Changing markets for Domain Names: Technical, Economic, and Policy Challenges

William Lehr
David Clark
Steve Bauer
MIT CSAIL

Abstract

The Domain Name System (DNS) is critical infrastructure for the Internet. The growth in size and complexity of the DNS, together with changes in its governance and operational practices, have raised numerous challenges. Herein, we provide a qualitative and empirical overview of DNS ecosystem economics and the relationships among key participants, including ICANN, the registries, registrars, and registrants. With this as background, the paper examines three key issues: (1) market power; (2) trends impacting the importance of domain names; and (3) concerns over DNS abuse and security.

Changing markets for Domain Names: Technical, Economic, and Policy Challenges

Table of Contents

1. Introduction

This paper provides a high-level characterization of the technical and economic role of the Domain Name System (DNS), which comprises an important part of critical infrastructure for the Internet.[1] Our particular focus is on the economic structure of the DNS ecosystem, how economic considerations shape that ecosystem, and several of the important challenges confronting the DNS, and hence, the future of the Internet.

The key technical role for the DNS is as identity and address management infrastructure for the global Internet. The DNS provides the technical mapping from human-understandable names to the network addresses that are used to forward traffic in the Internet. Moreover, because the names have economic value as intellectual property and marketing assets (e.g., they are associated with "brands"), legacy design decisions underlying the DNS give rise to its role in "routing money in the Internet," which is short-hand for explaining how the DNS impacts commerce.[2] In Section 2, we review the technical and economic role that the DNS plays in the Internet and describe the institutional governance structure controlling the DNS. This includes identifying the key stakeholders that comprise the DNS ecosystem. In Section 3, we summarize how dollars flow in the ecosystem.

An organization called the Internet Corporation for Assigned Names and Numbers (ICANN) has overall responsibility for the governance of the DNS. Since ICANN's creation in 1998, much has changed in the Internet ecosystem. At a high-level, the Internet has scaled exponentially in size (number of connected nodes, traffic volumes), economic importance (share of trade, sectors of economy and business functions dependent on Internet), and global scope. These macro-environment changes have substantially altered the stakeholder context in which discussions over Internet governance take place. For example, although the Internet first emerged to prominence in the U.S. and its technical design and governance have reflected that legacy, today's Internet is global and a larger share of future growth is expected to be associated with content, applications, and connected end-points that are not US-centric, English-language-focused, and not even human (e.g., IoT, M2M). It is unclear whether the Internet will continue to serve as a (relatively) open platform for global connectivity, or whether it will fracture into multiple regional, national or otherwise (partially) closed user groups as a result of changing geopolitics and market trends. One key factor in this is rising concerns related to cybersecurity (as more of the world is Digital, more of the crime will be there as well), and how these concerns about security will be addressed.

[1] This work was supported in part by the U.S. National Science Foundation under Grant C-ACCEL OIA-1937165. The authors would also like to acknowledge helpful comments from Michael Kende, Roslyn Layton, Milton Mueller, and Catherine Tucker. All views expressed and any omissions or errors are the authors alone.

[2] Because the Internet is widely used as a platform for search and contracting for trade in both real and digital goods (and additionally, as a platform for the distribution and consumption of digital goods), its impact is economy-wide, and not limited to on-line commerce, as narrowly defined and reported in government accounting statistics.

In Section 4, we summarize some of the ways in which the DNS ecosystem has changed so as to raise a number of important questions. In Section 4.1, we address the question of where market power concerns may arise within the DNS ecosystem, since that is a key consideration that colors calls for reforms to DNS governance. The concerns over market power are multifaceted. On one hand, concerns over too much market power by some stakeholders is used to justify enhanced regulatory oversight, with conflicting perceptions on the benefits from further expansion in the number of TLDs. On the other hand, the fragmented market and governance structure in the DNS raises questions as to whether *any* entity (or collection of entities) have sufficient governance power over DNS markets to effectively address abuses.

In Section 4.2, we speculate about how trends in market forces and alternative future visions of how the Internet may evolve would impact the DNS ecosystem. Changes in the nature of online commerce, search, how firms organize their marketing and branding efforts, and the changing nature of Internet applications are shifting the role of domain names. It is possible that these changes will render the DNS less technically and economically relevant in the future.

In Section 4.3 we address the challenges that the DNS faces in protecting against cybercrime and DNS abuse. Even if the economic concerns arising from market power abuses are resolved, the growing threat from cybercrime at all levels, including in the form of attacks on critical networking infrastructures like the DNS is a problem that will likely only grow in coming years.

Section 5 provides summary conclusions and suggests directions for future research.

2. Understanding the DNS and its Structure

The DNS ecosystem may be understood from a technical perspective (how it functions), from the perspective of industry structure (who are the actors involved in its realization), and an economic perspective (what are the incentives of those various actors). From a technical perspective, Geoff Huston has characterized the DNS as referring interchangeably to several distinct concepts: "It's a structured namespace, a distributed database, the protocol we use to query this database and the servers and services we use to make it all work."[3] The structured namespace allows domain names (arbitrary strings of characters and symbols) to be mapped to the IP addresses used to identify destinations for traffic in the Internet.

If users or their applications were comfortable working directly with IP addresses, there would not be a need for domain names. However, introducing domain names served a number of important technical functions. First, domain names can use human-meaningful strings of characters to identify destinations on the Internet such as the websites of particular organizations (www.ibm.com, www.mit.edu, or www.google.com).[4] Second, those domain names do not need

[3] See Huston (2019), page 1.

[4] One motivation for modifying the DNS was to include non-English character and symbol sets to include strings that are meaningful to non-English speakers. (However, it is worth noting that even when restricted to the ASCII character set, many feasible string combinations constitute nonsense combinations that are not be human-meaningful.) The introduction of Internationalized Domain Names (IDNs) was thought necessary to "prevent balkanization of the Internet" (see page 13, NRC, 2005). New gTLds support IDNs

to change when the mapping to the underlying IP addresses or network routing changes. As a result, human-meaningful domain names have economic value as identifiers that are independent of the underlying network address and routing infrastructure. Domain names can be associated with economically valuable assets in the off-line world and are inextricably linked to the "brand" value or intellectual property of the named assets (e.g., IBM or Google the company, MIT the university).[5]

The domain name space is hierarchically structured, which enables the DNS to be implemented as a distributed database, which in turn is key to ensuring good performance, scalability, and resiliency. The highest level element in a domain name is called the Top Level Domain or TLD. Examples include .com, .edu or .org. For each of the TLDs, there are a set of servers that record the addresses of the name servers for the next level domains registered in those TLDs, for example google.com or mit.edu. To find the servers that manage the TLDs, there are a set of *root servers,* which keep track of the location of the TLD servers. The addresses of the root servers in turn are well-known and unchanging, and are statically embedded into software that performs domain name resolution. There are 12 different organizations that maintain versions of the root database, and most of these organizations have many replicas of their service positioned across the globe, so that the data in the root servers is highly available and resilient.[6] Additionally, having the root servers distributed globally helps with performance (reduces latency) by reducing the distance queries may need to travel.

The TLDs are organized into several categories that differ with respect to how they are managed. The different categories of TLDs are generic TLDs (gTLDs), country code TLDs (ccTLDs) and brand/community (or sponsored) TLDs. The gTLDs will be the focus of most of our discussion here. The gTLDs include the legacy domains such as .com, .net, and .org. Prior to 2010, there were fewer than 20 gTLDs, but the number of gTLDs was greatly expanded in 2010 by the addition of 1,200 new gTLDs. The ccTLDs include the TLDs for sovereign governments such as .us (United States), .eu (European Union), .ai (Anguilla),[7] or .cn (China). The management of the ccTLDs is delegated to the sovereign governments. Finally, the brand/community TLDs include TLDs typically assigned to specific branded companies like .google or .ibm, while the community TLDs

---

in non-ASCII character sets, and so do a number of legacy gTLDs and ccTLDs. Introducing support for IDNs required software upgrades and is not without problems. For example, the International Chamber of Commerce raised its concerns in 2006 (see ICC, 2006) and others have noted how IDNs can be exploited to launch cyberattacks (e.g., see Liu et al., 2018). As of December 2018, only 2.5% of globally registered domains were IDNs and universal support for IDNs by browsers, email, and other applications continues to pose a challenge. The USAG, a multistakeholder body set up to promote universal acceptance of all valid domain names (including IDNs) provided estimates that total annual benefits globally of accomplishing that goal would be close to $10 billion (see Analysys Mason, 2017).

[5] Domain names may be trademarked, and thereby assume intellectual property protection in many jurisdictions. This status can be challenged and the linkage between intellectual property rights and the DNS is complex and evolving.

[6] There are 13 root servers spread around the globe, operated by 12 organizations (with Verisign operating two of the root servers) (see https://www.iana.org/domains/root/servers).

[7] This is of interest since Anguilla has realized revenue from allowing domain registrations in ".AI" for companies seeking to signal a connection to Artificial Intelligence (AI).

are for domains for a clearly delineated community of interest, like .broker or .beer.[8] There were several community TLDs (e.g., .org for non-profits and .edu for educational institutions) among the legacy (pre-2010) TLDs, but the increase in the number of TLDs that occurred after 2010 made it feasible to also have brand TLDs and additional community TLDs.

2.1. Industry structure

There are three parts to the industry structure that underpins the DNS: the firms that maintain the hierarchy of databases, the firms that manage the leasing of second level domain names inside TLDs, and the firms that support the action of performing the query.

The TLDs are managed by firms called *registries.* They are responsible for maintaining the information about the domains that are registered in the TLD, although they may outsource the actual management to a *back end* service provider.

For the generic TLDs, registries obtain the right to manage a TLD by licensing that name from ICANN, which involves the payment of considerable fees. We discuss the relevant economics in Section 3. The situation with respect to ccTLDs is different. The ccTLDs refer to internationally recognized geographic territories, which in many cases are sovereign states. In the case of sovereign states, ICANN typically delegates significant authority for the management of the ccTLD to the sovereign entity it is assigned to.[9] Some countries have used their ccTLD for domain names associated with their country, others have decided to monetize their names in a way similar to gTLDs because they seem to be popular TLDs in which to register second level domain names. Examples include .io, .ly, and .tk.

When an organization wants to obtain a second-level name in a TLD, that business is handled by *registrars,* who provide a retail front-end to the registries. Registrars must be accredited by ICANN to provide this service for gTLDs.[10]

---

[8] Community gTLDs may be sponsored or unsponsored. In both cases, the intention is that the registrants in the community gTLD comprise a common community of interest. The sponsored gTLDs have a well-identified sponsoring organization that is identified in the contract with ICANN and has a charter that specifies what entities are eligible to register in the domain. Examples of sponsored gTLDs include .aero (for members of the air-transport industry), .edu (for institutions of higher learning), .gov (for US local, state, and federal government), etcetera. With the expansion of gTLDs after 2010, community gTLDs no longer need to have a well-defined sponsoring organization (see https://www.icann.org/en/system/files/files/statement-sponsored-community-based-tld-models-06may11-en.pdf).

[9] See Mueller & Badiei (2017) for a detailed discussion of the complex history of ccTLDs and the legal issues that arise over the property rights to ccTLDs assigned to sovereign entities. Originally, the ccTLDs were assigned to geographically-meaningful territories, but over time disputes have arisen as to whether the right to control a ccTLD should be regarded as a sovereign right, a property right (that may be treated like private property), or perhaps, under some other model such as a public trustee model where ICANN is the trustee.

[10] ICANN does not provide accreditation oversight for registrars selling domain names in ccTLDs. Registrars are free to sell domain names for ccTLDs whether they are or are not accredited by ICANN.

In general, the domain name service that is responsible for a particular name is called an *authoritative* resolver for that name. The registries are the authoritative resolvers for the TLDs, those who have licensed second level domain names must set up authoritative name servers for those names, and so on.

When a user needs to translate a name into an address, the software acting on the user's behalf (e.g., a browser) will usually contact a local service called a *stub resolver* to start the process. The stub resolver may be software running on the same machine as the browser or on a home router. The role of the stub resolver is minimal, as all it typically does is hand off the query to a *recursive resolver.* The recursive resolver sends successive queries into the domain name hierarchy until it finds the desired answer. First, the resolver may have recently looked up the name and *cached* the answer to the query, in which case it just returns that answer. If not, it must forward the query to the right part of the domain name hierarchy to get the answer. To do this, it again refers to its cache. For example, to resolve the name www.example.com, the recursive resolver would first see if the address associated with this name is in the cache. If not, it would see if it has the address of the authoritative name server for example.com in the cache. If so, it would send the query for that name to the address of the example.com authoritative name server. If not, it would see if it had the address of the authoritative name server for .com in its cache, and if so, query that name server for the address of the authoritative name server for example.com. If it did not have the address of the authoritative name server for .com, it would consult its built-in list of root name servers, and send a query to get the address of the authoritative name server for .com. The term *recursive* is used to describe this service because of this recursive pattern of sending intermediate queries and getting intermediate responses.

The final element of this process is the step by which the host at the edge (or the stub resolver acting on behalf of the user) selects a recursive resolver to use. Over-time, the way that recursive resolvers have been provided has evolved.[11] Originally, the recursive resolver was a service provided by the ISP that provided the end-host with Internet access, and ISPs like Comcast, Charter, and Verizon continue to provide this service. To many, supporting DNS resolution simply represented basic functionality of the Internet and a necessary component of Internet access service. When a host first powers up and connects to the Internet, it runs a protocol called the Dynamic Host Configuration Protocol, or DHCP. DHCP sends a query to the Internet Service Provider for the host, and one element of that query is to get the address of a recursive resolver. Normally, the ISP would return the address of its own recursive resolver, and all the parts are now in place to resolve a query.

[11] See Huston (2020) for a readable explanation of how DNS resolution has evolved so as to challenge fundamental notions of what constitutes the Internet. Huston concludes: "If the characterisation of what makes the internet a single network is a single address space and a single name space, then it's pretty clear that we've dispensed with a coherent and uniform address plan already as NATs are just so pervasive. But if a coherent name space is all that's left of a single unifying Internet what happens when we tear that apart as well." (NATs refer to Network Address Translation which are boxes at the edge that take public IP addresses and map them opaquely to downstream end-nodes on the end-user's private network. NATs were introduced both to address the scarcity of IPv4 addresses and to provide a firewall to separate private networks from the public Internet). See Lehr, Clark, Bauer, Berger & Richter (2019) for further discussion of how conceptions of the Internet are evolving.

Over time, however, the Internet ecosystem has evolved to allow mix-and-matching of capabilities provided by different platform providers.[12] For example, in 2009, Google launched a public DNS resolver service located in the Google cloud to which stub-resolvers could point for recursive DNS resolution, thereby unbundling the function of DNS resolution from basic Internet access functionality.[13] To use the Google recursive resolver, the user would have manually installed the address of that server, overriding the address provided by DHCP. That address would typically be recorded in the operating system of the host, which would provide a means to change it.

One aspect of the DNS that has motivated economic tussles is the market-intelligence value of the meta-data provided by being able to observe DNS queries, which reveals what domains end-users are interested in communicating with. The entity best positioned to do this is the recursive resolver. Historically, DNS queries were sent in the clear, so any entity that could observe the traffic could also gather this data, and being able to observe enables the observer to profile users' on-line behavior. The data that can be gathered is potentially valuable market-intelligence, but at the same time capturing and using this data may pose a threat to end-user privacy.[14] Concerns over securing the DNS and protecting end-user privacy have helped motivate efforts to encrypt DNS queries and prompted changes in how DNS queries are resolved.[15] If the query to the recursive resolver is encrypted in transit, then only the recursive resolver can see it. Passive observers in the network cannot usefully observe it. Encryption of the query puts a premium on the question of who picks the resolver. Today, the typical browser sends a DNS query using the address of the resolver stored in the operating system, which the user can change. But it need not be this way. The browser could, for example, send a DNS query itself to a recursive resolver that it picks, perhaps using an address for a resolver that is "baked-in" to the browser software at the time it is shipped. It might or might not be possible for the user to change this address.

