on their domain name identities, are aware of the prospect of supplier opportunism. (Some small-scale casual consumers may not be.) Long-term contracts are an option, and have become more available as the market has become more competitive. The contractual terms would become more favorable to consumers with additional competition among registries. Competition among suppliers is potentially almost unlimited, although the regulatory restrictions

on entry discussed below limit it. And of course, in a subscription-based service such as domain name registration, repeat business and reputation are critical.

Number portability has been a major policy issue in telephony and there is a significant amount of economic literature on switching costs in telephone numbers. (Andeen and King 1997; NERA 1998; Reiko and Small 1999; Farrell and Klemperer 2001; Gans 2001; Viard 2001) In telephony numbers were assigned to connectivity providers (i.e., carriers) in blocks. Changing one's carrier meant changing one's number, creating a substantial switching cost for local subscribers. Implementation of number portability (in 800 numbers, Local Number Portability and Mobile Number Portability) required investment in and operation of a database-driven number recognition system common to all carriers. The key policy issue is how the costs of this added infrastructure are distributed. The empirical and theoretical literature indicates that number portability facilitates consumer choice and carrier competition. (Gans 2001; Viard 2001) Several authors, however, have raised doubts about the ability of mandated number portability to make the market as a whole more efficient, and have emphasized the need to vest users with property rights in their identifiers. (Reiko and Small 1999)

By way of contrast, under the DNS protocol names have always been virtual resources grounded in a shared database. Thus, domain names have always been portable across connectivity providers unless consumers were ill-informed and allowed their ISP to register the domain in the ISP's name. Moreover, customers have always had stronger, more transferable rights in domain names than in telephone numbers. Even when telephone number portability is implemented, customers cannot reserve them in unlimited quantities without use, cannot engage in free reassignment or sale of them, and have a very limited ability to select a specific number. Domain names in the most popular open TLDs, in contrast, can be reserved without use and resold, and the selection of the desired string is entirely in the hands of the customer. Thus, an important part of the demand for domain names stems from the ability they give consumers to obtain unique identifiers that are portable across connectivity providers, and whose value can be captured and exploited more easily by the user.

Nevertheless, regulation of the domain name market has provided consumers with an additional layer of portability. In the generic top-level domains and some country codes, regulations separate the wholesale "registry" part of the service from the retail "registrar" functions. This makes domain names

portable across registrars as well as across ISPs. (It is not possible to make domain names portable across registries without radical changes to the DNS protocol.) This aspect of the domain name market is a supply side phenomenon and will be discussed in more detail in section 4.2.2 below.

3.1.2 Network externalities (demand-side economies of scope)

Anyone with control over the configuration of a DNS client or name server has the ability to choose which computer they treat as the root of the hierarchy for resolution of domain names.⁵ To some, this suggests that the DNS root need not be a centralized monopoly and that competition could prevail in the provision both of the definition of the root zone file and the operation of root servers. There are, however, strong demand-side network effects in the choice of a DNS root that undermine competition. When Internet user A registers a domain name, she almost always wants that address to be compatible with the DNS implementations of all existing and future communication partners (which we will refer to as *N*). “Compatibility” in this case means that any *N* who encounters A’s domain name will, after sending the appropriate queries to the name servers to which *N* is connected, be returned the *correct* answer about the IP address associated with the A’s domain name. Incompatibility means that no answer, or an incorrect answer, is returned, which makes resolution of the name impossible.

Competing DNS roots can result in varying levels of incompatibility. In general, incompatibility can arise in two ways. First, if A’s and *N*’s name assignments are derived from different, uncoordinated roots, A’s domain name may not be globally unique. If it is not unique, there is a chance that the name servers that handle user *N*’s query will confuse it with some user other than A who has adopted the same name, and return the wrong information. Second, even if the name is unique, the name servers that handle user *N*’s query may not know where to find the information needed to resolve A’s name if it does not begin its query process at the same root as A.

Obviously, the value of a domain name to A increases as the scope of its compatibility widens. An identifier that cannot be resolved by other users is as useless as a telephone set without any connectivity. If only a fragment of A’s universe of communicating parties can resolve the name, A will have to carefully

⁵ In actual practice, default values are set in the implementation software, which for 90% of the world’s computers is a software known as BIND (Berkeley Internet Name Domain). See below.

manage and restrict his use of the identifier and utilize another identifier for use with communication partners who employ an incompatible root. Because so much of the communication on the Internet takes place among parties who do not know each other and have no prior relationships that would allow for coordination, the compatibility barriers created by fragmentation of the name space are high. There is very little reason to ever want to utilize a domain name that is not globally compatible. A single root increases the chance that all domain names will be compatible. Thus, network service providers who want to offer universal connectivity to their customers or clients have a strong incentive to converge on a single DNS root. Consumers of domain name services have a strong incentive to select their names from a single, name space provider with a dominant market share or a monopoly. Of course, competing roots could be coordinated in certain ways, and the effect would be identical to that of an interconnection agreement among competing telecommunication networks. (Mueller 2002c)

3.1.3 Pure Technical Demand vs. Semantic or Human Demand

To conclude the analysis of technical demand factors, in a market unsullied by human factors, the URL xyz.ispname.at/users/cs/hjfg23s.htm is a *perfect* substitute for www.mywebsite.org. In other words, there is no significant difference between a long and meaningless URL and a URL based on a domain name registered specifically to contain a meaningful reference to a person or a website's content. Similarly, in a purely technical market for domain names, the email address me@xx3c.gobbledygook.ispname.com is a perfect substitute for me@myname.com. Of course, when human factors are taken into account, the economic value of those pairs of names is quite different (see below). It follows that if demand for domain name services were based entirely on their technical functions, the market for domain names would be much smaller than it is now. Most individuals and organizations would need only one domain name, and even the largest organizations would need only a few. Domain name users would be willing to get all the additional names they wanted by extending the hierarchy under their existing domain name(s), or by extending the directories of web URLs to the right. There would be little incentive to register names specific to products, people, or a variety of other potential identities (if they were hosted on the same machine). There would be no incentive to speculate in domain names, no incentive to register multiple names for the same purpose, and no incentive to use top-level domain names to differentiate users or

applications. Indeed, there would be little need for end users to participate in the selection of their domain names at all. Assignments could be done by machines, as users would be indifferent as to the specific string they received as long as it was unique and exclusive.

3.2 Demand for Semantic Functions

Most of the market value – and the political and social complications – of domain names stem from the semantic aspects of demand and related human factors considerations. As explained below, the demand for a memorable identifier, the desire to economize on the length of the domain name, concerns about authenticity, and the demand for names that will attract web traffic or aid in search and navigation functions have been powerful influences on the market for domain names.

As the market for domain names has matured it has become increasingly clear that individuals and organizations demand multiple online identities rather than merely one. The multiplication of identities is not a function of imperfect compatibility among systems, as is often assumed. Rather, it reflects end users' different roles and their need to organize their increasingly numerous and complex online interactions. A person on the Internet may have an Internet identity at work, a different Internet identity at home, a hotmail address for when they are on the road, and so on. They probably have multiple identities within their home ISP account, corresponding to different family members or different online persona. They may want to use one identifier on public web sites and email lists and another one for private, priority email in order to foil spammers. They may want to disguise themselves to access illicit websites, or take on new identities for fun. A company wants distinct identifiers for its corporate name, its product and brand names, its departments, task forces, etc. Online identities are created and destroyed by the tens of thousands on a weekly basis. The demand for distinctive identity and the way in which computer technology makes it inexpensive to create, disseminate and erase identifiers combine to fuel semantic demand for names.

The discussion below isolates four aspects of demand related to semantic and human factors:

- Names as memorable identifiers
- Economizing on the length of the name
- Names as authenticators
- Names as search tools or navigation aids

3.2.1 *The Demand for Memorable Identifiers*

The primary driver of domain name demand is the need of Internet users for *memorable identifiers*. The registrant of a domain name uses the name(s) to advertise and communicate their Internet location, and the registrant's communication partners store them or remember them for use when they want to find the registrant's Internet location. Users select names to convey something about themselves; e.g., a cat-lover may adopt <paw> as a domain if given the choice, whereas a critic of corporations might value registering a business's name under a <.sucks> top-level domain to post critical information. The demand for meaning, memorability, or catchiness in domain names all assume that the names are visible to and used by members of the public, rather than by machines alone. The further domain names recede from public view and the more the recognition and storage of Internet identifiers is performed by machines, the less important these functions become as a determinant of demand. However, there are no known models for incorporating semantics into the demand function for identifiers. Here is an interesting problem for a theorist.