Another aspect of a recursive resolver that may not have been sufficiently appreciated in the early days of the DNS is that it is in a position to block a query or lie about the answer. Blocking can be beneficial if the blocking is done on behalf of the users to prevent unintended connections to a malicious site, or totally adverse to the interests of the user if the blocking is imposed by a state

[12] See Lehr, Clark, and Bauer (2019).

[13] Pointing your platform (e.g., network setting on your laptop) to the DNS Server 8.8.8.8 will make use of Google's open DNS resolver service. Other such services are provided by Cloudflare (1.1.1.1) or Quad9 (9.9.9.9), and a host of other options are available. Many users are unaware of their ability to manage the selection of DNS Server, and rely on their application or ISP to set default configurations and are unaware of how DNS queries are resolved.

[14] For example, it is useful for understanding user demand for products and services, including for directing targeted online advertising; it is useful for provisioning network resources and load-balancing; and it is useful for strategic business planning, including responding to competitors. Large datasets of end-user on-line DNS query behavior can be used to train machine learning programs and asymmetric access to such data may provide a source of competitive advantage.

[15] See Huston (2020) and Note 11 *supra* for a discussion of these, including the battle over DNS-over-HTTPs v. DNS-over-TLS, etc.

actor as a part of censorship. A resolver that lies could send the user to a malicious clone of a web page that then steals user credentials or commits other harms.[16]

One response to this tussle is to conclude that the whole idea of a shared recursive resolver is a bad one, and every host should run its own recursive resolver on behalf of its own users.[17] But what this set of examples illustrates is that the DNS (and in particular the recursive resolver) has evolved from a simple functional element that to its designers was just a means to avoid having to remember IP addresses to a space of power struggle shaped by concerns over security, privacy, control and economic advantage.

There are some other important industry players in the ecosystem. The registries run the authoritative name servers for the TLDs. But the lower-level domain names also require authoritative names servers. Thus, in the example of www.mit.edu, the registry can return the address of the authoritative name server for MIT, and then the recursive resolver must query the authoritative name server for mit.edu to find the address of the machine with the name www.mit.edu. Historically, most owners of domain names ran their own authoritative name server for their domain name, but the trend has been to out-source this function to a third party. Providers of authoritative name service today can provide very sophisticated options for their clients, such as directing a user to a copy of the service that is in close proximity to that user, by picking among a number of IP addresses based on the location of that user and other considerations.

A final class of players in the DNS ecosystem are the back-end firms that provide the technical platforms needed by the registries and registrars to make the DNS work. Many of the registries outsource the back-end services to other providers of these technical platforms, and many of those who provide those services in-house, also provide those services on behalf of other registries.[18] Some of the better-known backend service providers are CentralNIC, ZDNS, Donuts, Neustar, and Afilias, although this business has been changing rapidly.[19]

---

[16] The protocol extension, DNSSEC, was introduced in 1997 to help secure DNS data. DNSSEC facilitates authentication and data integrity checks for DNS queries by adding cryptographic keys to ensure that the responses are from a "real" DNS server and not a spoofed one. Although forged and spoofed DNS data continues to pose a problem, adoption of DNSSEC is still not universal, although ICANN and others continue to recommend using it. For further information, see https://www.icann.org/resources/pages/dnssec-what-is-it-why-important-2019-03-05-en and for statistics on the use of DNSSEC see https://www.internetsociety.org/deploy360/dnssec/statistics/.

[17] See Schomp, Allman and Rabinovich (2014).

[18] In 2009 (before the new gTLDs were created), KPMG (2010), estimated that 31% of registries outsourced their registry functions to backend providers; and 31% of those who provided their own-registry functions in-house, offered those services also to other registries. It appears that the outsourcing of back-end services by new gTLDs is more common and that just three providers (Rightside, Neustar, and Afilias) may account for as much as 90% of the registrations in those new gTLDs (see page 44 in CCT, 2018).

[19] See https://ntldstats.com/backend for the backend providers. For 33 million domains, associated with 1,178 TLDs, nTLDStats identifies 37 registry back-ends (visited November 20, 2020). However, it is worth noting that GoDaddy, the largest registrar, acquired Neustar in April 2020 (see https://www.prnewswire.com/news-releases/godaddy-acquires-neustars-registry-business-

Because of its role in supporting basic Internet functionality, the security of the DNS ecosystem is a key concern for cybersecurity. Protecting the DNS from being misused (e.g., to disrupt the resolution of DNS queries, return incorrect mappings, redirect traffic, or engage in other forms of cybercrime) is becoming an ever-more pressing concern. All of the authoritative name servers, because they store the address to which a name is mapped, could be a target for penetration. If a hacker can change the IP address associated with a name, it can redirect all the traffic intended for the service at that name to its (presumably malicious) server instead. We do not catalog all the attacks that have been launched at the DNS, which exploit actual system penetration as well as exploitation of configuration errors committed by the holders of names, but the situation is serious. Operators of authoritative name servers, including the registries for the TLDs, must operate their system at a high level of care with respect to security vulnerabilities.

2.2. Governance of the DNS

The entity responsible for managing the data in the root servers and the allocation of gTLDs, and hence, at the center of the governance structure for the DNS ecosystem, is the Internet Corporation for Assigned Names and Numbers (ICANN), which was established in 1998 as a non-profit corporation in California.[20] ICANN oversees the DNS via a series of contracts through which the actual management of the domain name system is delegated to *Registries* that are responsible for managing TLDs.[21] The gTLD registries may be controlled by for-profit entities like Verisign,[22] which is the authoritative registry for several of the largest TLDs, including .com, .net, .edu, and .gov,[23] and non-profit entities like the Public Interest Registry (PIR) which is the registry operator for .org.[24]

301036134.html); and Donuts acquired Afilias in December 2020 (see https://afilias.info/news/2020/12/29/donuts-acquires-afilias).

[20] See https://www.icann.org/en/history/icann-usg for a history of ICANN and ICANN (2013) for an infographic that lays out the different governing bodies that oversea the management of the DNS. These include the Internet Architecture Board (IAB) that helps set standards and direct research; the Internet Governance Forum (IGF) that provides multi-stakeholder input; and sundry other organizations that interact and collectively share responsibility for the on-going management of the DNS.

[21] Verisign (2021) reports that there were 366.3 million domain names as of December 2020, comprised of 158.9 million ccTLDs (43%), 26.0 million new gTLDs (7%), and 182.0 million legacy gTLDs (50%). For new gTLDs, the site nTLDStats lists 504 registries supporting 33 million domain names across 1,178 TLDs as of November 20, 2020 (see https://ntldstats.com/registry/group).

[22] See https://www.verisign.com/en_US/domain-names/domain-registry/index.xhtml.

[23] For .com, .net and several other gTLDs, Verisign is the registry operator, providing both front- and back-end services for domain registrants; whereas for .edu and .gov, Verisign performs the registration and resolution functions on behalf of the US Government.

[24] See https://thenew.org/. In 2019, the DNS ecosystem was roiled by the decision by the ISOC Board to approve the sale of the .org TLD to a private equity firm in November (https://www.icann.org/news/blog/icann-board-withholds-consent-for-a-change-of-control-of-the-public-interest-registry-pir) for what was later disclosed to be for $1.135 billion and in a transaction that had links to individuals that previously had been involved in ICANN governance. The furor over the way the deal

The entities that control a TLD registry may control multiple TLDs. However, in a number of cases, the controlling entities are not publicly-traded companies or are located in other countries with incompatible reporting requirements. This means that the precise structure of the registries and how they relate to each other is not always clear. For example, Donuts is the US-based business entity that manages 100s of gTLDs, but is privately held. Were one to fail to aggregate the domain names controlled by all of the entities that operate under the Donuts umbrella, one would get a misleading perception of the scale of Donuts operations within the DNS ecosystem.[25] Although Donuts is reasonably transparent regarding its role within the DNS ecosystem (although it does not publish its financials), the same cannot be said for many of the business entities that control gTLDs. One advantage of having non-profit or, if for-profit, publicly-traded business entities (like Verisign) manage the TLDs is that that provides better access to business information, including financial information.

ICANN delegates control over TLDs to registry operators via a standardized contract or Registry Agreement (RA)[26] that assigns to the registry an exclusive franchise to manage names in each TLD.[27] Back in 1998, the existing entity that functioned as both the sole registry and registrar (at the time, Network Solutions) was required to separate their wholesale, back-end operation of the TLD from the retail operations of selling domain registrations to end-user registrants. This history led to the current separation of *registrars* and *registries*. Presently, registrars are subject to accreditation by ICANN via accreditation agreements. The RAs require registries to only work with accredited registrars and the registries sign Registry-Registrar Agreements (RRAs) that certify the registrar's authority to register domain names in the registry's gTLD.[28] Registrars may provide retail registration for multiple TLDs and many registrars may compete to offer domain

was managed resulted in ICANN deciding not to approve the transfer of the registry agreement in May 2020 (https://www.eff.org/deeplinks/2020/04/victory-icann-rejects-org-sale-private-equity-firm-ethos-capital). In addition to concerns over the how the deal was structured, opponents raised concerns that transferring ownership of the .org TLD to a for-profit business was inconsistent with the .org mission of providing domain registry services to non-profits.

[25] See https://donuts.domains/what-we-do/top-level-domain-portfolio/.

[26] See https://www.icann.org/resources/pages/registries/registries-agreements-en. The template RA is close to 100-pages long and includes a number of standard articles and sundry attachments. The text and attachments may be slightly modified to fit the requirements of particular TLD registries. Each RA is for a duration of 10 years with subsequent renewals, and includes covenants detailing the responsibilities of RAs and laying out the basic operating principles. The RAs also include the fee terms which are typically $6,250 per year and $US 0.25 per DNS registration transaction (where transactions include domain name increments/renewals, transfers, once the first 50k transactions have occurred). The RAs also include clauses outlining termination and appeal procedures.

[27] Registries are granted technical exclusivity to assign names in the TLD to ensure that those names are uniquely assigned.

[28] The RAs constrain the registries to offer non-discriminatory RRAs to all accredited registrars, and there is a formal process requiring ICANN approval, if registries wish to amend their RRAs (see https://newgtlds.icann.org/sites/default/files/agreements/agreement-approved-31jul17-en.html#article2.9).

registrations in the same TLDs.[29] The registrars also typically offer additional retail services to domain registrants, including such services as email, search-engine optimization (for marketing), web-hosting, security, and other ancillary services. As with registries, registrars operate under a variety of business models and multiple registrars may be controlled by the same enterprise. Among the largest registrars are GoDaddy (US), NameCheap (US), Alibaba (China), and TuCows (Canada).[30]

Finally, the registrants of the domain names lease control of those domain names via lease agreements with the registrar that are subject to Acceptable Use Policies (AUPs) and other terms specified in the domain registration contract.[31] Registrants may transfer their registrations from one registrar to another, but a domain name registered in one TLD cannot be transferred to another TLD. Thus, a registrant who wants similar domain names in multiple TLDs needs to register the domain in each of those TLDs separately.[32]

[29] Accredited registrars may authorize resellers to sell on their behalf, but the resellers are not accredited by ICANN. The resellers are governed by their contract with the accredited registrar they resell for, and that accredited registrar is governed, in turn, by its RRAs for the various registries and by accreditation agreement with ICANN.

[30] For data about registrars of new gTLDs, see https://ntldstats.com/registrar. For 33 million domains across 1,178 TLDs, they list 367 registrars operating as of November 20, 2020. Alibaba Cloud (https://www.alibabacloud.com/en) is the largest registrar (in terms of domain names registered) and is part of the Alibaba Group of Chinese eCommerce companies, which operates the largest on-line shopping site in China (TaoBao). Alibaba (eCommerce), Tencent (which operates a social network and owns WeChat) and Baidu (search company) comprise the Chinese answer to Amazon, Google, and Facebook. NameCheap (https://www.namecheap.com/) is a privately-held US-based registrar; GoDaddy (https://investors.godaddy.net/investor-relations/overview/default.aspx) is a publicly-held US-based registrar; and TuCows (https://www.tucows.com/) is a publicly-held Canada-based registrar. Since new gTLDs represent less than 10% of total registered domain names (see Verisign, 2021), this data is not representative of all registrar activity (which includes ccTLDs and legacy gTLDs). A comparable set of data was not available at the time this report was prepared.

[31] Generally, these AUPs are not negotiable and are click-through agreements that registrants must accept.

[32] Businesses that want to protect their trademark interest in a domain name from being registered separately by another registrant in a different TLD may feel compelled to register their desired domain in multiple TLDs (e.g., register the string "ibm" in ibm.tldx, where tldx is another TLD where "ibm" might be registered by a party that uses it to threaten IBM's trademark). The idea that registrants might feel compelled to invest in registering in multiple TLDs solely to protect their interest in the trademark is referred to as *defensive registrations* and poses a hold-up threat for registrants with a brand to protect. Note, this can also extend to registrants anticipating a new branding campaign (or politicians planning an election campaign) that want to protect against counter-sites being set up with the same name in another TLD. Abuses of the DNS system to exploit such opportunities include "cybersquatting" where bad-faith registrations are undertaken to extract rents later from the registrant who really needs the site; "front-running" where pre-emptive registrations stake out an interest in a domain name ahead of a company's planned need for the domain name; and gripe-sites or otherwise offensive sites like "Xsucks.com" to direct attention to material critical of X.com (see RAPWG, 2010). Whether all of these uses constitute abuses is open to debate. For example, many might regard "gripe-sites" as allowed ways to express an opinion.

2.3. The value of names

One of the most significant actions that ICANN has undertaken is to greatly expand the number of TLDs that it licenses.

The arguments in favor and in opposition to this move were passionate.[33] One argument was that the limit on TLDs was creating an artificial scarcity in names. (Obviously, names are just strings of letters and numbers, so why limit their creation?) Artificial restriction on the number of TLDs has been cited as a cause for economic scarcity, and control of a scarce resource can give rise to economic power that may be reflected in prices being set substantially above cost-based levels or what would be sustainable in the face of effective competition. Even more concerning is the potential for abuse of a scarce resource to foreclose beneficial innovations, market growth, and/or competitive entry – or the ways in which markets work to ensure allocative, productive, and dynamic efficiency.

On the other hand, there was resistance to expansion of the TLDs, primarily from security practitioners concerned about the potential for increased DNS abuse. Critics called for (a) first demonstrating that there is a real economic need and that net benefits should be expected from further expansion of the TLDs;[34] and (b) recognizing that oversight is already overly challenging with today's TLD landscape. Critics motivated either by market power or security concerns feared that further expansion would either not address the relevant market power concerns (e.g., related to legacy control of .com) or would render efforts to establish needed oversight controls even more difficult (e.g., by providing pressure to relax price controls on legacy TLDs like .com). Thus, among those with concerns, there were and are conflicting views as to the benefits of expanding the TLDs.

Although it was not completely clear that the DNS ecosystem could technically handle the scaling of the DNS to several orders of magnitude in the number of TLDs, the growth from less than 20 in 2010 to almost 1,200 today has demonstrated that scaling the DNS by expanding the number of

[33] There is a voluminous public record documenting support and opposition over decisions to expand the number of gTLDs. For example, ICANN's website documents its successive decisions to add new gTLDs before 2010, as well as the significant increase in gTLDs that occurred after 2010. As part of its on-going process and internal rules, ICANN has created a series of review committees. One of these is the Competition, Consumer Trust and Consumer Choice Review (CCT) which issued its report in October 2018 that highlighted both positions in support and in opposition of further expansion of the gTLDs, which concluded that ICANN needed to collect more data and conduct further economic research before proceeding with further expansion in the number of gTLDs (see, ICANN CCT (2018), "Competition, Consumer Trust, and Consumer Choice Review Team (CCT) Final Report Now Available for Public Comment," 10 October 2018, available at https://www.icann.org/news/announcement-2018-10-10-en). For an earlier summary of comments opposing the decision to significantly expand the number of gTLDs in 2010, see the blog post at https://blog.caida.org/best_available_data/2011/01/19/thoughts-on-icanns-plans-to-expand-the-dns-root-zone-by-orders-of-magnitude/.