The demand for meaning or memorability profoundly affects the consumer's relationship to the available supply of names. It means that the vast majority of available unique strings are of no interest to an individual consumer. Consumers are only interested in combinations of characters that are semantically related to the use they plan to make of the name. We can refer to those clusters of meaningful combinations of strings as "semantic clusters;" e.g., a small-time restaurateur named Joe would be interested in <joe.com>, <joes.com>, <joes-eats.com>, <joes.food>, and many others. Demand-side competition for the right to occupy domain name assignments occurs when there is some degree of overlap between the semantic clusters of two or more users. The alternatives available to prospective registrants are not homogeneous, perfect substitutes. The logic of substitution is based on subjective factors of interpretation; it is fuzzy, not discrete.

The existence of meaning also leads to selection of domain names by customers rather than assignment by suppliers. Of course, in exceptional cases users of driver's licenses and telephone numbers are allowed to select vanity license plates or particular telephone numbers, too. Because the domain name

space is so vast, however, it can accommodate consumer selection of their preferred identifier much more readily than other identifier spaces. Consumer selection is the rule rather than the exception.

One of the key problems facing this aspect of demand is the existence of different language scripts. (Hotta 2001; Katoh 2001) For a very large and growing portion of the Internet's user base,⁶ ASCII-based domain names are hardly meaningful and memorable. This has led to demand for domain names based on the more complex Unicode standard, which can accommodate most other scripts. Experimental implementations of multilingual domain names, involving proprietary browser plug-ins, are operational now but they are not globally compatible. Prodded by market and political pressures, the Internet Engineering Task Force has reluctantly made an effort to figure out how to define a standard for internationalized domain names while retaining global compatibility.⁷ Adding scripts involves such basic changes to the DNS protocol, however, that severe compatibility issues arise in any transition. Internationalizing domain names involves a long-term standards migration. On the other hand, making available domain names in non-ASCII characters would greatly expand the demand for registrations, by expanding the population to whom memorable identifiers were available, and by increasing the usability of the names. The localization of domain names would also, however, create many new opportunities for conflicts over name rights. (See section 5.2 below)

Because the human capacity for memory is limited, and because time is scarce, shorter domain names tend to be more valuable than longer ones, all else being equal. Despite the attempt of some trademark lobby groups and trademark holders involved in domain name litigation to portray domain names as both powerful keywords and the legal equivalent of source identifiers, most companies utilize truncated, easier-to-remember and type-in domain names their primary online identity. Merrill Lynch may have registered <merrilllynch.com> for defensive purposes, but the domain name it really uses in its publicity materials is <ml.com>. Likewise, the Ford Foundation relies on <fordfound.org>, and has not even registered <fordfoundation.org>.

⁶ It is estimated that by 2003, two-thirds of all Internet users will be non-English speakers. Over 90 per cent of the world's population speaks a primary language other than English.

⁷ For the charter and schedule of the "idn" (internationalized domain name) working group of the IETF, see <http://www.ietf.org/html.charters/idn-charter.html>.

In this calculus of memorable identifiers, the supply of clusters of names meaningful to and desired by any individual user expands linearly with the addition of top-level domains. However, the meaning and hence the value of second-level domains is dependent upon the meaning or qualities associated with the top-level domain. E.g., the Disney Corporation is unlikely to be interested in any names under a <.sex> TLD whereas a porn site operator would have both legal and business reasons to avoid registering under a <.kids> TLD.

3.2.2 Names as authenticators

As a complement to and extension of memorability, names can aid in the performance of authentication functions. That is, in addition to establishing a mnemonic linkage between the character string and its owner, it can also suggest or guarantee that the owner is an entity of a certain type or status. The <.edu> top level domain, for example, is (mostly) restricted to organizations that are accredited four-year degree-granting institutions of higher education in the United States.⁸ Thus, if one sees a domain name with <.edu> at the end, and *if* one is familiar with the policies of the <.edu> registry, one knows more about the registrant than one would otherwise. An example of a more modest form of authentication would be a country-code top-level domain that requires that the registrant be domiciled in the country referenced by the TLD. Authentication is always a matter of degree, and the level of trust or credibility that can be provided by domain name registration authorities is necessarily limited.

Linking domain name registrations to verification or authentication functions makes domain name registration much more expensive and the number of registrants smaller, which is why restricted domains are not as common as one might expect. Registries that perform the additional functions are required not only to reserve a unique string and provide name service, but must also engage in verification and enforcement functions that require human judgment, manual processing of registrations, and enforcement procedures or dispute resolution mechanisms. The problem is complicated even more by political and regulatory issues. One could, for example, imagine a <.xxx> domain for pornographic content or a <.kids> TLD for children's content. But who will define and apply the criteria for deciding what is "pornographic" and what is "suitable for kids?" Of course, private sector firms have a long history of providing ratings of

⁸ The standards for <.edu> registration were just expanded to include Community Colleges.

video games and movies on a self-regulatory basis. Such signals lower the search costs of consumers by narrowing down the choices of consumers into appropriate ranges and converging the expectations of suppliers and consumers. However, in mass media industries these self-rating systems achieve uniformity of application by emerging from collective action among dominant industry players, often incited by the threat of government intervention. Internet domain names are global, not national, in scope, and the range of activities embraced by them are much more heterogeneous than movie or video game ratings. In a free marketplace, domains can acquire a reputation just as a brand name or a movie rating category can. But they can also lead to political friction over divergent expectations or attempts by politicians to impose the classifications on unwilling parties.

Domain names as authenticators overlap with efforts to make top-level or second-level domains some kind of taxonomy or categorization of registrants. Country codes are the most obvious example of a categorization scheme imposed on the domain name space. As an indication of the limitations of this approach, however, it should be noted that many country code TLDs do not limit registrations to businesses or individuals located in that country. Many country codes restrict second-level domains to categories such as <ac> or <edu> for academic institutions, <co> or <com> for commercial organizations, etc., and do not allow individual organizations or people to register in the second level. Most of the sorting of registrants into these categories is performed by the registrants themselves when they register, although some registries issue policies and may enforce them if deviations are brought to their attention.

The problem with taxonomic approaches to the domain name space is that the categories of interest are manifold and change constantly over time. No classification scheme will ever be exhaustive. New categories will always arise that will seem more meaningful or important than the old ones. For example, there is great interest now in a <.blog> TLD because of the proliferation of a new genre of communication known as “web logs.” But the “blog” category simply did not exist three years ago. Due to the switching costs discussed earlier, domain names cannot be quickly or easily re-sorted into new categories as things change. Thus, it may be unwise to attempt to link identifiers, for which persistence and memorability are key virtues, with content-based classifications, which might shift over time or be the site of competing approaches. Unless a central authority has complete control over who registers in which domain – something that seems undesirable to most Internet users – the mapping of domain name

registrants to specific categories will be imperfect, particularly given the large number of legacy registrations in existing domains. Moreover, it is unlikely that any categorization scheme will be truly global in effect – i.e., any scheme that works at the psychological and user behavior levels in sub-Saharan Africa or Indonesia will probably not work well in Scandinavia or the United States.

3.2.3 Demand for Search and Directory Functions

In addition to the memorable identifier function, part of the demand for domain names is derived from their ability to perform search functions. Domain names can perform search functions when they are “guessable” and can be used to locate web sites (or other online services) for the first time. In the paradigmatic case of using a domain name as a search aid, an Internet user who wants to find AT&T’s web site may guess that the address will be <att.com> and type that into the URL bar of her web browser. (Note that even this seemingly simple guess requires knowledge on the part of the user that the “&” character is not allowed in DNS.) In effect, the second-level domain name is used as a keyword. Early browsers reinforced this practice by automatically making <.com> the default TLD whenever a word was typed into the URL window. Obviously, search engines and portals are substitutes for these aspects of domain naming and improvements in the quality of Internet searching services has begun to greatly diminish the demand for domain names as search keywords.

Reliance on guessing was a response to unique and temporary conditions that emerged from 1994 to 1997 when the World Wide Web (WWW) was new.⁹ At that time the WWW was overwhelmingly American and Internet content was overwhelmingly English. The number of registered domain names was 1% of the number registered now. Moreover, during the initial rush of WWW growth, domain name registrations were concentrated in the <.com> top level domain (TLD) to a degree that was excessive and historically unique. This concentration happened because .com was the only commercially oriented top-level domain that accepted registrations from anyone at a low price. Most of the other TLDs were restricted or much more expensive at that time. At the end of 1996, more than 88 percent of all domain names in the USA, and just over 75 percent of all domain name registrations in the world, were under <.com>. The state

⁹ The WWW protocol was developed at a Swiss physics research institute (CERN) in 1990-92. The first free, easy to use Web browsers only appeared in 1993; Netscape appeared in 1994.

of search engine technology and service was much less developed than it is now. Under these historically unique conditions, the search strategy of “guessing” a domain name was reasonably effective (although there are no generalizable empirical studies on how often users did this).

Virtually every commentator who has remarked on the use of domain names for search and directory functions has agreed that domain names make lousy directories or search tools. Domain names require a precise match between the search term and the domain name to succeed; most searches involve guesses that are highly unlikely to be exact matches. Domain names must be unique, but there can be multiple legitimate users of the same names or search terms; e.g., “Apple” can refer to a music recording company, a fruit, a computer company. The sheer number of second-level domains makes guessing increasingly less viable. Instead of the 1 million or so registered domain names in 1997, there are now nearly 50 million registered domain names. As the number grows, the viability of guessing declines. Instead of 75% of all registered domains being concentrated in the <.com> top level domain, now less than half of them are. There are over 2 million registrations in the new generic top-level domains <.info>, <.biz>, and <.us>. These domains are growing more rapidly than <.com>. Expansion of the number of generic top-level names and distribution of more registrations into the country code TLDs undermines the ability of a second-level name to function as a keyword. As this happens, the importance of the search function of domain names declines in relative importance to the memorable identifier function. As the domain name universe and the Internet grow, guessing domain names may seem as benighted as attempting to guess phone numbers. Higher-level tools, such as search engines, act as direct substitutes for domain name-based guessing or trial and error. Here again, however, there is a paucity of empirical research. The cross-elasticity of demand for domain names and search engines has never been measured or even studied systematically.

Nevertheless, the function of domain names as search tools has had and continues to have a profound effect on the political economy of domain names. The prospect of capturing web traffic stimulated speculative registration of names, including some bad faith or deceptive registrations of well-known trademarked names. It also encouraged companies and brand owners to engage in systematic defensive registrations. On the assumption that large numbers of users might expect to find them at those domains, a company might try to pre-emptively register all of the names in the semantic cluster associated

with a particular concept or source identifier. It might also attempt defensive registration of the same names across multiple TLDs. Obviously this served as a major stimulus to the domain name registration market within the established TLDs. But it also led to efforts to regulate the market and restrict the supply of new TLDs.

4. Domain Name Supply

The supply side of the domain name industry consists of three basic elements: the root server operators, registries, and registrars.

4.1 Root servers

Root servers are the name servers for top-level domains. ICANN, with the US Department of Commerce looking over its shoulder, is the registry for top-level domains. The root servers' technical, political and economic distinctiveness is derived from the importance of the starting point, the root, in the name space tree. No TLD can operate unless the information needed to resolve it is entered into the root zone file and advertised to the public Internet via the root servers. Thus, the root zone file is an important point of policy control in the DNS, because it controls supply and market entry in the DNS name space. Adding TLDs (and authorizing registries to operate them) expands the supply of names and (possibly) the number of competitors in the domain name registration market.¹⁰

At the time of this writing, there is no centralized and uniform method of operating the root servers. The information contained in the root zone file is controlled by ICANN's decision making process and ultimately by the U.S. Department of Commerce. But the actual operation of the root servers is performed by a heterogeneous group of entities from the Internet technical community who inherited operational control of specific root servers prior to the Internet's commercialization. One of the root servers is operated by ICANN itself. Two are operated by VeriSign. ICANN and VeriSign are both directly or indirectly beholden to the U.S. government. Other root servers are in the hands of universities, federal military or scientific agencies, Internet technical veterans or ISPs. Most do not have formal contracts with ICANN or the US Government. All but three of the root servers are in the United States.

¹⁰ Note that new names can be assigned to existing registry operators. In that case, a new TLD would not expand the number of registry competitors, but it would expand the choice of names available to consumers.

(Table 2)

So far ICANN has failed to execute contracts with any external root server operators. The operators' compliance with policy changes ordered by ICANN is essentially voluntary, although compliance is strongly influenced by the need for a coordinated, unified name space. Any attempt by one root server operator (or a sub-coalition of them) to enter their own TLD names unilaterally or to refuse to accept changes proposed by the others would risk fragmenting Internet connectivity. There are strong professional norms among the involved communities against such a course of action, as well as the likelihood of more direct pressures from industry and government.

4.1.1 DNS Root as Commons

The DNS root is a commons. Anyone who sets up standard-implementation DNS can query the root without being charged. With the exception of ICANN, which taxes the registries and the registrars it accredits, none of the root server operators receive any compensation from the users of the DNS service. Financial support for the root servers comes from the institutions that operate them. The prospect of charging for DNS queries does not seem to be feasible, although this is a good topic for new economic research.¹¹

DNS root servers share all the familiar advantages and disadvantages of a commons. The absence of charging minimizes the transaction costs of interacting with the DNS root. The programs and networking structures associated with the Internet can (at least for now) take its availability for granted. But there is no incentive for Internet users to ration or control their usage of the root. A minor glitch in widely distributed Microsoft software generated millions of queries to the root that could not be resolved, wasting root server capacity. (Brownlee, Claffy et al. 2001) As the size of the Internet grows and new DNS applications such

¹¹ Technical experts insist that end users lack basic forms of control over when, where and how their resolvers issue DNS queries, making any charge-per-query scheme unfeasible. Users might be subject to massive manipulation and abuse, as operators who charged for queries could disseminate programs or links that generated large numbers of queries, or that provoked queries by the machines of uninvolved third parties. Aside from those problems, the value of any individual transaction might exceed the transaction costs of defining a pricing, exchange, and collection mechanism. However, these problems have been solved at lower levels of the hierarchy. Rather than saying it is "impossible" to charge for root service it would be more accurate to say that the DNS industry and many forms of Internet operation have been organized around the assumption that the root is a commons, and altering that would involve sweeping adjustments.

as ENUM are created, the query load on the root servers will increase. Because the number of coordinated root servers is fixed at 13 it is conceivable that growth in query load can outrun the technical capabilities of the root servers at some point. But this too is an area where new modeling and empirical research is needed.

4.1.2 Root Administration as Service to TLD Registries

While for end users the DNS root is a commons, the central coordinator of the root zone file must also provide critical services to registries. The root zone file must contain correct, up to date information about the IP addresses of the top-level domain name servers. When that information changes, the TLD registries must be able to request changes from the root authority and have them executed promptly. These exchanges of information must be highly secure to make sure that malicious redirection of DNS queries, capable of disabling entire top-level domains, does not take place.

At the lower levels of the hierarchy, e.g., in the domain name registrant – registry relationship, these record updates are contract-based services provided for a fee. Currently, however, ICANN does not treat this relationship as a service to paying customers. Instead, it seeks a governmental relationship between the TLD registries and itself, wherein the registries sign contracts recognizing and accepting ICANN a policymaking authority with the power to regulate and tax the registries. This conflict of visions over the role of the root administrator has led to considerable friction between ICANN and the country code registries.

4.1.3 Competing DNS Roots?

We have already noted that on the demand side, consumers of DNS services have strong incentives to converge on a single root. On the supply side, however, there are very few economic barriers to the creation of a competing root authority or root server operators. The root servers are just name servers that happen to be high-capacity because of the reliance of the global Internet on them. Any other network of 13 high-capacity name servers could perform the same function. The level of investment required to replicate all 13 would be in the tens of millions of US dollars.

There is one salient barrier to entry aside from demand-side network effects. BIND software contains a “hints” file that contains a list of the IP addresses of the 13 root servers. When a name server is

started, it needs to establish contact with a root server so that it knows it can send queries there. The hints file tells the ordinary name server operator where to look for root servers. It is possible to alter this file manually. Any attempt to mount a coordinated shift to a new DNS root, therefore, would have to convince those hundreds of thousands of name servers to manually reconfigure their software, or convince the stewards of BIND software to include the new root server addresses, or both. The hints files embeds the IP addresses of the dominant root in all BIND implementations, making the dominant root the default value and thus giving it the advantage of inertia as well as network externalities.

4.2 Registries and Registrars

Below the root level of the domain name hierarchy, registration and name resolution services are supplied by firms that operate on a subscription fee basis. Typically, domain names are sold in one or two-year subscriptions although as the market has become more competitive longer term contracts have appeared. Most of the market is supplied by firms that are commercial for-profits. Some registries are operated by non-profit consortia of Internet service providers or other members of the local Internet community. Still others (usually in developing countries) are run by government ministries or state-owned entities.

4.2.1 Registries

In essence, registries operate public databases that assign names under TLDs and provide, in real time, the name resolution data needed to use the names for communication over the Internet. Domain name registries maintain records, known as *zone files*, of second-level domain name registrations under their top-level domain. They record which second-level domains have been occupied and who has occupied them, and provide name service for those registrations. In addition, they maintain and may offer publicly a “Whois” service, which allows Internet users or private parties to type in a domain name and see the name and contact information of the person who registered it. Registries may also perform the functions of accepting customer orders for specific names, maintaining customer accounts, billing customers, accepting changes from customers, notifying them of expiration, and so on. In the major generic TLDs regulated by

ICANN, these “retail” functions must be separated from the “wholesale” functions of maintaining the zone files. (See section 4.2.2 below.)