[34] To date, there has still never been any comprehensive economic study undertaken to evaluate empirically the economic need for or net economic benefits realized from expanding the gTLDs either before 2010 or thereafter.

TLDs does not pose any significant technical challenges, at least with respect to the processing capacity of the DNS to support many more TLDs and many more domain names.

The market forces of supply and demand for domain names, competition among different players in the DNS ecosystem, technical innovations and changes in market tastes and conditions all contribute to shaping how markets for domain names are governed. However, the complex multi-stakeholder governance structures complicate the analysis of the markets for domain names. Although names at times appear to be amenable to being bought and sold subject to competitive market dynamics that balance supply and demand, domain names are not like typical real goods. The supply of desirable domain names is artificial in that additional TLDs can be created at will.[35]

Registrants may sell their domain names, and the potential to sell valuable domain names has created a secondary market for domain names in TLDs with attractive character strings, where what is attractive may depend on the buyer. Someone who wants a domain name can go to a registrar and see if it is available and if it is available, obtain it directly from the registrar. The standard RA and registrar agreements impose minimal restrictions on the wholesale and retail prices set by registries and registrars.[36] The absence of pricing restrictions allows registries and registrars to price differentiate based on the string of characters. For example, JoesGarage.*TLD* and Jegorsaaeg.*TLD*; or 12345.*TLD* and 1234.*TLD* may be priced differently. The former might be because JoesGarage.*TLD* represents a string that is meaningful to humans, while Jegorsaaeg.*TLD* is a random string.[37] The latter is meant to reflect different pricing for domain names of different lengths. Often shorter names are more valuable (e.g., because they require end-users to type fewer characters). Many registries and registars do set prices that differ based on the domain name character string and most have created so-called "premium" names that are priced above their general pricing. For most, the general pricing is offered in tiers or bands of pricing. In the case where a registrant finds that the desired name has already been registered, it is always possible to register that domain in some another TLD;[38] or if the registrant has a strict preference for a particular TLD, the registrant may be able to acquire the desired domain name on the secondary market. For example, for a registrant who wants a domain name in .com,

[35] None of the current legacy or new gTLDs have exhausted the supply of potential character strings.

[36] As explained further below, the wholesale pricing of .com domains by Verisign is subject to price regulations due to Verisign's contract with the US Department of Commerce. Those price restrictions were incorporated into the Verisign RA with ICANN.

[37] We use .*TLD* here to signal that this is the case for virtually all TLDs, whether new or legacy gTLDs or ccTLDs.

[38] For example, as we explain further below (JoesGarage.com is already registered, see Note 40 *infra*), but JoesGarage.ai (in the ccTLD ai for Antiqua, but often marketed as standing for "artificial intelligence") and JoesGarage.xyz (in the new gTLD .xyz) are both available from GoDaddy.com for $99 (for first two years) and $0.99 (for first year) – thereby illustrating the wide range of retail pricing available for similar domain names. An interesting side-story is that .xyz attracted a lot of attention when Google's Alphabet selected .xyz as the TLD for its new domain ("abc.xyz") in 2015 (see "Thanks to Google's Alphabet, .Xyz will end .Com dominance," Wired, August 11, 2015, available at https://www.wired.com/2015/08/alphabet-rewrites-the-domain-name-game/ (visited March 10, 2021)). This illustrates the significant and often misleading hype that is common in media coverage of the DNS ecosystem.

JoesGarage.*com* has already been registered (not surprisingly), while Jegorsaaeg.com is available.[39] Although JoesGarage.com is taken, it may be possible to contact the existing registrant and purchase it directly from them. The existing registrant may be a business that acquired the domain name for its own use (and may no longer need it) or a domain-name speculator who acquired the name for the express purpose of selling the name.[40] There are a number of secondary market sites that are designed to facilitate the purchase and sale of domain names that have already been registered.[41]

In addition to setting different prices based on the composition of the character string, registries and registrars may set different prices for initial registrations and renewals, and varying models exist both for registries and registrars.[42] Sometimes the initial registration is priced higher than the renewal price, which would be consistent with the need to recover the non-recurring, one-time costs of establishing a registration. In other cases, the initial registration is priced lower than the renewal price. That would be consistent with discounts offered to attract new demand.

Thus, governance of the DNS marketplace involves a complex collection of actors whose behavior is jointly controlled by a set of overlapping contracts (from ICANN to the registries to the registrars and then to the registrants). In the context of these contracts, registrants obtain names under renewable leases. The control of the terms and pricing of those renewable leases depends on a

[39] GoDaddy will license Jegorsaaeg.com (the same 10 character string but in random order so meaningless) under a variety of terms based on the services with which the domain name registration is bundled. The most basic registration is $17.99/year, with a discount offered for the first year. A more comprehensive bundle that includes email and additional eCommerce and security services is only $240 for the first years, so still a relatively inexpensive cost for even a relatively small legitimate business seeking a branded online presence. (Above obtained by going to https://www.godaddy.com/ and entering domain names into the "Find your perfect domain" search box, November 20, 2020.)

[40] For example, for a one-time fee of $69.99, GoDaddy will undertake to contact the owner of "JoesGarage.com" and negotiating a purchase price that if accepted will enable the new registrant to acquire the domain name, with a 20% broker service fee appended. Alternatively, if you want JoesGarage.net, GoDaddy can provide it for a one-time fee of $1,200 and then $19.99/year annual fee because GoDaddy can get it from the domain-name reseller (www.enom.com) that is the owner of the registration for JoesGarage.net. (This information is as reported at https://www.godaddy.com/ after entering the desired domain name into the "Find your perfect domain" search box, November 20, 2020). Interesting, if you go directly to www.enom.com and enter the for JoesGarage.net, the price is $1,380 plus renewals at $18/year [eNom is registry of record for this domain name but does not own].

[41] For example, SquadHelp.com offers a number of "premium" domain names for sale. SquadHelp.com describes itself as being the "world's largest platform for company naming and branding." It was founded in 2011 and is headquartered in Chicago area and provides a range of digital branding services, including selling so-called "premium" domain names, but its claim to being the "world's largest" is impossible to verify and is highly suspect. Nevertheless, under "brand names for sale," Squadhelp.com offers Mathematics.com for $523,750 (highest) and Ignitement.com for $750 (lowest). (Visited site November 20, 2020).

[42] Per domain wholesale and retail prices also may differ because of volume discounts or transaction cost differences. Moreover, it may be hard to infer the per domain prices if the domains are sold bundled with other services or are sold for different license durations.

mixture of instruments and forces. First, there are the rules and policies specified in the RAs and registrar accreditation agreements that commit registries and registrars to running a well-behaved market that includes obligations for reporting to ICANN, as well as payments to ICANN. Beyond the terms of these contracts, national governments that regard the Internet as important infrastructure may regulate many of the key Internet players under national telecommunications, competition, security, and other regulatory interests that constrain how those entities may act.

3. Follow the Money

Sizing the economic activity and tracing the flows of money within the DNS ecosystem is challenging for several reasons. First, much of the business activity that supports the DNS is undertaken by firms with broader interests in the digital economy, and the boundary between DNS-related activities and the rest of the Internet infrastructure ecosystem is indistinct. Second, the DNS ecosystem is global and many of the important players are privately-held or overseas (e.g., in China), complicating efforts to standardize their financials. However, in contrast to many other equally challenging industry segments, sizing the DNS ecosystem is aided by the fact that ICANN, as part of its organizational structure is required to transparently and publicly report lots of quantitative and business metrics; and further by the fact that the DNS ecosystem is heavily skewed toward a relatively small number of large businesses.

3.1. Sizing the ecosystem

In this sub-section, we summarize a number of data points that lead us to estimate that the total annual amount of DNS ecosystem-related revenue is around $8 Billion. Figure 1 provides a map of what we think we know about the money flows and the following describes some of the key inputs to this figure.

The earlier literature and sundry other sources provide some glimpses into the size of the DNS ecosystem.

1. For example, in 2003, Mueller and McKnight (2004) estimated that total global DNS service revenues were about $2 billion.
2. An OECD (2004) study reported information on the growth in domain name registrations and pricing trends, noting that actions taken by ICANN to reform the operation of the DNS ecosystem to promote competition (discussed further below) had resulted in price reductions that "saved consumers and businesses over USD 1 billion annually in domain registration fees."[43] If annual savings could exceed $1 billion, then that provides one indication of the aggregate level of spending.
3. A KPMG (2010) study looked at the operating costs of registry operations and determined that after a few years, those costs vary from $1 million to $2 million per year, depending on the size of the registry; but that costs may be several times that during the early years when the registry is setting up.

[43] See OECD (2004), p. 26. The OECD report noted that domain registrations had increased to 57 million by December 2003, reflecting a 40% growth rate since 2000; while a domain name that would have cost $35 per year in 2000 could be acquired for less than $6 in 2004 (see Tables 2 and 12). The OECD report also noted the wide dispersion in domain name pricing across different gTLDs, which remains characteristic of today.

4. Verisign (the registry for .com and .net) and GoDaddy (the largest registrar) earned 2019 annual revenues of $1.2 billion (159 million domain names in .com and .net) and $1.35 billion (79 million domain).[44] Those imply a wholesale price of around $7.50 for Verisign and a retail price of around $17 for GoDaddy.
5. The Public Interest Registry (PIR), which is the non-profit registry set up to manage the .org gTLD in 2003, earned $94.7 million (10.1 million domain names) in registration fees in 2019, which implies a wholesale price of $9.38 per domain name. PIR expected to send $67.5 million to the Internet Society and Foundation to fund its Internet-related outreach efforts in 2020.[45]
6. According to nTLDStats.com, there are 367 registrars and 504 registries active in the markets for the new gTLDs created after 2012, however the size of registries and registrars is highly skewed and those counts may not accurately reflect the universe of registries and registrars because it does not include data on legacy domain name sales and the business relationships among registries and registrars are often unclear.[46]
7. ICANN reports receiving $143 million in funds during FY2019, of which $84 million is paid by registries as required by the 1,200 Registry Agreements (RAs) and $48 million by the registrars according to the 2,459 Registrar Accreditation agreements. Payments associated with the authorization of new gTLDs and sundry other sources make up the difference.[47]
8. Sundry marketing research reports estimate the size of the global market for DNS services to be $350 million dollars. [48]

[44] See Verisign (2020a) and GoDaddy (2020). GoDaddy's total revenue is $2.988 billon, of which 45% is associated with GoDaddy's domain name products. GoDaddy also provides hosting/digital presence products (38%) and business cloud applications (17%).

[45] See PIR (2020).

[46] nTLDStats.com reports statistics on a per gTLD, per registry, per registrar, and per back-end provider basis for domain name registration activity associated with the new gTLDs. This excludes the registries and registrars that are only active in the legacy domains like .com and .net so represents an incomplete count.

[47] These numbers are for FY2019, which ends in June 2019, as reported in https://www.icann.org/resources/pages/governance/financials-en. For example, this provides link to ICANN financials (annual reports and various compilations of data by ccTLDs and by funding source – see https://www.icann.org/en/system/files/files/fy19-funding-source-spreadsheet-01nov19-en.xlsx) for current and historical years.

[48] The "DNS Services market" refers to the market for managed DNS services, whereby cloud-based service providers host authoritative DNS servers for enterprise customers to ensure DNS queries to their corporate networks are resolved correctly (see https://www.thousandeyes.com/learning/techtorials/managed-dns). Some of the larger providers in the managed DNS services market include Dyn (now part of Oracle), Cloudflare, Amazon Route 53, Cloud DNS, UltraDNS, Verisign Managed DNS, Neustar UltraDNS, and Akamai. Often the managed DNS services include additional services like management of email services and the provision of security services to protect enterprise websites from DDoS attacks and other types of cyberattacks. A number of market-research reports estimate the global size of this market at between $349 to $372 million (see https://www.industryarc.com/Research/Dns-Service-Market-Research-500909, https://www.zionmarketresearch.com/report/dns-service-market, https://www.marketsandmarkets.com/Market-Reports/dns-service-market-240632025.html).

9. As of June 2020, Verisign reported that there were 370 million total domain names registered globally and that 210 million of those were associated with gTLDs.[49]
10. BCG (2021) estimates the total annual revenue generated by secondary market activity at $2.1 billion.

The above data points present a partial picture of how to size the DNS ecosystem. If one had the financial statements of all of the registrars and registries, it might be possible to add these all up, after accounting for payments among different stakeholders, to arrive at a bottom-up estimate. Unfortunately, as noted earlier, that is not possible since a census of registrars and registries with comparable financials is not available.

A more promising approach may be to estimate the total flows in the ecosystem from the top-down to arrive at an estimate. Using the Verisign estimate of 370 million total domain names in gTLDs and ccTLDs and assuming a retail price of between $10 and $20 per domain name yields a total revenue estimate for domain name registrations of between $3.7 to $7.5 billion today. Computing the price for domain names is challenging because many registrars lease the domain names at a steep discount initially and may use domain sales as a loss-leader to sell other ancillary services like web-hosting, managed DNS, email, and other services. Also, the average prices for domain names are different across TLDs.

Taking the average of this top-down estimate gets us a ballpark estimate of the DNS ecosystem revenues on the order of $6 billion.[50] Adding in a recent estimate for the secondary market increases this to $8 billion.[51]

These estimates are summarized in Figure 1, which represents our stylized and preliminary attempt to map the money flows withing the DNS ecosystem.

≪ INSERT FIGURE 1 HERE≫

At the top, are the registrants, who are the ultimate source of demand for domain names. That includes large enterprises with extensive on-line presences, as well as small and medium businesses, government agencies and non-profits, and individuals that want to registrar a domain name. Although most of the domain names appear to be active, there are a significant number for which registrations exist (and for which fees are being incurred) that do not appear to be active.[52]

[49] See Verisign (2020a). Verisign runs two of the root zone servers and regularly publishes its "Domain Name Industry Brief" reports. This is one of the more widely used sources for sizing the domain name system.

[50] This should be taken as an order of magnitude estimate. Moreover, it fails to include the private investment by registrants in supporting domain name infrastructure, ancillary services (like the aforementioned global DNS management services that are estimated to be about $350 million globally), and secondary market activity.

[51] See BCG (2021).

[52] The universe of active domain names is observable from observing the root zone files and by checking to see if the entry points to an active name server for the domain; however, queries to many domain names are not resolvable to an IP address, which indicates they are not in active use.