Each TLD is by definition a unique string. Under the current technical implementation of the DNS protocol, there can only be one <.com> and only one registry operator for <.com>. This raises the question whether registries compete with each other or are “natural monopolies.” There is a need for greater empirical research on the cross-elasticity of demand across different TLDs. However, experience with both new TLDs and country code registrations indicate that Internet users contemplating a *new* registration do view registries as competitive alternatives to each other. There is also anecdotal evidence of gradual shifts of web sites from one TLD to another.¹² Registries who wish to compete with <.com> can offer TLD names that either convey the same basic semantic meaning as <.com>, such as <.biz>, or that target specific subsets of the commercial world, such as <.shop> or <.law>. Registry competition for new registrations thus involves a form of product differentiation; competition is imperfect but there is sufficient cross-elasticity of demand to make a registration in one TLD a competitive alternative to registration in another. Indeed, competition across registries may exist even when the semantic similarity between TLDs is minimal. Although no formal study has been done, observation of country code registration data suggests that countries that impose numerous restrictions on registering in their ccTLD have a much higher portion of organizations who register <.com> names. In Japan and France, country code registries known for their highly restrictive policies toward registration, more firms are registered in <.com> and other gTLDs than in the country code TLDs. Countries with more accommodating policies, such as <.de> and <.uk>, have greatly outstripped the French and Japanese TLDs in number of registrations and have fewer companies per country registered in <.com>. Thus, it is clear that many Japanese and French registrants view gTLDs and specifically <.com> as an alternative to their country code, despite the absence of any semantic relationship.

Generic TLD registries, such as <.com>, <.net>, and <.info> are regulated via “contracts” with ICANN. As there is no alternative to the ICANN-managed DNS root, it would be more accurate to describe this as a *licensing* regime similar to the licensing of broadcast stations by the U.S. Federal Communications

¹² E.g., the popular “State of the Domain” web site operated by Snapnames, which supplies domain name industry data, has shifted from <sotd.com> to <sotd.info>, and the former domain no longer resolves.

Commission. When a registry is assigned a TLD, it signs a lengthy contract with ICANN that contains, among other things:

- A cap on the wholesale price that can be charged per name-year
- Commitments to abide by ICANN-developed regulatory policies, especially those regarding dispute resolution, Whois availability, and other regulations favored by intellectual property interests
- Detailed regulations regarding its technical and commercial relationship to registrars (see section on Vertical Separation of Registry and Registrar, below)
- Commitments to reserve or remove specific names from the market
- Prohibitions on certain kinds of conduct; e.g., interconnecting with alternate roots, introducing new services without ICANN approval
- Rates of taxation to support ICANN

On the other hand, most country-code TLDs have no formal contracts with ICANN, having received their TLD assignments prior to its existence. ICANN and the US lack the authority to compel their participation in the “contractual” regime, given that any threat to eliminate an entire country code from the root would not be credible. While ICANN and national governments have begun to develop a method for re-delegating country code TLDs, this is still an open, and highly sensitive, area of policy.

Although there has been no nonproprietary empirical research on this topic, an intuitive understanding of the DNS protocol and database operation suggests that there may be economies of scale up to a certain point in the operation of a registry. Economies of scale would come as increases in the number of registrations under a TLD consumed unused capacity in the existing server and operations infrastructure. As the peak load of the server infrastructure was exhausted, however, new investments would have to be made. As a TLD becomes widely used worldwide, geographic distribution of the server infrastructure must take place to maintain quality of service; this may lead to diseconomies of scale. TLD name servers face the same 13-server constraint as the root servers. The dominant registry operator VeriSign has developed a 13-server, geographically distributed infrastructure to support <.com>, <.net>, and <.org>. It has leased capacity on that infrastructure to some country code operators and ENUM service

providers, indicating that it is more efficient to use existing capacity more intensively than to build a new one.

There are also likely to be significant economies of scope when new TLD names are added to existing registry infrastructure. Any given TLD name is likely to appeal to only a portion of the marketplace. A diversified bundle of TLD names, appealing to different types of registrants and different usage patterns, is likely to make more intensive usage of the fixed investment. Adding an additional TLD name to the list of names served is a trivial expense; it involves a one-time entry in a list. Just as magazine and music publishers find it efficient to develop a repertoire of titles and artists to distribute their risk, so it is likely that registry operators would benefit from spreading their load and risk over multiple TLD offerings. But this, too, is just speculation as no scholarly empirical research or modeling has been done in this area.

4.2.2 Vertical Separation of “Registry” and “Registrar”

In some cases the wholesale and retail functions of the registry are separated. This may happen voluntarily, e.g. when the registry is organized as a consortium of retail domain name registration businesses or ISPs who want to keep the database operator in a subordinate role and removed from the retail market,¹³ or when registry operators do not want to enter the registrar market.¹⁴ In other cases it is required by regulation; e.g., both the U.S. Department of Commerce’s contracts with VeriSign and ICANN’s contracts with most new generic TLD registries required the implementation of a “shared registration system” (SRS) that allowed multiple competing registrars equal access to the registry database. Moreover, the ICANN regime makes it impossible for customers to deal directly with registries; they must go through a class of intermediaries known as “accredited registrars.” In a shared registry, registrars are retail firms that do not operate the database but can enter registrations into the registry database. Registrars handle the customer relationship (billing, receiving and executing changes in the customer account, notifying the customer of pending expirations, etc.)

¹³ Examples are Nominet UK, the registry for the country code <.uk>, and Afilias, the registry operator for <.info> and <.org>.

¹⁴ Neustar, the registry operating <.biz> and <.us>, does not operate in the registrar market.

Vertical separation is supposed to improve the competitiveness of the market by giving customers a choice of retail service providers and introducing price competition for the right to service accounts. Customers can transfer names from one registrar to another, so that registrars compete actively for customer accounts. Under the ICANN regime, the price of the registry (wholesale) is regulated but the price of the registrars (retail) is not. This form of regulatory mandated competition is quite similar in form to the mandated separation of local and long distance telephone service in the United States following the AT&T divestiture. Both required the creation of a new technical interface to provide “equal access” to what is perceived as a monopolized facility (the registry database). In both cases, prices for a required input are regulated while retail prices are not. In both cases, the vertical separation greatly reduced barriers to entry, allowing hundreds of firms to enter the market and bid retail prices down much closer to the floor set by regulation. But in both cases, the ease of entry and the artificiality of the interface created problems of “slamming,” or unauthorized transfers of the customer’s account. (See section 5.3 on policy) Just as the local-long distance separation was promoted as a remedy for the historical dominance of a single supplier (AT&T), so the registry-registrar split was a response to the historical dominance of the domain name market by Network Solutions, Inc. (NSI). In the late 1990s, NSI (now VeriSign) controlled nearly 90 percent of the global domain name market, and its ability to adequately service millions of <.com> registrants was questioned, with many reports of poor service. The major difference, however, is that while there are major economic and technical barriers to new entry in the local exchange market, there are minimal economic or technical barriers to new entry in the registry market. ICANN’s reluctance to permit new TLDs and new operators to enter the market is the only real barrier. Once the registry market becomes more competitive, the need for compulsory separation of the registry and registrar functions becomes questionable. Customers who prefer the additional choice and control would be able to select registry services that include a variety of competing registrars; others may prefer integration of the functions for simplicity or efficiency.

(Table 4; Table 5)

4.3 The Secondary Market

Semantic demand for domain names led to the development of an increasingly well-organized secondary market; i.e., a market in which domains are registered by one party and warehoused rather than used, and then auctioned or sold for a negotiated price to a third party. The practice of name speculation is clearly identifiable as a product of semantic demand because the additional value of the name over the subscription price is based entirely on the name's meaning. Also, the market has also moved to capitalize on the residual or inertial traffic-generating capabilities of existing domains, leading to battles over the control of expiring domain names. Data from Matthew Zook on the rate of name expiration indicates that in the second half of 2001, about 10% of all registered domain names changed ownership annually, and about 18% expired annually. (The expiration rate would be highly sensitive to industry conditions and cannot be extrapolated.)

(Table 6)

The emergence of a secondary market for domain names was a predictable product of basic economic characteristics of the supply of registry services. Due to semantics, the value of any individual unit within the vast pool of names under a TLD such as <.com> will vary greatly. And yet, registries with their automated procedures of registration were for the most part unwilling and probably unable to engage in accurate price discrimination reflecting variations in value. Thus, domain name registrations were sold in the primary market at a low, uniform price. Any user willing to invest a few thousand dollars could register multiple names and attempt to capitalize on the widespread gap between cost and potential value.¹⁵ Add to this the artificial restriction on the supply of new top-level domains that was in place from 1996 until 2001, which both restricted alternatives and reinforced the search value of domains within <.com>, and the recipe for name speculation was complete. Most of the speculative activity was in the <.com> domain, although copies of important or popular <.com> domain names began to be registered in <.org> and <.net> and even some country codes as speculative activity increased.

It is theoretically possible to capture the gap between cost and value by holding auctions for initial assignments in the primary market. (For a proposal along these lines, see Kahin 1996) One registry, the

¹⁵ When the "hunch" about the name's potential value was more than a guess name speculation could become an abusive practice. For example, a certain class of name speculators would comb business journals for news of mergers and acquisitions, and as soon as the deal was publicly announced they would register the name of the combined company; e.g., "ExxonMobil.com."

<.tv> top-level domain, experimented with auctions in the primary market. It allowed prospective registrants to enter their name of choice and then returned a price that purported to reflect interest in the name. This pricing mechanism did not work because the auctions were not thick, simultaneous, nor organized by a neutral party. Rather than organized auctions, the <.tv> initial assignments resembled more an algorithm for negotiation with the registry, with asymmetric information due to the registry's collection of past bids or expressions of interest. Indeed, the registry operator could (and for all we know, did) "bluff" about the auction value of names, as there was no transparency.

The value of expiring names is another area of the secondary market. Many Internet users would like to have a secure, predictable method of queuing up for expiring domain names. As an example, if a company has just trademarked a new product it wishes to name "Agile" and <agile.com> has been held for two years by a home user who thought it was a cool domain name to have, the company wants to stand first in line when or if the name expires. The company may not want to contact the home user to negotiate a transfer due to the uncertainty and potential for opportunism surrounding the market. The company does not know whether the registrant plans to renew the name or not. If it contacts the registrant and offers to buy it before it expires, the registrant has been tipped off that the name has value, and the registrant might exploit that to renew the domain and set a high sale price. The waiting registrant thus cannot know whether it stands to win or lose by initiating a transaction directly with the registrant.

Prior to mid-2002, the market for expiring domain names was a common pool wherein competing registrars sold domain name recovery services to end users and then stormed the registry with "add" commands when the name expired. The same recovery service could be sold by any competing registrar, making attempts to capture the name a matter of brute technical force. The "add storms" that afflicted the registry infrastructure were a pure example of a tragedy of the commons. The expiration system gave secondary market makers an incentive to inundate the registry with requests for the expiring name regardless of its effect on the registry infrastructure.

Economic analysis would suggest holding an auction for the name upon its expiry. An auction would assure that the name went to whoever valued it the most. The problem is: who should serve as the auctioneer? The prior registrant is not the best choice due to the information and opportunism problems mentioned earlier. Aside from that, efficient auctions require organization and publicity, both of which cost

money. Unless the incumbent registrant is an organized domain name brokerage, she will not be prepared to operate an auction. If the registrant uses the WHOIS record to advertise that the name is for sale, he runs afoul of ICANN's trademark-protecting dispute resolution rules, which make it possible to take away names from people who registered them purely for resale. Probably the most efficient and direct entity to serve as the auctioneer is the registry itself. The registry knows when the name expires and has direct control over when it is released. A centralized auction that incorporates all prospective bidders is more efficient at equilibrating supply and demand than a bunch of smaller, fragmented auctions. However, such an auction would step over the boundaries of the artificial separation of registry and registrar functions, which is a linchpin of the ICANN/US Department of Commerce regime. In other words, registries would be selling names at retail, and making higher margins to boot. A registry auction would create a distribution of wealth favorable to registries rather than registrars, yet registrars are more numerous and politically in favor than registries in ICANN's regime.

In Fall 2002 ICANN permitted VeriSign to enact a mild version of a registry auction known as the Wait Listing Service (WLS).¹⁶ In this service, companies interested in an expiring domain can pay their registrar a market price to gain "first dibs" on a domain name when it expires. The registrar in turn will pay VeriSign registry a \$35 per name wholesale rate (with rebates of \$7 or \$11 depending on volume) to put their customer first in line upon expiration. Almost a year after it was first proposed, ICANN approved the service, but not without adding a regulation that may have completely undercut its business value. ICANN required the registry to inform existing domain name subscribers when their names had been waitlisted (thus encouraging them to renew the name and bargain directly with the interested party).

A great deal of the interest in expiring names comes from Internet services that make their money selling "hits" to advertisers. These companies have recognized that for months and possibly even years after the expiration of a domain name, there may still be users or outdated links that direct traffic to that domain. By gaining control of a once-popular domain name that has expired, these businesses can capture a traffic stream and expose them to advertisements that may result in billable hits.

¹⁶ For a description of the WLS services, see "Domain Name Wait Listing Service," VeriSign Global Registry Services, March 2002. <http://www.icann.org/bucharest/vgrs-wls-proposal-20mar02.pdf>

5. Economic Policy Issues

Building upon the analysis in the previous sections, the framework can now be applied to the analysis of current policy issues. It should be noted that some of the most important policy problems are *institutional* in nature. The DNS industry is relatively young, and the attempt to create a global, private sector “self-regulatory” system around it is even younger. Unlike many policy analyses, therefore, this discussion cannot take for granted an established, stable governance framework that uses known methods to establish policies or known techniques for defining and enforcing regulations. There are also problems and issues regarding the formation of the institutions themselves: how they make decisions, the proper scope of their authority, the incentives to participate in the regime, and so on. Such a discussion falls outside the scope of this paper, however. The next sections survey the economic policy issues in the domain name industry in a way that basically takes the ICANN-based institutional framework for granted. Some indications of the institutional problems that pervade the treatment of many of the policy issues are provided.

5.1 New TLDs: Expanding Supply

ICANN was created to manage the root of the domain name system. That function implies making policy decisions about how many TLDs are added, what pace they can be added at, what criteria will be used to determine who gets the available assignments, and how to resolve competing applications for the same new TLD, and so on.

Policy conflict over adding new TLDs is one of the issues that led to the creation of ICANN in the first place. (NTIA 1998; Mueller 2002, Chapter 6) Calls for name space expansion are not, however, based on a “shortage” of domain names *per se*. As noted earlier, the DNS name space can handle all conceivable technical demand without any change. If users were willing to register and use extremely long names, meaningless names, and names that go five or more levels down into the name space hierarchy, there would be no need for expansion of the supply of TLDs. Likewise, if the ownership and operation of current registry services were perfectly satisfactory to everyone, there would be no need to permit the entry of new registry operators. The debate over new top-level domains is really a debate about the degree to which DNS administration should respond to human factors, user demand, and competition policy concerns.

There are many potential sources of demand for new TLD names. Users may want more choices regarding the identity they project online. For example, a <.blog> top-level domain might attract numerous registrants from the growing web loggers community. It is also possible that new entrants might have valid ideas about how to provide targeted services better than incumbent operators. For example, the <.name> TLD is targeted at personalized domain names, but their business model and policy restrictions are unattractive to many registrants. An alternative TLD for personal names would add competition and choice to the market. Or users may want to move their name up in the DNS hierarchy for ease of use and/or greater control over the management of their name and the name space below it. A corporation to whom online identity is essential, such as AOL or Amazon.com, may decide that it wants to “in-source” its DNS management functions completely by running its own TLD, just as many companies privatized their network management functions in the 1970s. Some groups of organizations may want to establish a controlled name space, analogous to <.edu> for US universities, to promote authenticity of online identity. Thus, adding top-level domains will have a major impact on a) the variety and usability of identifiers, b) competition among registries, and c) the ability of firms or industries to control their digital identity.

Surprisingly, after 4 years of existence ICANN has not defined a timetable or a rule-based method for responding to requests for new top-level domains. Indeed, it has not even set in motion a proceeding to define a method of routinizing this function. Instead, it has approached additions in an ad hoc manner. In year 2000 it held a long and controversial process that added 7 new TLDs. That round received 44 applications from bidders willing to pay a US\$ 50,000 nonrefundable fee to be considered. Before making its 7 awards, mostly to firms with close political and organizational connections to ICANN’s management and Board, ICANN announced that the whole new TLD process was an experiment or “proof of concept,” the results of which would be monitored. No timetable for the beginning or end of this monitoring process was fixed. Late in 2002 ICANN’s CEO suddenly announced that he thought it would be a good idea to add three more restricted TLDs. No one knows where this number came from.

ICANN’s unwillingness to tackle the problem of new TLDs in a systematic manner is caused by three factors: a) it is a self-regulatory system in which incumbent operators have a great deal of influence over policy making, and incumbents have little incentive to define an open entry process that would subject

themselves permanently to the prospect of additional competition;¹⁷ b) brand-name holders, trademark interests and holders of <.com> domains that are important and valuable fear that additional TLDs will only result in the need to defensively register their existing second-level domain names in order to pre-empt their occupation by others, including trademark infringers; c) it emerged from a technical culture in which policy decisions affecting the demand and supply conditions of an industry are often confused with design decisions, which can be centrally planned to create a local optimum. All three pressures act to maintain artificial scarcity in the supply of TLDs.