Figure #1: Sizing the DNS Ecosystem

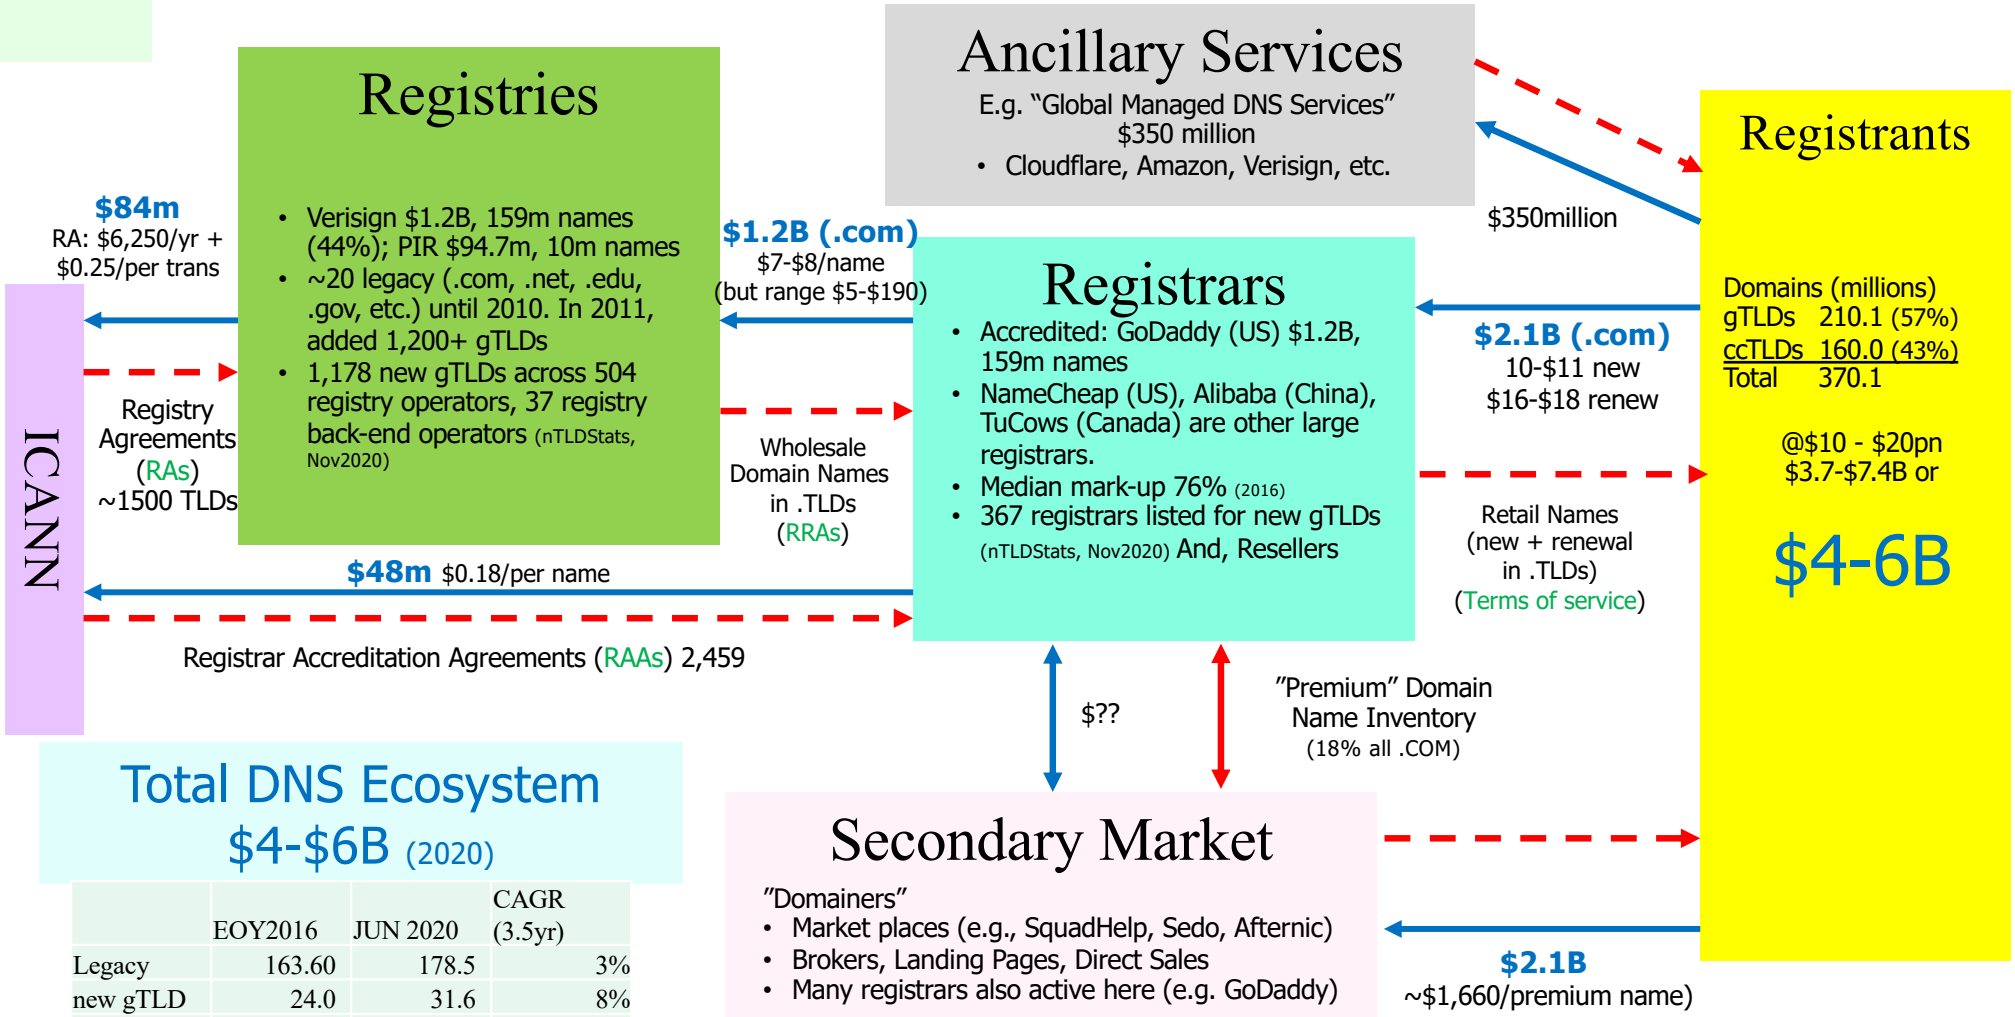

Total DNS Ecosystem
$4-$6B (2020)

	EOY2016	JUN 2020	CAGR (3.5yr)
Legacy	163.60	178.5	3%
new gTLD	24.0	31.6	8%
ccTLD	140.1	160.0	4%
Total	327.70	370.10	4%

Based on the total number of domain names that Verisign reported as active in June 2020 in all TLDs and in the .com/.net gTLDs and depending on the average price per domain name one wishes to use, the total registration fee payments are on the order of $3.7 to $7.4 billion.

Those payments are made to the registrars which are the second cloud from the top. The registrars are responsible for the retail sales of domain names and include companies like GoDaddy and many others that vary significantly in size and may have complex relationships involving outsourcing and reselling activities among registrars. The registrars make payments to ICANN as part of their registrar obligations. The registrars pay the registries in the various gTLDs that they offer second-level domain name registrations in for the right to registrar those domain names.

Below the registrars are the registries that provide the wholesale registry services that enable the registrars to undertake their retail sales and meet their service obligations to the registrants. The registries set the wholesale prices for domain names in their gTLDs. The largest registry is Verisign (by number of domain names registered).[53] A number of registries, like Verisign, provide both the wholesale operations and the back-end registry services that are necessary to make the DNS work, while others outsource their registry back-end functions to other back-end service providers. (Outsourcing appears to be especially common among the operators of new gTLDs.) Some registry operators, like Verisign, also provide back-end services to other registries.[54] As with the registrars, the business relationships among registry operators and across the operation of multiple gTLDs is not always clear. The registry operators constitute a mixed bag of public and private for-profit companies and non-profits. For example, the Public Interest Registry (PIR), which outsources its back-end TLD functions to Afilias, is able to send $67.5 million or 71% of its registry revenues to ISOC.[55]

At the bottom is ICANN which collects registration fees from both the registries and registrars as noted above. The funds it receives are reported on the basis of fund source, but the entities that report making payments to ICANN do so from multiple entities, which makes it challenging to uncover some of the relationships. For example, Verisign's fees paid to ICANN amounted to $45 million (31% if ICANNs' total FY2019 revenues), spread over nineteen different fund sources, while the registrar entities managed by the registrar DropCatch are spread over 1,205 sources that

[53] The largest registry in terms of number of TLDs managed is Donuts.

[54] In addition to providing back-end services for the gTLDs it operates, Verisign provides back-end services for the .edu and .gov gTLDs.

[55] The KPMG (2010) study noted earlier estimated the operating costs of a registry as between $1 to $2 million per year, which is significantly lower than the costs reported by PIR or Verisign, each of which report operating costs that look to be an order(s) of magnitude larger. Additionally, Verisign reports having a headcount of 872 in 2019 (Verisign, 2020a), whereas the average registry in the KPMG study had employment of 25. We suspect there are several reasons for this. First, the size of registries is likely heavily skewed with Verisign the largest and with a long tail of much smaller and less capable registries. Second, the KPMG study focused solely on the registry functions, whereas Verisign and PIR's financials reflect their involvement in other activities.

average $6k per source, but totaled $7.39 million in 2019, making DropCatch the third largest source of funds to ICANN, behind GoDaddy at $10.6 million.[56]

Finally, near the top but on the right is a cloud representing the activity of secondary markets for domain names. These markets allow existing registrants to buy and sell domain names. As will be discussed further below the activity of these is not known since there are so many secondary market players and the transaction volume of buy-sell activity is not reported.

As one can see, the picture of money-flows in Figure 1 is incomplete, but hopefully, it provides a useful starting point for following-the-money, which is an important step toward a better understanding of the economic incentives for different participants in the ecosystem. In the following, we highlight several caveats that ought to be considered when evaluating money flows within the DNS ecosystem.

3.2. Challenges with forecasting the growth in domain names

The key driver we use to size the total revenues for the DNS ecosystem are the number of domain name registrations and the average price per domain name. As we explain further below, domain name prices may vary significantly and be heavily skewed, meaning that the aggregate growth in domain names applied to the average price per registration may prove misleading.

Since 2010, the total number of domain names registered has grown at a CAGR of 6.8%, which is relatively modest compared to other changes in the ecosystem, and significantly slower than the growth experienced earlier in the 2000s.[57] As the Internet has grown to global scale, the slowing of its growth rate might be expected. However, with the expansion in gTLDs the potential for registrants to register in multiple TLDs and for new business models to entice new types of registrations makes it difficult to rely on past growth trends in domain names to forecast where domain registrations and the associated revenues they may garner will go in the future.

For example, before 2012, there were less than 25 gTLDs, but since then over 1,200 new gTLDs have been added.[58] Although domain registrations in the new TLDs have been growing at double digits, they still account for less than 15% of total gTLD domain names (with .com and .net still accounting for over 76%).[59]

---

[56] See https://www.icann.org/en/system/files/files/fy19-funding-source-01nov19-en.pdf. The fourth largest source of funds was another registrar, Donuts, Inc. ($5.3 million) which operates via 198 subsidiaries with the names Binky Moon, LLC.

[57] See OECD (2004) and Note 43 *supra.*

[58] The decision to expand the gTLDs was made in 2010. ICANN received almost 1,930 applications for new gTLDs and by the end of its FY2019 (June 2019), ICANN had 1,222 RAs with registries and had accredited 2,249 registrars. Each of those applicants was required to pay a non-refundable $5,000 ($9.65 million) and approval of an application was expected to incur a total one-time charge of $185,000 ($357 million if all are approved).

[59] Verisign's (the registry for .com and .net) share of total gTLD domain names in 2004 was 85% (OECD, 2004) and 87.2% in 2014 (Tucker et al, 2016). Its share as of 2019 is above 76%, because it is the registry

One reason for expanding the number of new gTLDs was the hope that the expansion will drive price competition and new business models which may tend to drive average domain prices downward, while also expanding registrant choices. Indeed, one of the largest TLDs is the ccTLD for .TK which is associated with the Tokelau Island group off New Zealand. The .TK domain has attracted a lot of attention because it accounts for 7% of the total TLD domain names registered (27.5 million). The .TK TLD operates under a novel business model in which it provides free domain names for one year, relying on its ability to reclaim unused domain names and use them to drive on-line advertising revenues.[60] (The .TK TLD has also attracted attention as a leading source of DNS abuse.[61]) Moreover, the growth and shift to cloud-based services and applications and the growing importance of managing an enterprise's on-line presence will likely help drive demand and competition for DNS services. For example, the market for managed DNS services is expected to grow at a CAGR of over 18%.[62]

3.3. Challenges with forecasting the average price for domain names

Offsetting the potential impact of competition on domain name pricing (and growth) is the fact that historically, price regulations limited the ability of registry operators to raise prices. However, even though wholesale prices were subject to price regulations, the retail prices charged by registrars were not subject to price regulations. Because most registry operators may bundle registration pricing with other services and with a suite of pricing options (e.g., reflecting different levels of customer support, contract durations, volume discounts, etc.),[63] resulting in quality-of-service differentiated registration services, it can be tricky to compare registrar domain name pricing and retail mark-ups to determine whether those may be excessive. Moreover, in addition to registering previously unregistered domain names, many registrars also act as secondary market brokers, offering for resale previously-registered "premium" domain names. The resale prices for premium names are significantly higher than the prices for new registrations in a gTLD. For example, one recent study estimated that upwards of 18% of all domain names registered in .com are held by "domainers" that are holding those names for resale and the average prices of domain names being traded on secondary markets was $1,660 in 2020. Total retail revenue for .com names

---

for additional domains beyond its two largest, .com and .net, which had a 76% share in 2019. To get a better picture of how competition from new gTLDs may be impacting domain name markets, it would also be useful to compute the share of growth in domain names that is being captured by different gTLDs and the registries and registrars that manage them.

[60] In 2013, Freenom launched its free-domain-name business model (see https://www.businesswire.com/news/home/20131216006048/en/Freenom-Closes-3M-Series-Funding#.UxeUGNJDv9s).

[61] For example, the DNS abuse tracking effort, SpamHaus (see https://www.spamhaus.org/statistics/tlds/) identified .TK as one of the "10 most abused Top Level Domains" in April 2020. The .TK TLD is used heavily by email spammers.

[62] See Note 50 *supra*.

[63] Verisign has contractual restrictions on pricing and tying and additional non-discrimination constraints associated with its management of the .com and .net TLDs via its Department of Commerce contract. The operation of other legacy TLDs such as .org, .biz, and .info, as well as the ccTLDs and new gTLDs, do not have those constraints.

in the primary market (new domain names $0.4B plus renewals $1.7B, or $2.3B) is comparable to the revenue generated in secondary market sales ($2.1B), on the basis of much lower sales but at several orders of magnitude higher prices. In the primary retail market, the average price for .com new domain names is ranges from $10-$11, but is $16+ for renewals.[64]

In any case, with the expansion in the number of gTLDs, wholesale price regulations were relaxed, leaving only Verisign's management of its .com and .net gTLDs still subject to price regulations. The further relaxation of price regulations is consistent with the desire to transition to increased reliance on market competition within the DNS ecosystem.[65] Relaxing price regulations on registries for new gTLDs (which are expected to compete for wholesale domain registrations in a much more competitive marketplace than prevailed when the legacy domains were established) is consistent both with allowing the new gTLDs to explore novel business models and giving competitive market forces more scope to operate. However, in the face of increased competition for growth in new domain registrations, legacy registry operators have lobbied to be similarly relieved of their pricing constraints. For example, Verisign, the registry for the largest of the TLDs, .com (with 71% of the domain names in the gTLDs), is price-regulated by the Department of Commerce which set an annual wholesale price of $7.85/domain name since 2012. In 2018, the Department of Commerce permitted Verisign to increase a maximum of 7% in each of the final four years of each six-year period. This provision was incorporated into Verisign's agreement with ICANN. Verisign must continue to provide advance notice of any price increase and registrants can extend their registration for up to 10 years, meaning that for existing .com registrants, if there were a price increase, those registrants could lock in prices with no change until after 2030. (It is worth noting that the special constraints imposed on Verisign via its RA with ICANN and through its contract with the Department of Commerce impose constraints on Verisign's pricing that are not applicable to the registry operators of other smaller or newer TLDs.)

Since renewal rates for legacy and other domains are typically quite high (e.g., Verisign reports .com and .net renewal rates of 74% while GoDaddy reports renewal rates in excess of 85%)[66] and there are significant scale/scope economies associated with operating a registry or registrar (with respect to legacy names), most of any increase in the annual price for a domain name will flow to the registry/registrar's bottom-line.[67] Thus, an increase of 7% to Verisign's fee of $7.85 would add

[64] See BCG (2021).

[65] Although prices for legacy domains were regulated at levels reflected in RAs, ICANN viewed the setting of those price levels as deriving from others, such as the Department of Commerce which regulates prices for .com wholesale domain pricing by Verisign.

[66] See Verisign (2020b) and GoDaddy (2019). Renewal rates are also reportedly high for ccTLDs, with one source reporting that in 2019, renewal rates for European ccTLDs were 82% (see CENTRstats Global TLD Report, Q4 2019 – Edition 30, January 2020, available at https://www.centr.org/statistics-centr/quarterly-reports.html#) and even new gTLDs report high renewal rates. For example, Rightside reported that 81% of domain names that were renewed after one year, were also renewed at the second anniversary (see "Unsurprisingly, re-renewal rates on new TLDs higher than original renewals," Domain Name Wire, October 15, 2016, available at https://domainnamewire.com/2016/10/15/unsurprisingly-re-renewal-rates-new-tlds-higher-original-renewals/).

[67] The annual incremental contribution for domain name registration revenue is quite high, likely in excess of 90%.

$0.55/domain-name per year (a trivial increase from the registrant's perspective if the registrant only has a few domain names), but would add $85 million to Verisign's bottom-line (or 7%).[68,69,70]

Since those higher costs would be paid by registrars in the form of higher wholesale registry pricing, it is possible that registrars might raise their prices.[71] Indeed, because registrars are not subject to price caps and retail prices are already significantly above wholesale prices, it is unclear how much if any of the higher wholesale pricing may be passed through to registrants, and to the extent retail prices are adjusted, they are likely to be adjusted differentially for new domain sales and renewals. For example, Roslyn Layton commented that GoDaddy was able to increase its price 20% from $14.99 in 2018 to $17.99 in 2019,[72] although wholesale prices did not change during that period. That added $237 million or almost 17% to the revenues for GoDaddy's domain business.[73]

These are big numbers for the registry and registrar operators that have large installed bases of domain names. Even if their installed base growth is relatively slow, even modest increases in their pricing could significantly increase their revenues.[74] For example, were the average price for a

[68] Verisign's 2019 revenue was $1,232 million. As of Jun2020, Verisign had 162.1 million domain names registered in .com and .net(which is more than it had at the end of its FY2019 so the estimate here is on the high side). An increase of $0.55 for each of those domain names would add $89 million in revenue (7% of total 2019 revenue) or (assuming a 90% incremental operating margin), $80 million to Verisign's bottom line. Since most of Verisign's costs are fixed, its overall operating margin is much lower. For example, in 2019, it was 65% (see page 25 in Verisign 2020b).

[69] Of course, this calculation ignores the potential that Verisign may need to increase its marketing expenses to attract new registrants with a wider-array of TLDs to choose among or that registries may incur higher costs in other areas (e.g., due to increased cybersecurity threats).

[70] During the public comment period preceding the approval of the new amendment to the Verisign-ICANN RA in March 2020, over 9,000 comments had been received with the vast majority of those arguing against relaxing the price-cap, but many of those having been solicited by registrars or the domainer lobby on behalf of their customers or constituents (see ICANN, 2020; Verisign, 2020c)

[71] Registrars that are already charging above-cost prices may be able to absorb the higher wholesale costs by reducing their margin.

[72] See ICANN (2020).

[73] GoDaddy earned $2,988 million in total revenue in 2019, spread across domain products (45%), hosting and digital presence services (38%), and business applications like email and Office 365 (17%). As of December 2019, GoDaddy had 79 million domain names registered. The price increase of $3 per domain name for 79 million domain names produces additional revenue of $237 million which is 17.6% of the $1,346 million of GoDaddy's total revenue that is attributable to its domain products (see GoDaddy, 2020).

[74] For a registrant seeking to register one or a few domain names, even a price increase of $50 to $100 is unlikely to represent a serious financial burden for most serious registrants, but that would represent a significant incremental contribution to the bottom-line for any register or registrar able to impose such a price increase. On the other hand, registrants do not want to pay more than they need to and may be willing to switch to a new TLD to save money, and such switching is more likely among new registrants than those seeking to renew domain names. While competition is likely to impose significant competitive discipline on the pricing of new registrations, the high renewal rates (due to high registrant switching costs) may

domain name to rise to as much as $50 per year, that would be unlikely to reflect a sufficiently significant cost increase to impact most serious businesses decisions regarding whether or not to register a domain name, but it would add significantly to the bottom-lines of any registrars or registries with millions or tens of millions of domain name registrations.

3.4. Skewed distribution of domain names by TLD

However, the distribution of domain names is highly skewed. Among the legacy TLDs (which includes .com, .net, .org, .edu, .gov, and others), Verisign has in excess of a 76% share as already noted. Among the new gTLDs, the top 10 have 88% of the domain names and only 28 out of the 521 registries of the new gTLDs have more than 50,000 domain names.[75] Determining market shares among registrars is more complicated because many of them are not publicly-traded and they do not report the domains they have registered by TLD. However, the data on domain shares among the new gTLDs shows that the top-10 account for 65% of the total number of domains in the new gTLDs. Moreover, there are only 55 registrars with over 50,000 domain names that account for 95% of the total domains in the new gTLDs.[76] Thus, most of the domain name activity is associated with registries and registrars with relatively large installed bases of legacy names.

3.5. Skewed pricing across TLDs complicates forecasting average price

However, the pricing across these registrars and registries varies substantially. There are some gTLDs that are setting high prices for domain registrations, while others are providing domain names for free (e.g., .TK). Thus, estimating the appropriate average price-per-domain name to apply to the growth rate in total domain names is challenging. (The dispersion in domain name pricing is discussed further below.)

---

provide much less discipline for renewal pricing. One way registrants can address this challenge is by taking advantage of their right to register domain names for ten years under contracts that limit future price increases.

[75] ICANN charges gTLD registries $6,250 per year and then $0.25 per domain name for all domain names in excess of 50,000. Fees for ccTLD registries are voluntary. Of the $137 million in total funds received by ICANN in 2019, $86 million (63%) was from registries. The 28 registries with 50,000 or more domain names accounted for 97% of all of the domain names in the new gTLDs. A large number of the new gTLDs have a single domain name and are associated with gTLDs that are for brands (e.g., .GOOGLE, .TELEFONICA, .BESTBUY, etc.) (As of October 2019, 333 of the 1,201 new gTLDs had a single DNS name registered according to https://ntldstats.com/.)

[76] Registrars often sell domain names in multiple registries, however there must also be quite a few ICANN accredited registrars that only sell in the legacy TLDs. The nTLDstat.com site reported only 521 registries and 371 registrars active in the new gTLDs; however, ICANN reports having RAs with 1,222 registries and RARs with 2,459 registrars as of FY2019 (see nTLDstat.com as of October 2019, and ICANN (2019)). However, many of the 10 largest registrars are also active in legacy domains, like GoDaddy. Other large registrars (as evidenced by the size of their payments to ICANN) include TuCows (Canada), Alibaba (China), Afilias (US), and NameCheap (US) (see ICANN sources of funds for 2019 at https://www.icann.org/resources/pages/governance/financials-en).

3.6. Providing a tentative upper bound for the DNS ecosystem

Nevertheless, were one to assume that the average price were to rise to $50, which seems like a reasonable upper bound, that would imply that the total market of the 370 million TLD domain names (including both gTLDs and ccTLDs) that are active today would result in an estimate of total industry revenues of $18.5 billion. If the growth in domain names is modest, then about $20 billion provides a reasonable upper bound benchmark; however, even if the growth in domain names is much faster, that may provide a reasonable upper bound since it is hard to imagine that the average price would rise above $50 per domain name and many of those domain names propelling the faster growth would likely have significantly lower pricing (perhaps, free).

3.7. Investor Interest in DNS Ecosystem as indicator of value potential

The prospect of earning significant incremental revenue from DNS-ecosystem related investments has attracted interest from the investment community, including private equity, which offers another signal of the value-generating potential of the DNS ecosystem.
Private equity firms have demonstrated an active interest in firms in the domain name space. For example, Siris Capital owns Web.com;[77] Neustar, a large player in the managed DNS services space, is privately held;[78] and GoDaddy was previously privately owned before its $4.5 billion IPO in 2015.[79]

It is also worth highlighting the example of Donuts.com.[80] On the eve of the expansion in domain names, Donuts was identified as one of the Wall Street Journal's Top 50 Start-ups for 2012 and has grown to control a large number of the new gTLDs, although its relative role in the DNS marketplace is obscured by its business structure. Donuts.com is privately held, and prior to 2018, paid its domain fees to ICANN via 198 separate but affiliated entities that if taken together raises Donut.com to the fourth largest source of funds for ICANN. (Furthermore, it is worth noting that, recently, GoDaddy acquired Neustar, and Donuts acquired Afilias.[81])

While private equity has played a large role in the DNS ecosystem for a long time,[82] its role became a major focus of attention within the DNS governance world when the ISOC Board approved plans

---

[77] Web.com is the parent company of Snapname which is another one of the registrars that obscures its significance in the domain name marketplace by spreading its payments to ICANN across 191 sources which together contribute $0.5 million to ICANN – significantly larger than the average source of funds. Snapname is a Florida-based eCommerce/website company that is active in the secondary market for premium domain names.

[78] Neustar acquired Verisign's security business that provided managed DNS services and protection services for DDoS attacks in December 2018.

[79] See https://www.reuters.com/article/us-godaddy-ipo-exclusive/godaddy-ipo-values-company-at-4-5-billion-idUSKBN0MR2S220150401.

[80] See https://domainnamewire.com/2020/05/04/icann-not-opposed-to-private-equity-owned-registries/.

[81] Note 19 *supra*.

[82] Donuts controlled the TLDs licensed to Binky via a complex business structure that makes it hard to uncover what Donuts revenues are.

to transfer ownership of the registry of the .org TLD to a private equity firm in November 2019 for $1.135 billion. With an installed base of 10.1 million domain names, that implies a valuation of $112/domain name.[83] The deal raised a lot of opposition, with a number of folks arguing that transferring .org to a for-profit private equity firm and under an RA that did not constrain price increases was inconsistent with the mission of .org (which had originally been set up to provide service to the non-profit community and although much smaller than .com or .net, is still one of the largest TLDs in terms of registered domain names).[84] Other complaints challenged the way the deal was done, in so far as it was negotiated without the benefit of transparency and involved individuals with prior business relationships with ICANN's management.[85] Some argued that a public auction would have provided a better mechanism, but sponsors of the proposed transaction argued that the rapid timing and unusual process were necessary to take advantage of an attractive opportunity to accelerate the transition towards marketplace economics governing the DNS instead of legacy regulatory structures. ISOC had planned to set up a foundation with the funding from the deal that would expand ISOC's capacity to pursue its mission. In light of the opposition, however, ICANN decided to block the transfer in April 2020.[86,87]

In addition to the significant interest that the history of such deals implies for the value potential of the DNS ecosystem, they also pose a challenge for efforts to trace the money flows. When private equity takes a public registry operator or registrar private, financial information on the performance and business operations of the registry operator are no longer publicly available, rendering the economic implications of their domain name market activity difficult to assess.

---

[83] For ISOC's explanation of the deal in November 2019, see https://connect.internetsociety.org/events/community-webinar-pir-ethos-isoc and for .org domain name count see Verisign (2020a). Additionally, for pointers to opposition comments see http://blogs.harvard.edu/sj/2019/11/23/a-tale-of-icann-and-regulatory-capture-the-dot-org-heist/ and https://www.eff.org/deeplinks/2019/12/we-need-save-org-arbitrary-censorship-halting-private-equity-buy-out. Proponents of the deal were also active (see https://www.keypointsabout.org/).

[84] Although .org was originally set up to serve the non-profit community, there is no requirement that registrants be non-profits. To address the concern that existing registrants might be exposed to excessive price increases, the private equity firm, Ethos Capital, proposed to lock in prices for ten years for existing registrants and committed not to raise prices for eight years.

[85] Additionally, shortly before announcing the proposed transfer of .org, ICANN removed existing pricing restrictions on .org, which raised fears that registrants might be subject to large price increases, further incensing opponents to the deal among the ISOC community.

[86] See https://www.icann.org/news/blog/icann-board-withholds-consent-for-a-change-of-control-of-the-public-interest-registry-pir. Some of the pressure on ICANN came from the California Attorney General. ICANN is a non-profit based in California and the California Attorney General wrote a letter advising ICANN to oppose the deal in April 2020 (https://www.icann.org/en/system/files/correspondence/becerra-to-botterman-marby-15apr20-en.pdf). The fact that an attorney general in California could exert regulatory pressure on ICANN revived long-standing questions about ICANN's independence (see https://www.eff.org/deeplinks/2020/01/after-nonprofits-protest-icann-californias-attorney-general-steps-org-battle).

[87] Since Ethos Capital's effort to acquire .org was blocked, Ethos Capital has become the largest equity holder in Donuts, which also now owns Afilias, which is the back-end provider for .org and hence the recipient of tens of millions of dollars from .org in return for its back-end services (see Alleman, 2021).

3.8. Rise of secondary markets for domain names

Another important phenomenon that makes it challenging to follow the money flows in the domain name market ecosystem is the rise of domain name secondary markets. The participants in this space include registries, registrars and others that, in some cases, have acquired large stockpiles of registered domain names that they hope to re-sell to end-user registrants.[88] The potential for secondary markets arises because of heterogeneity across both the common and private valuation of domain names.[89]

Generally, economists favor robust secondary markets because they help make investments reversible (i.e., the ability to sell an asset reduces the prospect that the original purchase price will be sunk) and secondary markets that provide ancillary services like web-presence brand management can help registrants acquire domain names better suited to their needs than the registrants might select on their own. However, much of the activity in the secondary markets for domain names appears are speculators. A moderate amount of speculator activity can assist in providing market liquidity, but too much speculative activity may increase price uncertainty and might result in upward pressure on the cost to final registrants of obtaining suitable domain names. For example, the opportunity to resell domain names on active secondary markets may induce domain name speculators to buy up attractive character strings, resulting in artificial scarcity for so-called "premium names." Unfortunately, the lack of good data on the transaction activity across the many different secondary markets (which include private transactions between registrants), makes it difficult to assess what the true role(s) of the secondary markets really are.

Earlier we discussed how brokers and secondary markets have created a mechanism for buying and selling domain names. While a price increase of even up to $50 per year may seem like a small relative cost for a legitimate business with a single domain name,[90] the same cannot be true for domain name speculators sitting on inventories of domain names that may number in the 100s of thousands. For example, GoDaddy.com reports that it had an inventory of premium domain names with 750,000 names. Assuming a registration fee of $10 to $20 per year, the carrying cost for that

[88] The operators of new gTLD registries are not restricted from operating registrars, and may engage in complex dynamic business strategies in efforts to maximize the joint value of their registry and registrar businesses. Those may range from trying to sustain high prices for an exclusive gTLD reputation, low prices for a discount gTLD reputation, or some mix in between that may change over time.