Earlier in the debate there were attempts to suggest that there are technical obstacles to the addition of new TLDs. The Internet Engineering Task Force has never created a working group to systematically analyze or define any technical constraints on TLD additions or their impact on the operation of the root zone. It is an irrefutable fact, however, that the original root administrator and one of the designers of the DNS, the late Dr. Jon Postel, proposed adding 50 new TLDs a year for three years in a row back in 1996. While this plan was rejected, its failure was not attributable to technical concerns about expansion of the root zone. During the early and mid-1990s, as country code TLDs were being delegated, the root zone was expanding by at least 10 TLDs per year, and for several years in a row more than 20 were added each year. Key figures within the IETF, such as Paul Vixie and Karl Auerbach, have pointed out that the root servers use the same technology as the name servers for the top-level domains. In that respect, the root zone file is no different from any other DNS zone file. As there are millions of functioning registrations in the .com, .net, .org, .de, and .uk zone files and those zones work reliably, there could be millions of top-level domains. However, the scalability of the system depends in large part upon the existence of hierarchy, so almost no one supports having that many TLDs.

There is one important difference about the root zone, however: errors or corrupted files at the root level could have more harmful consequences for Internet users than mistakes that occur lower in the hierarchy. An erroneous root zone file could result in the inaccessibility of entire TLDs until the problem

¹⁷ In an article in Computerwire in October 2002, ICANN CEO Stuart Lynn was reported as saying “the last round of new TLDs, which included .info, .pro and .name, were introduced largely due to market demand, but some people think this is no longer an appropriate reason to introduce new domains.” The “some people” referred to are incumbent registries, who believe that the market is saturated and do not want additional competitors.

was fixed, whereas the effects of a corrupted TLD zone file would be more localized. Thus, there is a valid technical concern about limiting the *rate at which the root zone changes* in order to minimize the risk of errors in the root zone. While some Internet technologists believe that root zone changes could and should be automated, more conservative traditionalists believe that the root zone should continue to be altered by hand and subject to human inspection before being released and propagated to the root servers. Even the adherents of this most conservative view, however, believe that 30–90 additions and changes in the root zone file made in batch mode at a specific periodic rate, such as annually or every six months, are safe. Thus, it is conservative to note that 50 or so new TLDs could be added annually to the root zone. (See for example Hoffman 2002)

The only other technical concern raised by the addition of new TLDs is the degree to which TLD additions increase the query load on the root. Because the number of root servers is fixed at 13 and there are technical limits on the capacity of each root server, there is some cause for concern about increasing root server load. However, because the DNS relies so heavily on caching at lower levels of the name server hierarchy, there is no simple, linear relationship between adding TLDs and increasing root server load. The consensus position within IETF seems to be merely that the number of TLDs should be finite rather than infinite, such that the hierarchical character of name resolution and assignment is maintained. As IETF Chair Fred Baker put it,

If we can add one TLD (and we obviously can), we can add 1000 TLDs to the [root zone] table. How that relates to [root-server] lookups for those TLDs is indeterminate from the fact that they are there. How many lookups, for example, do <.tv> or <.info> get? The fact that we added seven TLDs does not mean that we have even *changed* the root server load, much less multiplied it by something. How much additional load we would get is a *business* question: how many new computers, with people using them (and what would they be doing, and what sites would they therefore be accessing), will be added to the Internet because this TLD is in operation?¹⁸

¹⁸ Email to author, October 11, 2002

It is widely agreed that ICANN's initial addition of 7 new TLDs and the re-vitalization of the already existing <.us> TLD have led to no discernable change in root load or root server behavior. In sum, TLD additions have little direct bearing on what is a more fundamental question, which is how the DNS is able to scale with the growth of the Internet. The growth of the Internet, and new DNS applications such as ENUM, have more relevance to root server load than TLD additions *per se*.

Expansion of the number of TLDs is really a policy issue rather than a technical one. ICANN's unwillingness to define a stable process for adding TLDs, and its reliance on sporadic "beauty contests" (i.e., merit assignment procedures) to assign new domains are produced in part by amateurism on the part of ICANN management and in part by political pressures.

5.2 Domain Name – Trademark Conflicts

When the WWW accidentally made domain names into global public identifiers, it fomented property rights conflicts over who had the right to specific name assignments. The most important form of conflict occurred over trademark rights. National laws and international treaties recognize various forms of exclusivity in names, based on registration and (sometimes) use of the name in commerce as a source identifier. It is possible to register and use domain names in a way that constitutes trademark infringement or passing off. A domain name corresponding to a trademark, registered and used by an unauthorized person, may lead a user to a website that deceives or confuses her into thinking that she is dealing with the trademark owner. E.g., if someone other than America Online registers www.aim5.com and displays information and logos related to AOL Instant Messenger (AIM) Version 5 in order to attract advertising revenue or sell goods, classic concepts of trademark infringement apply.

There were two problems with applying trademark law to domain name registration, however. First, the costs of initiating a problem were absurdly low, while the costs of prosecution were high. Domain names in most registries were assigned on a first come, first-served basis at low annual fees. The economics of domain name supply did not support pre-checking of applications; indeed, most registration processes were automated after 1995. Thus, a potentially abusive registrant could acquire a domain name corresponding to a trademark in a matter of hours for less than US\$100. The cost of initiating a legal dispute, however, started in the tens of thousands of dollars. Second, the scope of a domain name was

global, but trademark jurisdiction is mostly national. Thus, the transaction costs of enforcing property rights in the domain name space were high. High transaction costs often harmed innocent registrants as well as trademark owners. Many legitimate registrants of generic terms or names that happened to be trademarked by someone (e.g., <prince.com> or <clue.com>) found themselves assaulted by lawsuits that required hundreds of thousands of dollars to defend. This threat, which came to be known as “reverse domain name hijacking” was cynically employed by some trademark counsel in order to acquire valuable DNS real estate at below market cost.

Domain name – trademark conflicts led to an institutional innovation: the creation of a new global system of dispute resolution over name assignments. (Froomkin 2002) ICANN’s Uniform Domain Name Dispute Resolution Policy (UDRP) created a mandatory system of arbitration wherein multiple dispute resolution service providers compete for the business of complainants (i.e., trademark holders initiating a dispute). The policy emerged from a World Intellectual Property Organization (WIPO) report and ICANN’s domain name supporting organization. (WIPO 1999) The UDRP dramatically lowers the cost of initiating a dispute and is often (though not always) faster than traditional courts. As basic economics would suggest, lowering the cost of dispute resolution led to huge increases in the quantity supplied. Whereas domain name -related court cases number in the hundreds at best, as of November 2002 approximately 7,500 proceedings involving 12,000 domain names have been initiated under the UDRP.

The UDRP now functions as a kind of global law in name -rights, at least within the domain name space. While the stated goal of the new system was to translate *existing rights* recognized by law into a new arena, the move from territorial jurisdiction based on national court systems to global jurisdiction based on contracts of adhesion with registrars has necessarily led to substantive changes in the nature of name rights.

The new regime has strengthened the exclusivity of trademark owners and sometimes coined entirely new rights to names. Research has shown that the right of complainants to select the dispute resolution service provider leads to forum-shopping which rewards trademark-friendly providers and weeds out any providers who interpret the policy in a pro-defendant way. (Geist 2001; Mueller 2001) Research also confirms that the UDRP has created new rights in names, notably in the area of rights of personality. (Mueller 2002b) WIPO and various national governments have been eager to use the new regime to create

expanded rights to names in various areas, such as geographical indicators, country names, nonproprietary pharmaceutical names, and names of international governmental organizations. (WIPO 2001)

5.3 Competition Policy

Some of the most difficult and technical policy issues in the domain name industry involve competition policy. Concerns about fostering competition and ameliorating market dominance have shaped the U.S. government's approach to the problem since 1997. At that time, the primary concern was the monopoly on commercial generic TLDs by the U.S. government contractor Network Solutions, Inc. (NSI). (NSI was later acquired by VeriSign and henceforth will be referred to as VeriSign. VeriSign's dominance was an unanticipated result of a decision by an educational/scientific agency, the National Science Foundation, which authorized it to begin charging for domain name registrations in 1995 with no thought for competitive alternatives.

The U.S. government could have responded to VeriSign's dominance of the market in two distinct ways. One would have been to authorize the entry into the market of new registry operators with exclusive control of new TLDs. The other would have been to regulate it as a dominant provider and impose upon it the vertical separation of the registry and registrar functions, leaving VeriSign's dominance of the registry market in place but creating managed competition in the retail market for registrar services. For better or worse, the US chose the latter path. It failed to authorize new TLDs (and hence stifled nascent registry-level competition) for more than five years of the Internet's most rapid commercialization and growth. Instead, the USG opened up the dominant <.com> domain to registrar competition, thereby making <.com> names even more popular and thus prolonging and supporting VeriSign's dominance of the registry market. Due to the stickiness of domain name registrations, the effects of that delay of competitive entry cannot be reversed.