[89] The common value refers to the valuation that would be placed on the domain name by a competitive market, which may differ from the private valuation of actual or potential bidders. Different domain names (i.e., different character strings) have different common values. A secondary market can enhance economic efficiency by facilitating the matching of buyers and sellers that can jointly realize gains-from-trade by exchanging ownership of a domain name for which their private value differs (e.g., an enterprise that is exiting a market may find it desirable to sell its domain name, or online "brand" to a prospective entrant). Also, secondary markets can assist in search so that buyers can better identify a domain name that better matches their private value/cost trade-off calculus (e.g., opting for JoesGarageMA.com which may be less expensive to acquire than JoesGarage.com).

[90] For example, a one-person business is likely to have a telephone bill north of $100 per month, so $50 per year equates to less than $5 per month. Although no one wants to spend more than they have to, it is hard to imagine any potential average price rise in domain registration fees to pose a serious impediment for most businesses seeking to manage their on-line presence.

inventory is $7.5 million to $15 million. [91] The prices for premium domain names vary significantly but many are priced above $1,000. If GoDaddy could sell on average 1-2% of its inventory at that price, that would be sufficient to cover the registration fees for the inventory. These "guesstimates" suggest that the secondary markets may be sustainable with relatively low transaction volumes.

Nichols (2013) analyzed the domain name markets and concluded that the potential character spaces of most TLDs are mostly empty – that is, most TLDs (other than a few) have most character strings available (at least that was the case in 2013, before the growth in secondary market speculation). Nichols defined the set of desirable domain names as being comprised of the population of domain names that are registered in at least one TLD and found that only 0.523% were registered in all domains, which means a registrant who was willing to go to another TLD could register that "premium" domain name there. Indeed, the only TLD where the set of desirable domain names might be deemed to be scarce is .com, where over 90% of the identified premium names were already registered. Thus, if one accepts Nichols analysis, it is unreasonable to conclude that the inability to get a desired character string requires further expansion in the TLDs than had already occurred by 2010 (i.e., if not available in .com, it probably is in .net, .BIZ, or .US).

If expanding the number of TLDs is not needed to address fundamental scarcity in the availability of meaningful character strings, then what is the effect? One concern is that it would drive the need for significant defensive registrations by legitimate registrants seeking to protect their brand image by protecting against their domain being registered in another TLD. This risk raised another significant concern about the potential role of secondary markets and domain name speculators in pursuing hold-up strategies to extract surplus from legitimate registrants. To address this concern, ICANN helped put in place process rules to protect domain name registrants' trademarks by requiring that new registrations pass a waiting period during which the registrations may be challenged by existing registrants and has put in a series of protections to allow new registrations to be screened for trademark violations.[92]

---

[91] GoDaddy claims to offer "one of the world's largest domain after-markets" and "over the last five years… acquired more than 750,000 domain names." GoDaddy operates a "cross-registrar network that automates transaction execution across registrars" and receives "a percentage of the sales price for each domain sold." (See page 10, GoDaddy, 2019). Namecheap is another registrar that also sells premium domain names and runs an aftermarket for buying and selling registered domain names (see https://www.namecheap.com/domains/marketplace/buy-domains/). As of November 2020, Namecheap was the third-largest registrar in the new gTLDs with 3.1 million domain names registered in the new gTLDs (behind Alibaba and GoDaddy, according to nTLDStats.com). Even a relatively small increase in wholesale pricing could substantially damage Namecheap's after-market business so it is not surprising that they have been actively opposed to relaxing price-caps on registry wholesale pricing (see https://www.namecheap.com/blog/icann-allows-com-price-increases-gets-more-money/ for a February 2020 blog post).

[92] As part of the RA, ICANN charges registries a $5,000 one-time fee to become members of the Trademark Clearinghouse and a $0.25/per transaction fee to check new registrations against existing trademarks. Trademark Clearinghouse is a dbase of validated and registered trademarks created by ICANN (see https://www.trademark-clearinghouse.com/). It is essential part of gTLD program. Design originally proposed in 2009 and ICANN published draft mandatory RPMs April 2013 (see http://newgtlds.icann.org/en/about/trademark-clearinghouse).

Although studies of trademark infringement costs have alleged significant risk, empirical studies suggest the actual costs are likely modest.[93] This is not surprising since as noted before the costs of multiple registrations are unlikely to be cost-prohibitive for most businesses and the data suggests that even if initially many .com registrants did engage in defensive registrations back in the early 2000s, many of those were not renewed and there does not appear to have been a stampede to register in the new gTLDs.[94] Thus, further study of the costs and benefits that expanded choice and competition delivers as a consequence of the large expansion in the number of gTLDs is warranted.

3.9. Assessing the magnitude of economic harms associated with threat to the DNS

Finally, while the above analysis suggests that the entire DNS ecosystem represents a relatively small part of the global Internet ecosystem (maybe $6 billion today with the potential of growing to $10 billion in a few years), it is important to remember that the DNS is a core component of the infrastructure that keeps the Internet running. Since the Internet is at the heart of the Digital Economy transformation, anything that seriously threatens the Internet poses a threat with the potential for global, socio-economic harms.[95] The RAs impose obligations on the registries to work to secure the DNS and the fact that operators like Verisign and others have managed to support the DNS with highly reliable service as the Internet has scaled in size exponentially (e.g., supporting over 100 billion daily DNS queries) is a testament to the success of the DNS.[96]

3.10. The effect of DNS Policy Changes on Money Flows

Until 2010, there were less than twenty TLDs.[97] In the 1980s, the DNS was managed by SRI, a non-profit under contract to the U.S. Department of Defense, and during those early days, the DNS

[93] See Katz, Rosston, & Sullivan (2010); Krueger & Van Couvering (2010); and Nichols (2013).

[94] Although the costs of acquiring a new gTLD (at $185,000) is sizable even for large companies, a number of well-known brands did invest in acquiring gTLDs. However, as noted earlier, most of these new brand gTLDs have yet to be utilized. The expectation is that companies may have acquired these gTLDs as options for the future that will give them flexibility in how they may wish to manage their on-line presence. As we discuss further below, while it is difficult to move from xxx.TLD1 to xxx.TLD2 (large switching costs), the owner of TLD xxx does not need to move.

[95] See Lehr, Clark, Bauer, Berger, and Richter (2019) and Clark & Claffy (2015).

[96] Verisign operates two of the thirteen Internet root servers that are the authoritative sources for the global Internet root zone and critical components of the DNS infrastructure. See "DNS Outages: the Challenges of Operating Critical Infrastructure," Verisign Blog, April 15, 2014, available at https://blog.verisign.com/security/dns-outages-the-challenges-of-operating-critical-infrastructure/.

[97] As of 2010, there were 21 gTLDs. See page 8 of Katz, Rosston & Sullivan (2010) for a list a chart of the 21 gTLDs that existed as of 2010. The original seven gTLDs created in 1985 were .ARPA, .com, .edu, .gov, .MIL, .net, and .org; in 1988, .INT was added; and in 2001-2002 seven more gTLDs were added (.AERO, .BIZ, .COOP, .INFO, .MUSEUM, .NAME, and .PRO); and from 2005-2007, six more gTLDs were added (.CAT, .JOBS, .MOBI, .TRAVEL, .TEL, .ASIA).

was virtually under the sole control of Dr. Jon Postel.[98] In the 1990s, the NSF had assumed management responsibility for the Internet and contracted with Network Solutions, Inc. (NSI) to manage the DNS. In 1998, under President Clinton's administration, the U.S. National Telecommunications and Information Administration (NTIA) issued a statement setting forth plans to privatize administration of the DNS and announcing the creation of ICANN, a California-based non-profit that would take over management responsibility for the DNS under a Memorandum of Understanding (MOU) with the U.S. Department of Commerce with a key mandate to create competition in the domain name market.[99]

Management of the gTLDs represented an exclusive right to sell domains registered in their TLDs. As a first step toward expanding competition in the domain name marketplace, ICANN took steps to separate the registrar functions (retail sale of domain names) from the registry functions (management of the gTLDs). Until relatively recently, the RAs limited registry operators, who retain an exclusive right to sell domains registered in their TLDs, from having more than a 15% ownership interest in any registrar and from discriminating in the services provided to unaffiliated registrars. Additionally, registries were subject to price-caps that limited the wholesale prices they could charge registrars. Following the introduction of vertical separation, retail registration "prices went from $35/yr to about $10/yr,"[100] however it would be incorrect to attribute the price drop solely to the decision to require vertical separation.

Although retail prices for domain names fell precipitously as a consequence of introducing retail-level competition among registrars, the wholesale registry market remained highly concentrated. Verisign, which acquired NSI in 2002[101] and replaced NSI as the registry for the .com and .net gTLDs was the registry for over 86% of the domain names in the early 2000s.[102]

Today, although there are now over 1,200 gTLDs, the wholesale market for gTLDs remains highly concentrated with Verisign's .com and .net gTLDs still accounting for 77% of all gTLD domain

[98] For notes on the early history of DNS, see https://www.cybertelecom.org/dns/history.htm, https://cyber.harvard.edu/icann/pressingissues2000/briefingbook/dnshistory.html, and https://www.internetsociety.org/internet/history-internet/brief-history-internet/.

[99] See https://www.ntia.doc.gov/page/1998/memorandum-understanding-between-us-department-commerce-and-internet-corporation-assigned-; and for policy statement by NTIA on the management of domain names see, https://www.ntia.doc.gov/federal-register-notice/1998/statement-policy-management-internet-names-and-addresses. The latter articulated "four principles to guide the evolution of the domain name system: stability, competition, private bottom-up coordination, and representation," which were further elaborated on in the NTIA's "green paper" which was published in the Federal Register (see https://www.govinfo.gov/content/pkg/FR-1998-06-10/html/98-15392.htm).

[100] See http://www.circleid.com/posts/20081204_reexamining_domain_registry_registrar/, and CRAI (2008) report that it references. CRAI (2008) concluded (p2) "ICANN's policy of fostering registrar competition has been extraordinarily successful. ICANN estimates that registrar competition reduced gTLD domain name registration fees by 80%, saving registrants more than $1 billion annually."

[101] Verisign announced plans to acquire NSI for $21 billion in 2000 (see https://money.cnn.com/2000/03/07/deals/verisign/, and the deal closed in 2002, https://www.cnet.com/news/verisign-buys-network-solutions-in-21-billion-deal/).

[102] See Figure 2, OECD (2004).

registrations, or 44% of all TLD registrations (after including ccTLD registrations).[103] However, a more meaningful metric might be to focus on the shares of new domain names since that is where the competition is most intense. By this metric, .com and .net accounted for 43% of new gTLD adds, or 39% of total TLD adds.[104]

The lack of comprehensive data on registrar pricing makes it difficult to assess what the distribution of actual offered or adopted pricing is. Registrars offer domain names under a wide variety of terms and conditions, quite typically with significant discounts for longer-period registration agreements which can range from one to ten years (with ten years set as the maximum by ICANN standards). And, as already noted, in many of the other TLDs (like the ccTLD .TK and many of the newer gTLDs), registrar offers for low-priced domain names are common. Those low retail prices are enabled by similarly low wholesale pricing by some of the new gTLDs as well as by lower mark-ups by many of the registrars selling domain registrations in those new gTLDs. The low prices may be justified from a business perspective as either penetration pricing to establish a brand image and build an installed base of legacy registrants or as loss-leaders to enable the registrars to sell additional complementary services like web-hosting or on-line brand-management services. Since much of the costs of establishing a new registry and of operating a registrar are fixed (and perhaps sunk),[105] such pricing strategies make business sense.

However, at the same time that the expansion in the gTLDs that occurred after 2010 enabled the introduction of new business models with lower priced domain name options, the dispersion in domain name pricing has increased at both the wholesale and retail levels. Many of the new gTLDs have wholesale prices that are significantly higher than the prices set by many of the legacy gTLDs, and in part due to this, but also due to differences in business models, many of the registrars also have retail pricing for domain names that are significantly higher. For example, Tucker & Rafert (2016) surveyed wholesale prices and retail prices for domain names by gTLD in 2015 and 2016. They found significant dispersion in wholesale and retail pricing (tracked as mark-ups). Whereas wholesale prices for gTLDs were relatively flat or trended up over the two samples, retail mark-

[103] See Verisign (2020a). At the end of June 2020, Verisign reported there were a total of 370.1 million registrations, consisting of 148.7 million (.com), 13.4 million (.net), 16.4 million (other legacy gTLDs), 31.6 million (new gTLDs), and 160.0 million (ccTLDs).

[104] As of June 2020, Verisign reported that new registrations had increased 15.3 million (4.3% growth) year-of-year, with 6.0 million new registrations in .com plus .net, 8.6 million new gTLD registrations (37.4% growth), and ccTLDs 1.4 million (0.9% growth) (see Verisign, 2020a).

[105] For example, acquiring the authorization to establish a new registry incurs a one-time application fee payable to ICANN of $185,000. This may not be sunk, however, since the rights to operate the registry may be sold to another registry. Additionally, much of the technical infrastructure to establish a registry (e.g., the back-end operations) needs to be sized to handle peak transaction loads and so a significant share of the costs is likely fixed based on the capacity of the system, and adjusting the capacity at the margin is likely relatively low cost given the dynamic availability of cloud infrastructure in today's global Internet ecosystem at low incremental costs. However, since many registries outsource their backend operations to backend providers like CentralNIC, Donuts, Neustar, Afilias, or others (see https://ntldstats.com/backend) and the market for such services appears to be quite competitive, it is likely that such services are available under highly scalable and competitive pricing terms. Thus, the entry costs for becoming a new gTLD registry operator appear to be relatively low, and while non-trivial, do not seem likely to preclude entry by a large number of enterprises with relatively modest business resources.

ups appeared to have declined significantly over the two sample. Also, they found domain name prices for gTLDs ranged from $5.00 to $190.00 across legacy and new gTLDs, but the highest wholesale price observed for a legacy gTLD was $80.00 in 2016.[106] They found retail mark-ups for legacy gTLDs (median, min, max) were 76%, -2%, 170%; and for new gTLDs (median, min, max) were 74%, -44%, and 186%. These price data demonstrate the brand/business model differentiation that exists across the gTLD domain name markets, and also the significant fluctuations that may occur over time due to changes in the mix of gTLDs and domain name activity across registrars and to registrars changing their pricing and marketing strategies. Although not conclusive, these results are consistent with vigorous competition in the domain name marketplace. However, the constraints Tucker & Rafert faced in collecting their sample data demonstrates that the lack of comprehensive data on the pricing practices of registries or registrars makes it hard to fully assess price trends.[107]

3.11. ICANN's Efforts to Increase Competition and Transition to Market Forces

One reason why wholesale price dispersion is higher is because the new gTLDs are not subject to price-cap regulations. And, ICANN has moved to relax price controls on legacy gTLDs as well.[108] From the perspective of expanding marketplace competition, the move to get ICANN out of price regulation makes sense, and indeed, a number of researchers have long argued that ICANN is ill-constituted to undertake that responsibility.[109] ICANN lacks both the resources and mandate to regulate prices in the much more dynamic DNS marketplace that has been created with the expansion of gTLDs from just over twenty to over 1,200 today, and prospects for even more gTLDs in the future.

From an economic perspective, ICANN's prior policy for increasing the number of gTLDs only after significant review and consideration imposed a regulatory-induced artificial scarcity on the DNS ecosystem. As long as there was no underlying technical reason why the number of gTLDs

[106] For additional data on wholesale and retail pricing, see OECD (2004).

[107] Although Tucker & Rafert's 2015 and 2016 studies offer the best publicly available analyses of pricing trends we are aware of, they had to rely on voluntarily self-reported registrar pricing data and a review of posted pricing scraped from websites. They had to reply on interpreting offered prices in the absence of actual sales data, from registrars with multiple pricing offers, and offers which often differed with respect to non-price terms – all of which can contribute to potentially large measurement errors.