Registry – registrar relationships are now a sensitive area for regulation that affects both consumer protection and competition policy. The separation of registry and registrar was supposed to make domain names easily portable across registrars and thereby make the market more competitive. However, ease of movement of the account across registrars also makes it relatively easy for a domain name version of “slamming” to take place. Unethical registrars can transfer domain name registrations without the

authorization, knowledge or consent of the registrant. The threat of slamming in turn makes it possible for dominant registrars that are losing market share to impose more restrictive practices on requests for transfers of accounts. As a result, ICANN (or some other regulator) must play a significant role in defining procedures that consumers can rely on when they want to transfer domain names from one registrar to another.¹⁹

Interesting competition policy questions also arise in relation to the secondary market. Several major registrars have through acquisitions gained control over major domain name resale brokerages. By forward integration into the secondary market, a major registrar can undermine the registry-registrar split, because expiring names might be retained by the registrar and sold on the secondary market for an auction price instead of being released into the common pool where any registrar has an equal chance to sell it. VeriSign as noted earlier promoted a “wait list” service as an alternative to the “first come, first served/brute force” method of grabbing expiring names. This, too would tend to favor a supplier with a larger market share because it would allow a supplier with a large but declining market share to generate additional revenue from registrations they already controlled. However, these alleged distortions in the market are nowhere near as significant as the long-term restriction on competitive entry into the registry market caused by the de facto moratorium on new mass market-oriented TLDs. Indeed, one could argue that with a more competitive registry market regulation of the registry – registrar interface would be entirely unnecessary.

As part of its agreement to allow VeriSign to continue in control of the <.com> domain in 1999, the U.S. Commerce Department required VeriSign to divest itself of its registrar subsidiary after a certain period of time. VeriSign’s market share in the registrar segment had declined so rapidly, however, that when the period expired VeriSign and ICANN entered into a new bargain, under which VeriSign was permitted to keep its registry business but was required to divest itself of the <.org> TLD and possibly also the <.net> registry. In divesting VeriSign of <.org> ICANN’s management decided at the outset that it would be safest not to award it to any new entrant, but only to applicants that already ran a registry with

¹⁹ See the report of the Names Council Task Force on Transfers, “Policies and Processes for Gaining and Losing Registrars,” which states in its executive summary that “actions by 2-3 of the largest Registrars means that choice is simply ‘on hold’. Despite several attempts at forging voluntary agreement on the necessary changes, the record shows that there are still problems with the “portability” of domain registrations – customer choice remains limited.” <http://www.byte.org/nc-transfers/final/>.

500,000 or more registrants. That policy decision effectively restricted the field to two incumbents, Afilias, the registry consortium that it had already awarded the <.info> TLD in the first round of expansion, and Neustar, the registry operator that it awarded the <.biz> TLD in the first round of expansion. In 2002 ICANN awarded the <.org> registry to a partnership of the Internet Society and Afilias. The registry market remains highly concentrated. (Table 4)

5.4 WHOIS and Privacy Policy

Associated with the DNS protocol is the WHOIS service, a protocol that allows one to look up the name and contact information of the registrant of a domain name. The WHOIS service was developed in the early days of the Internet when contact information associated with a domain might be needed to resolve technical problems.

As domain names became economically valuable after 1995, WHOIS also became a popular way to find out which domain names were taken, who had registered them, and the creation and expiration date of the registration. Economic value also brought with it domain name – trademark conflicts. Trademark holders discovered that they could perform searches for character strings that matched trademarks, and pull up many of the domain name registrations in the generic top-level domains that matched or contained a trademark. This automated searching function proved to be so valuable in instigating litigation and/or UDRP complaints that the trademark interests began to demand that the WHOIS functions be institutionalized, expanded, and subsidized. The first WIPO Domain Name process recommended that the contact details in a WHOIS record be contractually required to be complete, accurate and up to date, on penalty of forfeiture of the domain name. (WIPO 1999, para. 73) The intellectual property interests also demanded “bulk access” to the WHOIS data of domain name registrars; i.e., the right to purchase the complete list and contact data for all of a registrar’s customers in one fell swoop.

When ICANN was created its power to accredit registrars gave it the ability to implement the transformation of WHOIS demanded by the intellectual property interests. The registrar contract was drafted in accordance with WIPO’s policy suggestions. By signing the accreditation contract, registrars

commit themselves to the provision of a public WHOIS service for free, and to the supply of bulk access to their WHOIS records for a regulated price.²⁰

Intellectual property holders now want WHOIS functionality to be expanded so that data can be searchable by domain name, the registrants' name or postal address, technical or administrative contact name, NIC handles,²¹ and Internet Protocol addresses. They also want searches to be based on Boolean operators or incomplete matches, as well as exact string matches. Further, they are requesting that the results of searches not be limited to a certain number (VeriSign only returned 50 records at a time). Moreover, they want this expanded capability to be subsidized; i.e., they want it to be considered a part of the public Internet infrastructure and not a value-added service that they would have to pay for. Not content with the already massive reduction in transaction costs brought about by the existence of a single, integrated name space that can be searched using automated tools, they want to shift the costs of policing and monitoring the trademark-domain name interface onto users, registries, and registrars.

The issue is no longer exclusively one of trademark surveillance and protection. Copyright interests now view expanded WHOIS functionality as a way to identify and serve process upon the owners of allegedly infringing web sites. Moreover, public law enforcement agencies, notably the U.S. Federal Bureau of Investigation (FBI), have become deeply interested in the use of WHOIS to supplement their law enforcement activities. The intent is to make a domain name the cyberspace equivalent of a driver's license. Unlike the drivers' license database, however, this one would be publicly accessible to anyone and everyone.

The growth and commercialization of the domain name system, and political pressure from intellectual property interests, continue to have a major impact on the evolution of WHOIS. The technical challenges involved in providing an integrated lookup system across the gTLD name space(s) have been magnified by the introduction of competing registrars who share access to the same registry, and by the addition of new top-level domains. For the first year or so after the introduction of registry sharing, WHOIS service was fragmented across registrars; i.e., one had to know which registrar a customer used to register a

²⁰ See Section 3.3, ICANN Registrar Accreditation Agreement.

²¹ The NIC handle is a short, unique alphanumeric code that a registry assigns to a domain name holder when the registrant registers a name. People who use different names might use the same NIC handle in the WHOIS record.

name in order to be able to look up the name. Some limits have been placed on technically abusive data-mining techniques that have been directed at the WHOIS databases of registries and registrars.²²

In response to pressure from major corporate trademark and intellectual property holders, the US Commerce Department is pushing for universalizing and globalizing the level of WHOIS service formerly associated with the VeriSign monopoly. Efforts are underway to create a “universal WHOIS” that would provide access to registrant information within country code TLDs as well as the generic top-level domains. The revised Verisign-Commerce Department contract of May 2001 requires the company to continue “development and deployment of a universal Whois service that allows public access and effective use of Whois across all registries and all top level domains at the direction of the Department.”

6. Conclusion

A rich series of economic and policy problems have been raised by the emergence of a domain name market. Economists could profitably devote more attention to issues of naming and identity and the associated industries of registration and name service. There is a need both for modeling and for empirical research. ICANN’s loosely organized and amateur policy development apparatus is badly in need of professional research into policy options and the consequences of prior policies.