[108] As already noted, Verisign's new amendment to the .com RA resumed the relaxation in the price-caps that had been permitted until 2012 essentially restoring the ability for Verisign to increase wholesale .com prices in the final four of each six year period up to 7% per year and in March 2021 Verisign announced plans to raise prices from the current price of $7.85/year by 7% on September 1, 2021. From about 1999 until 2006, the price for a .com domain name was $6.00. In 2006, the Department of Commerce (DoC) permitted pricing flexibility of up to 7% in 4 of 6 years (tied to the term of the registry agreement with ICANN). In 2012, the DoC withdrew that flexibility. In 2018, the DoC restored it once again

[109] Milton Mueller, one of economists who has been most closely engaged with analyzing the challenges of governing the DNS ecosystem has long argued persuasively that ICANN's functions should focus on the technical requirements of managing the DNS and should refrain from straying into industrial policy or economic regulation of the Internet. See, for example, Mueller (1998, 2002a, 2002b), Mueller, Mathiason, and Mcknight (2004), Mueller and McKnight (2004), Mueller (2006), and Mueller and Badliei (2017).

could not be increased, the artificial scarcity appeared to impose an arbitrary constraint on competition and its potential to deliver the benefits of both lower pricing and innovations that might more closely match end-users' demands for expanded choice in domain registration options. Although there was some concern about the ability of the DNS to scale technically to an order-of-magnitude larger number of gTLDs, the experience since 2010 demonstrates that increasing the number of gTLDs poses no significant technical nor cost challenges (from the perspective of putting in place the necessary DNS back-end infrastructure). This suggests that further expansion in the number of gTLDs is feasible without causing significant technical capacity problems.

In addition to the technical concerns, there were concerns among existing registrants of the risk that secondary markets might pose for pre-emptive registrations by speculators seeking to capture premium names for future resale and the need by incumbent registrants to engage in defensive registrations in new gTLDs to forestall those being acquired by others and used in ways that may damage the registrant's brand (either because those alternate registrations might result in confusion for customers or, worse, might be used maliciously). The first problem has already been touched on and will be discussed further below.[110] The second problem induced ICANN to adopt procedures for allowing registrants to challenge new registrations that may adversely impact a registrant's trademark interest in its registered domain name if a similar domain name were to be activated in another gTLD.[111]

However, as the above discussion makes clear, the analysis of the economic need for expanding the gTLDs and relaxing the prior vertical integration and price-cap restrictions that existed before the gTLDs were significantly expanded in 2010 remains far from adequate or complete. In 2009, ICANN, as a requirement of the process put in place to evaluate ICANN's governance practices, commissioned a study by the well-known economist, Dennis Carlton, that analyzed the basic economic case for expanding the gTLDs. Professor Carlton concluded that expanding the gTLDs was a good idea.[112] However, Professor Carlton's analysis was based on qualitative arguments since the data to conduct a substantive empirical analysis of expanding the gTLDs was not available (even though the number of gTLDs had been growing since the original seven gTLDs in 1985 to the 21 gTLDs that existed by 2010). Significant criticisms from many in the technical community challenged the analysis as overly superficial, especially with respect to its failure to consider the potential costs and threats to the stability and operation of the DNS that may arise as a by-product of the increased risk for DNS abuse if the number of gTLDs was substantially expanded.[113] A more extensive study by Katz, Rosston, and Sullivan (2010) highlighted the need to consider spillover effects on third parties and noted that the available empirical data did not provide a sufficient basis to conclude that the net economic benefits overall from expanding the

[110] See the discussion of domain name retail pricing in section 3.3 and secondary markets in section 3.8.

[111] See discussion in section and 2.3 and footnote 32 regarding the need to protect existing registrants' trademark interest in prior registrations.

[112] See Carlton (2009).

[113] For example, see Kende (2009). One of the economic justifications for expanding the gTLDs was to offer viable competition for .com, however, as Kende points out, the limited success of .info and .biz in capturing growth in new registrations suggests that new gTLDs may not be close substitutes for .com and to the extent that is true, expanding gTLDs would exert less competitive discipline on .com.

gTLDs would be positive, and that further research was needed. The most recent ICANN-sponsored studies by Tucker and Rafert (2015, 2016) which were undertaken as part of ICANN's mandatory review process did provide a partial empirical assessment of the competitive impacts of the post-2010 expansion in gTLDs, and those studies provided provisional evidence that legacy gTLD market shares and retail price margins have declined.[114] However, those studies also reiterated the partial nature of the conclusions that could be drawn based on the continuing lack of adequate data on domain marketplace dynamics, and especially with respect to the pricing and operating behavior of registrars and registries.

Thus, although the number of gTLDs has been substantially expanded, there has never been a comprehensive evidence-based review of the economic cost-benefits of expanding the gTLDs. Even if one concludes that there is more competition in both the wholesale and retail markets for domain names, there is no good empirical basis for concluding that the Internet ecosystem and domain name registrants are better off in the new world.[115]

4. Changing DNS Ecosystem and new Challenges

In the following sub-sections, we examine three issues of relevance to the future of the DNS and offer some speculations as to what they may mean. First, we discuss the relevance of market power concerns for the DNS ecosystem. Second, we consider the implications of changing technology and market trends for the DNS. And, third, we address the threat of DNS abuse for the security and stability of the DNS.

4.1. Market power

As noted above, Verisign and the .com registry remans the largest TLD with a unique position within the DNS marketplace; and across the realms of legacy and new gTLDs, we see a wide array of wholesale and retail pricing behaviors. What does this tell us about the existence of market power or the policy concerns that may raise?

One goal of promoting competition is to limit the threat that market power might be accumulated and abused. Economists have long-recognized that focusing on market shares alone is not an adequate basis for assessing market power and antitrust policy accepts that the position of market power is not illegal but only its abuse. If entry barriers are sufficiently low, then the potential of

[114] See earlier discussion around and included in Note 106 *supra*.

[115] The concerns of many across the DNS stakeholder community with the lack of adequate economic analysis of cost-benefit case for expanding the number of gTLDs has been amply documented across multiple reports prepared by ICANN stakeholders. The ICANN bylaws require ICANN to establish a process for periodic reviews and the reports of those sundry reviews are instructive. See for example, CCT (2018), which documents the long-standing history of repeated calls for more extensive analysis and data collection to support the evaluation of the economic impacts of expanding the gTLDs. One oft-repeated call has been for collecting more extensive data on registry and registrar business practices – a call which has been opposed by registry and registrar operators. The tussle over what data ought to be collected and how it should be analyzed or used is complex and its full consideration is beyond the scope of this paper. Suffice it to say that our current empirical understanding of the economics of the DNS ecosystem is inadequate and the lack of necessary data is part of the problem.

new entry may offer sufficient competitive discipline to limit the ability of incumbent firms from exploiting their market power no matter how large their market share.

A firm with market power may abuse its market power in two classic ways: first by setting prices higher than would be the case in the face of competition; and second, by engaging in strategies designed to raise the costs of potential rivals (and thereby protect the firm's monopoly pricing). To assess the first threat, antitrust economists typically begin by defining the relevant market for which prices are thought to be too high. That is far from simple in situations where the products on offer differ along multiple dimensions. In the domain name space, one of the key dimensions of differentiation relates to the relative value of different character-string combinations as already noted. Another dimension that is important are the terms under which different domain name registration offers may be made. Those terms include a complex array of price and non-price related terms and conditions, including the duration over which the registration is to be provided, the pricing structure (one-time, fixed, and variable pricing elements), the prospects for renewal (and the associated pricing), the reputation (or brand) of the registrar and registry, and the options for purchasing the domain names (e.g., are bulk registrations offered? How much before/after-registration customer support is provided, et cetera). Across different registries and registrars, these terms may differ widely.

Thus, observing that the prices for domain names may differ does not allow one to conclude that those prices reflect any abuse of market power. Indeed, observing that the average or median wholesale or retail prices for domain names has increased does not imply that registrants are necessarily worse off or that market power problems exist. Without further analysis (some of which we address further below), one cannot conclude that the higher pricing may not be justified by higher costs associated with enhancements in quality (e.g., more premium domain names either due to their intrinsic characteristics or bundling with valuable other terms and conditions)[116] or new sources of costs (e.g., increased cyber-security risks).[117]

What is true about individual domain names, is also true about domains registered in particular TLDs (e.g., .com) which may be sufficiently popular that the registry can charge a premium price for registration. Providers of higher-quality products have always been able to charge a premium price and the existence of such options does not immediately justify a concern that consumers are suffering any abuse.

[116] The advent of on-line book sales appeared initially to confound economic predictions that expanded options for purchasing books would drive market forces to assert the "law of one price" and narrow the price dispersion among book pricing. That did not happen, but on closer inspection, part of the issue was that inappropriate price comparisons failed to adequately account for difference in the terms and conditions associated with different bricks-and-mortar versus on-line book sales offers.

[117] Generally, given the trajectory in information technology costs (for networking, computing, and storage), the costs of supporting registry and registrar functions (if not the marketing and other customer-related costs) are likely decreasing over time. The need to respond to changes in the cyber-security landscape and other exogenous shifts in the Internet, however, may be sources for rising costs that could justify higher domain name registration pricing.

However, once a registrant has committed to a domain name, a different set of considerations apply. Once a registrant has committed to a domain name, and embedded that name in their online presence, they may face substantial switching costs if they were to move to a new domain name. That creates an opportunity for the registrar, and by proxy, the registry, to seek to hold-up existing registrants for additional revenues when it comes time to renew the registration. Although this issue of potential market power abuse is sometimes framed as a distinction between legacy TLDs and new TLDs, and by extension to the registries that operate the different sorts of TLDs, it applies with respect to all existing registrants that have significant switching costs associated with moving to another domain name or registry. Indeed, the threat of hold-up exists regardless of the TLD at issue, but the incentives to take advantage of registrant switching costs will hinge on the balance between a registry's interest in future growth (i.e., ability to attract new registrants to the TLD) versus the potential to harvest value from the registry's installed base.

Verisign, as the operator of .com which accounts for the largest number of domain names, is sometimes singled out as requiring special treatment because it hosts those long-standing TLDs. But if the new TLDs are successful in attracting domain registrations that persist, the same issue will arise there as those names come up for renewal. It would seem that the correct approach to deal with the concern with "renewal hold-up" would be to protect the process of renewal from abuse, not single out one or another registry for special treatment.

One regulatory response that would somewhat protect domain name owners would be to impose a rule that the price for a renewal registration cannot be higher than the current price for an initial registration. If we believe that the market for new registrations is competitive, and the presence of many choices of TLDs where any desired character-string might be registered and competition among registrars will drive the price of new registrations down, then this rule would tend to protect the holders of long-standing domains from abusive pricing.[118] Note, with the relaxation of the restriction against vertical integration between registries and registrars for most TLDs,[119] targeting such a price rule may require applying it both to registries (via the RAs) and to registrars (via their accreditation agreements with ICANN).[120] Also, figuring out the best approach to enforce such

[118] Many registrars compete for new customers by offering attractive up-front or initial period discounts or by bundling with other attractive features or services. As already noted, such strategies can reflect sound business and pro-competitive strategies, but such strategies may also be used to bypass price regulations. For example, a registrar (and the associated registry) may offer sufficient discounts for initial registrations and then a uniform (but much higher) renewal rate for initial customers that would effectively price discriminate between new and legacy customers. Some sort of maturity-weighted price-cap might be needed to protect against differential pricing for new and legacy registrants and by the burden of this constraint on registries/registrars with different mixes of new and legacy registrations (as evidenced by the maturity profiles). For example, a registry/registrant that was willing to forgo future growth, would be better able to exploit its ability to engage in hold-up pricing.

[119] The RA under which Verisign operates the .com TLD retains the restriction against vertical integration.

[120] For example, if applied only at the wholesale level, the affiliated registrar may be able to undo the pricing constraint via its retail registration terms. However, if registries remain subject to non-discrimination rules in their dealings with affiliated and unaffiliated registrars, then inter-registrar competition may limit the effectiveness of such regulatory bypass strategies. The well-known challenges of relying on wholesale price regulation to discipline market power when the wholesale firm is able to compete in downstream markets highlights the complexities that enforcing any such pricing regulation will confront.

any pricing rules will prove challenging since ICANN is ill-suited to be a price regulator as already noted.[121]

Returning to the question of the benefits of expanding the number of gTLDs, it is useful to consider the rationales that were offered. One is to alleviate a presumed scarcity. Whether or not there is actually a scarcity in names is debatable.[122] The second is to allow innovative experimentation in the value of new TLDs as part of branding. We discuss the possible erosion of this benefit in the next section. The third reason is to create competition across registries to create further downward pressure on domain name pricing (not already induced by retail competition among registrars). However, if competition among registries is the goal, what we must track is not the number of new TLDs but the number of firms offering those names—the number of registry operators. If we see consolidation in the registry marketplace (no matter how many TLDs are being offered), this might signal a reduction in competitive pressure, as well as a signal of cost pressures on the less successful firms that leads to mergers.[123] Operating a registry is an activity with mostly fixed costs. Systems do need to scale with the number of names hosted and the rate of queries, but many of the expenses are unrelated to those factors. This fact would suggest that over time firms will consolidate to share the fixed costs across more TLDs, and competitive pressures may diminish. Tracking the industry structure is an important part of assessing the effectiveness of competitive pressure, but it turns out to be difficult to track down the details of the industry structure behind the registries since not all of them are publicly-traded and a number of the registries that provide their own back-end services also provide back-end services to other registries. Moreover, as already noted, the relaxation of the vertical separation restriction between registries and registrars and similar consolidation at the registrar level make it difficult to trace the business relationships across registries, registrars, and registry-registrar functions. The registrar retail-level costs represent a mix of fixed and variable costs. The establishment of national/global brand are largely fixed, but customer acquisition and support costs are variable.

Although there are valid concerns that the market structure may tend toward enabling increased market power at the registry and/or registrar level, it is also important to consider what the potential for harm is from increases in domain name registration pricing. As noted earlier, when ICANN approved the new amendment to the .com RA with Verisign for the .com gTLD in 2020, the price-cap that previously limited Verisign's pricing flexibility over .com was relaxed.[124] However, under

[121] See Note 109 *supra*.

[122] See the earlier discussion and Nichols (2013), where he argues that the scarcity of premium character strings was already addressed by the expansion in gTLDs that occurred before 2010.

[123] On the other hand, a well-functioning oligopoly might be vigorously competitive. Cellular markets have become more concentrated over time in most countries, and much of that consolidation was to expand coverage and to realize scale and scope economies. Whether it has been beneficial or harmful to competition continues to attract significant disagreement among industry analysts.

[124] The price constraints imposed on Verisign related to the operation of the .com gTLD are complicated. Verisign's operation of the .com gTLD is also governed by an agreement with the U.S. Department of Commerce (DOC) which limits the percentage wholesale price increases Verisign is allowed to impose and track the limits included in the ICANN RA. In addition, Verisign is required to provide advance notice of any planned price increases and to allow registrants the option of extending their current registration

any reasonable scenario for the future trajectory of Verisign pricing in which Verisign sets higher domain name prices, it seems very unlikely that businesses with the need for one or a few domain names would incur cost increases that would be sufficiently large to significantly alter their business plans. However, even if the ultimate registrants' business plans might not be impacted, the potential ability of DNS registry/registrar operators to extract significant surplus (in aggregate) from registrants raises equity concerns. Absent a justification founded on either an increase in the quality of service being provided or the costs of providing services,[125] it may strike many as unfair to allow DNS service providers to benefit from the potential windfall of being able to extract rents associated from incumbent registrants that are prevented from abandoning or switching their domain names due to large switching costs.