²² See “Investigation Report by Verisign Global Registry Services,” In the matter of Register.com, Inc. v. Verio Inc., 00-Civ-5747 (BSJ) <http://www.icann.org/registrars/register.com-verio/registry-report-30jan01.htm>

Table 1: Growth of <.com>, <.net>, and <.org> domain name registrations

Date	# Registered Domains
Jul-96	488,000
Jan-97	828,000
Jul-97	1,301,000
Jan 98	2,292,000
Jul-98	3,282,117
Jan-99	5,504,151
Jul-99	9,098,066
Jan-00	13,402,448
Jul-00	23,864,611
Jan-01	33,045,397
Jul-01	37,539,541
Jan-02	30,225,000
Jul-02	29,255,166

Table 2: Root Servers

Name	Geographic Location	Operating institution
A	USA – East	VeriSign Global Registry Services
B	USA – West	Information Sciences Institute
C	USA – East	Cogent Communications (as ISP)
D	USA – East	University of Maryland
E	USA – West	NASA Ames Research Center
F	USA – West	Internet Software Consortium
G	USA – East	U.S. DOD Network Information Center
H	USA – East	US Army Research Laboratory
I	Stockholm, Sweden	Autonomica
J	USA – East	VeriSign Global Registry Services
K	London, Great Britain	Réseaux IP Européens (RIPE) Network Coordination Centre
L	USA – West	ICANN
M	Tokyo, Japan	WIDE Project, Keio University

Table 3: Compatibility Relations among Competing DNS Roots

3a - Competing Root (Root-C) Supports Established TLDs of Incumbent Root (Root-I)

	<i>Origin of domain name query</i>	
<i>Origin of domain name assignment</i>	Users of Root-I	Users of Root-C
Root-I	Compatible	Compatible
Root-C	Incompatible	Compatible

3b - Competing Root (Root-C) is Uncoordinated with Incumbent Root (Root-I)

	<i>Origin of domain name query</i>	
<i>Origin of domain name assignment</i>	Users of Root-I	Users of Root-C
Root-I	Compatible	Incompatible
Root-C	Incompatible	Compatible

Table 4 Registry Market Share

Name	Share	Registrations (millions)
Verisign (.com, .net)	53.78%	25.04
DENIC (Germany)	11.73%	5.46
Afilias (.info, .org)	7.10%	3.32
Nominet (.uk)	7.10%	3.30
Neustar (.biz, .us)	2.40%	1.11
Total for top five	82.11%	38.23
World total		46.55

Table 5 Registrar Market Share

Name	Rank		Market Share		Number of Registrations	
	Q4 2000	Q3 2002	Q4 2000	Q3 2002	Q4 2000	Q3 2002
VeriSign Registrar	1	1	53.0 %	30.1 %	14,474,754	8,795,276
Tucows (OpenSRS)	3	2	7.4 %	10.4 %	2,011,880	3,044,787
Register.com	2	3	12.3 %	9.8 %	3,379,237	2,863,004
MelbourneIT	5	4	3.5 %	5.6 %	969,423	1,619,257
BulkRegister	4	5	6.6 %	4.7 %	1,812,582	1,372,454
CoreNIC	6	12	3.5 %	2.0 %	967,185	1,290,309
All Others	--	--	13.5 %	37.5 %	3,748,857	10,206,005

Table 6: Data Relevant to the Secondary Market

CHANGES IN DOMAIN STATUS	Three Months	Monthly	Yearly
Changes in Ownership			
No Change in Ownership	93.0%	-	-
New Owner	2.5%	0.8%	9.8%
Expired Domains	4.6%	1.5%	18.3%
	100.0%	2.3%	28.1%
Changes in Registrar			
No Change in Registrar	92.7%	-	-
New Registrar	2.7%	0.9%	10.7%
Expired Domains	4.6%	1.5%	18.3%
	100.0%	2.4%	29.0%
Ownership & Registrar Changes			
Same Owner, Same Registrar	91.6%	-	-
Same Owner, New Registrar	1.4%	0.5%	5.5%
New Owner, Same Registrar	1.1%	0.4%	4.6%
New Owner, New Registrar	1.3%	0.4%	5.2%
Expired Domains	4.6%	1.5%	18.3%

Source: Matthew Zook, Zooknic Internet Intelligence. Based on a sample of 50,000 randomly selected names with observations in June and September 2001.

References

- Albitz, P. and C. Liu (1992). DNS and BIND in a nutshell. Sebastopol, CA, O'Reilly & Associates.
- Andeen, A. and J. L. King (1997). Addressing and the Future of Communications Competition: Lessons from Telephony and the Internet. Coordinating the Internet. B. K. a. J. H. Keller. Cambridge, MA, MIT Press: 491.
- Bechtold, S. (2002). Governance in Namespaces. TPRC 2002 The 30th Research Conference on Information, Communication, and Internet Policy, Alexandria, VA.
- Berners-Lee, T. (1996). The Myth of Names and Addresses, World Wide Web Consortium. **2002**.
- Berners-Lee, T. (1998). Cool URIs Don't Change. **2002**.
- Brownlee, N., k. Claffy, et al. (2001). DNS Measurements at a Root Server. San Diego, California, Cooperative Association for Internet Data Analysis (CAIDA).
- Cannon, R. (2001). ENUM: The Collision of Telephony and DNS Policy. TPRC 29th Annual Research Conference on Information, Communication, and Internet Policy, Alexandria, VA, USA.
- Farrell, J. and P. Klemperer (2001). Coordination and Lock-In: Competition with Switching Costs and Network Effects. **2002**.
- Froomkin, M. (2002). "ICANN's "Uniform Dispute Resolution Policy" -- Causes and (Partial) Cures." Brooklyn Law Review **67**: 605.
- FTC (1998). Comment of the Staffs of the Bureau of Economics and Competition of the Federal Trade Commission on Improvement of Technical Management of Internet Names and Addresses. Washington, DC.
- Gans, J. S., King, S. P. and Woodbridge, G. (2001). "Numbers to the People: Regulation, Ownership and Local Number Portability." Information Economics and Policy **13**(2): 167-180.
- Geist, M. (2001). Fair.com?: An Examination of the Allegations of Systemic Unfairness in the ICANN UDRP. Ottawa, Canada, University of Ottawa Faculty of Law.
- Hoffman, P. (2002). Reforming the Administration of DNS, Paul Hoffman. **2002**.
- Hotta, H. (2001). Technology and Policy Aspects of Multilingual Domain Names. Geneva, International Telecommunication Union: 56.
- Huston, G. (2002). "ENUM - Mapping the E.164 Number Space into the DNS." The Internet Protocol Journal **5**(2): 13-23.
- Hwang, J. and M. Mueller (2002). The Economics of ENUM Over Cable Broadband Access Networks. Syracuse University School of Information Studies, The Convergence Center.
- Kahin, B. (1996). Auctioning Global Namespace: A New Form of Intellectual Property, a New Vision of Universal Service. CIX/Internet Society workshop on Internet Administrative Infrastructure, Washington, DC.
- Kahn, R. and R. Wilensky (1995). A Framework for Distributed Digital Object Services, Corporation for National Research Initiatives. **2002**.
- Katoh, M. (2001). Report of the Internationalized Domain Names Internal Working Group of the ICANN Board of Directors, Internet Corporation for Assigned Names and Numbers.
- McTaggart, C. (2001). E Pluribus ENUM: Unifying International Telecommunications Networks and Governance. 29th Research Conference on Information, Communication, and Internet Policy, Alexandria, VA.
- Mockapetris, P. (1987). Domain Names--Concepts and Facilities, Internet Society. **RFC 1034**.
- Mueller, M. (2001). "Rough Justice: A Statistical Assessment of ICANN's Uniform Dispute Resolution Policy." The Information Society **17**(3): 153-163.
- Mueller, M. (2002). Ruling the root : Internet governance and the taming of cyberspace. Cambridge, Mass., MIT Press.
- Mueller, M. (2002a). Ruling the root : Internet governance and the taming of cyberspace. Cambridge, Mass., MIT Press.
- Mueller, M. (2002b). Success by Default: A New Profile of Domain Name Trademark Disputes under ICANN's UDRP. Syracuse, NY, The Convergence Center.
- Mueller, M. (2002c). "Competing DNS Roots: Creative Destruction or Just Plain Destruction?" Journal of Network Industries **3**(3).
- NERA (1998). Feasibility Study & Cost Benefit Analysis Of Number Portability For Mobile Services In Hong Kong. HK, National Economic Research Associates Economic Consultants (NERA) and Smith System Engineering for the Office of the Telecommunications Authority (OFTA).

- NTIA (1998). Management of Internet Names and Addresses. Washington, DC, National Telecommunications and Information Administration, U.S. Department of Commerce.
- Paskin, N. (1999). Toward Unique Identifiers. IEEE, IEEE.
- Reiko, A. and J. Small (1999). The Economics of Number Portability: Switching Costs and Two-Part Tariffs: 1-30.
- Rood, H. (2000). "What's in a Name, What's in a Number: Some Characteristics of Identifiers on Electronic Networks." Telecommunications Policy **24**(6-7): 533-552.
- Rutkowski, A. (2001). "The ENUM Golden Tree: The quest for a universal communication identifier." Info **3**(2): 97-100.
- Saltzer, J. (1993). On the Naming and Binding of Network Destinations, Internet Society. **RFC 1498**.
- Viard, V. B. (2001). Do Switching Costs Make Markets More Or Less Competitive? The Case Of 800-Number Portability.
- Vixie, P. (1994). External Issues in DNS Scalability. The Global Information Infrastructure, Annenberg Washington Program, Washington, DC, Science, Technology and Public Policy Program John F. Kennedy School of Government, Harvard University.
- WIPO (1999). The Management of Internet Names and Addresses: Intellectual Property Issues. Geneva, World Intellectual Property Organization`.
- WIPO (2001). The Recognition of Rights and the Use of Names in the Internet Domain Name System: Report of the Second WIPO Domain Name Process. Geneva, World Intellectual Property Association.