This might seem especially unfair with respect to the legacy domains that were granted the exclusive franchises without having been required to pay for those franchises. The owners of new gTLDs that had to pay at least $185,000 for the right to operate a new gTLD and are entering a marketplace crowded with two orders of magnitude more gTLDs than existed when the legacy gTLDs were building their installed bases may have a slightly better argument for being relieved of pricing constraints that might prevent them from maximizing the value of their investment. On the other hand, the legacy registries had to bear the burden of building much of the back-end DNS infrastructure that registries of new gTLDs benefit from. In any case, however, the basis for setting the franchise price by ICANN is unclear (i.e., how it relates to the costs ICANN incurs in approving a new gTLD or to some reasonable proxy for a market-based valuation is not documented). As long as ICANN retains control over the supply of gTLDs, some have argued in favor of adopting a mechanical process for expanding the number of gTLDs to render the future supply more predictable and then relying on an auction process to assign those gTLDs to registries.[126] This would mirror somewhat the trajectory of reforms in radio-frequency (RF) spectrum management, by which regulators have moved from "beauty-contest" assignments of RF licenses to auction-based assignments of tradable RF licenses.[127]

---

licenses for an additional ten years at the current price. These rules are intended to alleviate concerns that Verisign might seek to exploit its enhanced pricing freedom to impose excessive prices on incumbent registrants.

[125] Increased demand for reliability in the DNS in the face of growing cybersecurity threats and increased strategic complexity due to the changing dynamics of DNS markets may provide a cost-based justification for increasing registration pricing. Ensuring that the DNS backend services continue to operate reliably is a significant challenge in the face of DDoS attacks that can drive high-levels of peak traffic. In spite of the growth of such problems, it is worth noting that Verisign has ensured robust and uninterrupted DNS resolution services for over two decades. Moreover, registries differ in the quality of their infrastructure and business operations and higher quality may come with higher costs that may contribute to pricing differences. Furthermore, it is worth noting that Verisign's wholesale pricing is significantly lower than many other gTLDs and is below levels that Verisign could set under its current contracts with the Department of Commerce and ICANN.

[126] See Manheim & Solum (2002) and Mueller & McKnight (2004).

[127] The regulatory debates over clearing broadcasters and satellite users (which received their RF usage rights for free) from spectrum in order to expand spectrum access to mobile network operators (which increasingly have moved to operating under spectrum licenses purchased via spectrum auctions) has similarities to the debates justifying differential regulatory treatment for legacy and new gTLD users.

In the following sections, we raise two issues that might affect the pricing of domain names. In the next section we raise the question of whether, especially in the future, domain names will be effective as elements of branding. Then, in the subsequent section on security, we discuss how the changing landscape of cybersecurity concerns may impact the DNS ecosystem. Meeting those challenges may raise costs for registries and registrars that choose different approaches to address those challenges. If this is the case, the conception of "quality" domain names may change, and come to be associated with registries and registrars that undertake operational practices to better secure their security and reputation.

4.2. Changing role of DNS

In the early days of the DNS, the designers conceived of it as having a very simple purpose. IP addresses were hard to remember and type correctly, and having a name for a machine that would remain constant if the address changed was useful. So, the original purpose was to translate a name to an address, which seemed at the time like a very simple and uncontentious service. There was, however, a poorly documented debate among the designers as to the structure of the names. A minority of the designers argued that the names should be meaningless—perhaps strings of numbers. If the purpose of the DNS was just to provide a name that did not change when the address changed, why was it necessary that the names express meaning? To others who were doing the design, it was obvious that the parts of the domain names should capture some meaning, and considerable effort went into picking the initial TLDs: .com, .edu, .org, .net and so on.

Once the decision was made that the elements of a domain name could be chosen to have meaning, issues of branding and trademark instantly arose, and to many people the most important aspect of the DNS was its possible potential for branding and differentiation.

The next step in the evolving role for the DNS was more operational. Initially, the mapping from name to address was fixed, and indeed the original designers thought that stability of the DNS depended on it always giving the same answer to a lookup, no matter where in the world that lookup occurred. This view has now been totally reversed—as applications are designed to be highly distributed, with many points of entry across the Internet, the DNS is being used to direct a user to a nearby version of a service, so the DNS now often gives a different answer to a query depending on where the query originates. Indeed, firms are now offering very sophisticated versions of the DNS resolution service that take into account locality, loads on servers, and network performance, as well as policy considerations specific to the owner of the name. Fulfilling this need created a significant business opportunity.[128]

Finally, the DNS is coming to play a growing role as a point of control over what happens in the Internet. Since essentially all services offered on the Internet depend on the DNS, any actor who can block or modify the resolution of a name into an address can disable the associated service. To the initial designers of the DNS (and to the early engineers who made it operational), the idea that it would be used as a point of control was offensive. The DNS was a critical service on which

[128] Addressing this need is a focus of the service providers in the global market for DNS management services discussed earlier (see Note 48 *supra*).

all services depend, and since the Internet should offer a dependable and stable service, "messing with" the DNS was a terrible idea that would erode its stability, and that of the Internet.

However, the temptation to use the DNS as a point of control is irresistible. Content owners concerned with content piracy, sovereign states concerned with the control of subversive or destabilizing content, those who would protect users from malicious services, and a host of other actors have asserted their right to disable or modify the answer the DNS gives to a query. Fights over the right of control are now a major part of the policy debates around the Internet, with some advocates claiming that the DNS should not in principle be used for control, and others arguing that it is a valuable point of control.

Fights over control of the name bindings in the DNS may be the most destabilizing force on the DNS today. If users cannot count on being directed to the place they wanted to go, this may lead to the design of systems that bypass the DNS or attempt to wrest control back to the owner of the name.

4.2.1. The DNS today

We identify four considerations in trying to map out possible futures for the DNS:

- (1) importance of branding : how important are domain names for enterprises seeking to brand their products and how may branding strategies change over time?
- (2) user behavior : how are user's behavior and options for relying on domain names changing?
- (3) usability : what changes are occurring in how domain names may be used or resolved?
- (4) security: what changes in the security environment are impacting the DNS?.

We will elaborate on each of these as we explore the forces that will shape the future. Before considering these questions, however, it is important to review current usage of the DNS, including:

- How are domain names actually being used today?
- How are the patterns of use changing over time, and how will this influence how names are picked?
- How does the different patterns of use affect the switching costs if a domain owner wants to move to a new name?

Different domain owners may be using their names in very different ways. For example, as consumers become more technology savvy and comfortable with navigating the on-line world, companies may find it desirable to brand new products with domain names in new gTLDs. New ventures may like the idea of locating in a new gTLD that helps differentiate them precisely because is not ".com" or in another well-known legacy gTLD. An entity wishing to compete with Amazon.com's shopping platform, or more likely, with a vendor that is already well-known on Amazon.com's platform, may find it advantageous to locate in the .SHOP gTLD.

Today, DNS names are primarily used as the first part of a Web URL. They also can be used in other ways, for example the name of a machine on the Internet that is reached using a remote login

or remote file transfer protocol, but the predominant use is in URLs. URLs were originally conceived as a way to link web pages together -- those links are what made the web a *web*. URLs were also used as entry points into the web, by making them available in other forms of content, including (online) embedding them in email and other documents (e.g., PDF and Word documents that are not necessarily named by a URL), and (offline) including them in printed advertising or other materials. Table 1 summarizes important uses of the DNS.

≪INSERT TABLE 1: IMPORTANT USES OF THE DNS≫

Search: Today, the most common entry point to the web is probably a search result provided by a tool such as Google. Search results include a URL, which can be used to get to the actual material. As we explain further below, research shows that users often do not pay attention to the actual URL in the search results, but rather depend on the other information provided in the results. When this is the case, the domain name in the URL is not important, so there is no branding benefit from picking a particular domain name.

Links in web pages: This is the original context in which URLs were conceived. Users are well-trained that links, which are highlighted in some way, can be "clicked" to jump to the related page. In this context, the domain name in the URL may be of branding value to the owner of the target web page, in that it may help users understand where they are about to go. But the actual link text on a page pointing to the target page need not be the URL itself, which would hide the URL from the user. Hidden URLs can confound any branding objectives of the domain name.

URLs in email: Email is another common vector by which users obtain URLs. Because of the insecurities associated with email, including (in many cases) lack of trustworthy identity information, email has been a vector for delivery of URLs that point to malicious web sites. For this reason, those who try to protect users from email with malicious intent scan for such emails, and try to block the URLs or the domain names in those URLs. There is thus a cat-and-mouse game between attackers and defenders for URLs embedded in email. Again, the actual URL in the email can be hidden, so some special action by the user is required if the user wants to be sure they are seeing the actual name.

URLs in other documents: Many documents today, including reports, scholarly papers, and the like, contain embedded URLs, which may, for example, point to cited work. These documents may or may not be "in the web"—that is, they may or may not themselves be findable using a URL. They are not coded as a normal web page using HTML, but may be PDF or word files. Document readers today have been instrumented to behave somewhat like browsers, in that they allow the users to click on these links and follow them. (This action is usually done by invoking a browser to do the actual link resolution.)

We make two points about these sorts of links. First, while some links may point to the actual location of the document, a growing trend is to link to a document using what is called a Digital Object Identifier (DOI), which is a namespace unrelated to the DNS. DOIs have the following appearance: doi.org/10.1145/3402413.3402421. They are a form of *indirection*. We will return to *name indirection* below.

Tables for DNS Economics Paper

Table 1: Important uses of the DNS (mostly in context of URLs) and our four considerations about these cases.

	Branding	Security	User behavior	Usability
Search results	Hypothesis: domain name may not be important.	Users assume that URLs returned by search are trustworthy. Manipulation of results may be a concern.	Question: how often do users look at the URL?	User need not consider the URL at all. They can just click on the search results.
Following links on web pages	Some benefit. However, the visible text of the link need not be the actual URL.	A URL hidden behind deceptive link text can misdirect a user. Most users assume that links on web pages are trustworthy.	Question: do users inspect the URL as part of deciding whether to click it?	Users are well-trained to click on links.
URL in email	DNS name in URL can be helpful to users, so branding may be relevant. But actual URL can be hidden.	URLs in malicious email are prime vector for abuse. URLs hidden behind innocent-looking link text contribute to deception.	Question: Do users understand that the URL can be hidden?	Should users have to inspect the URL in an email and make judgement based on domain name? Seems unrealistic.
URL in document	Benefit of branding not clear. URLs often do not convey meaning.	Such links are usually trustworthy.	How do users search for documents?	Links may become invalid over time.

Second, when a reader of a document wants to follow a reference to another document, they have the choice of following a provided link, or using a search engine to find the document. Many users have learned that using a search engine may be more effective for several reasons. First, the link in the source paper may be out of date. The link cannot be updated in the document, once the document is published, in contrast to a web page that can be updated to point to a new version of linked content. Second, search may find multiple locations for the document, which may differ in useful ways such as not being behind a paywall.

While this list of common uses for domain names is certainly not complete, it is sufficient as a foundation to talk about how the use of domain names (and URLs) is changing.

4.2.2. Future trends for the DNS

Indirection and name resolution: In the previous discussion of Digital Object Identifiers, we mentioned that they were an example of *indirection.* A more detailed discussion of the DOI ecosystem will illustrate the general concept of indirection, as well as some of the design goals of the DOI system.

Name *resolution* is the process of making a query to a service that returns the result associated with the name. A DNS name resolves to an IP address, which is the actual location of the named object. In contrast, some names, when resolved, yield another name.[129] One might ask why this sort of name is useful? An overly simplistic example may help. Imagine that you want to meet up with someone. You could tell them that when they want to meet with you, they should come to a specific coffee shop. They can look up the name of the coffee shop and find its address. But this approach would mean that you have to be at that coffee shop whenever they set out to meet with you. A different approach would be to give them the telephone number of your office, and tell them that when they want to meet, call and your assistant will tell them where you are. The analogy is imperfect, but the resulting process has two steps. In the first step, the caller contacts the "assistant" name resolution service and presents your name, which is mapped to the name of one or another place, which the caller can then resolve into an address.

The computer science term for a name that maps to another name is an *indirect* name, and indirection is a powerful general concept. It allows the binding between the two names to be changed from time to time, and if the first name is resolved just before it is needed (so-called *late binding),* the user gets the very latest information. The other point about an indirection, of course, is that the party that has control over the binding can determine what the result of the resolution will be.

The DOI system: A DOI is of the form: doi.org/10.1145/3402413.3402421. Those names convey nothing about the origin or the owner of the document. There is no meaning to the elements, and no branding. They serve only as a means to find the content. For this reason, there is no speculative or defensive registration of DOI names. DOIs are hierarchical, like domain names, but have only two components, separated by a slash: the organization issuing the DOI and the identifier issued

[129] For further discussion of how identity management may be changing in the Internet, see Sollins & Lehr (2021).

by that organization for the object. The service that resolves DOIs is the International DOI Foundation (IDF), which describes itself as a not-for-profit membership organization that is the governance and management body for the federation of Registration Agencies providing Digital Object Identifier (DOI) services.[130]

As an aid to finding the DOI resolution service, the name can be embedded in a URL which then looks like this: https://doi.org/10.1145/3402413.3402421. This URL will enable a browser to resolve the name "doi.org" to the location of the DOI resolution service. In this way the DOI system is distinct from but depends on the DNS for its functioning.

When a user queries the DOI system to resolve a DOI, the result is a URL, but that URL does not point to the document itself, but to a web site that provides information about how to get to the document. One must go to that web site and take a step to go from that web site to the object itself. Why does the DOI return the URL of a page that provides information about the object, rather than the URL of the object itself? One of the goals of the DOI system was to provide a control point for the rights-holder for the object to exercise those rights. The page to which the DOI system directs the requestor might, for example, implement a paywall. By adding this indirection, the DOI system creates a point where the content owner can intervene. As well, the content owner can move the content at will, since the location is not embedded in a URL, but that feature may be less important than the creation of a new control point for the content owner. The DOI system could also in principle be used as a point of control, but the IDF has located itself in Switzerland, and made strong assertions about its intention to disavow this role.

Just as this paper looks at the economics and stability of the DNS, one could look at the economics and stability of the IDF. It is set up as a centrally controlled non-profit organization, which is a distinct difference from the organizations that make up the DNS.[131]

≪INSERT TABLE 2: EVOLUTIONARY FORCES THAT MAY SHAPE THE FUTURE OF THE DNS HERE≫

Link shorteners: Link shorteners are an indirection service, and at first glance, an odd phenomenon. They take a link that may have many characters in it and provide a short link that looks like bit.ly/1c92x72 or ow.ly/uK2f50RFTC9.[132] An early motivation for a link shortener was to make it easier to fit a link into a tweet. But providers of link shorteners have realized that because they are a form of *indirection*, and points of indirection create points of control, the providers can offer other services, such as analytics. The services allow the creator of the shortened link to manage the binding to the original URL, get data about usage, and so on. The ecosystem of link shorteners is now quite rich with options. The bitly and owly shorteners, by default, have no components with any semantic meaning. There are no branding or other implications in the names.

[130] See https://www.doi.org/idf-member-list.html and https://www.doi.org/registration_agencies.html.

[131] See Farhat (2017) for a discussion of why DOI is unlikely to supersede DNS.

[132] At the time we write this, these are not valid shortened links. We offer them as examples. But someday they might be. We have no control over that.

Table 2: Evolutionary forces that may shape the future of the DNS[1]

	Branding	Security and stability	Trust and control	Usability
DOI system (Indirection scheme)	DOI names and related DNS names provide no opportunity for branding.	So far, no high-payoff vulnerabilities have emerged.	Content owners must trust in stability of DOI service, and its ability to prevent others using it as point of control. System provides a point of control for rights-holders.	Users have to take multiple steps, including interaction with content description page.
Link shorteners	Most but not all link shorteners remove any branding ability	Users cannot realistically tell where the link points. They must trust the source of the link.	Links depend on availability of the resolution service. The resolution service is a point of control, which could be contested.	Some people consider short links to be more appealing and less suspicious. Will users trust them?
Specialized search	Some benefit, depending on design. The visible text of the results need not reveal the actual URL.	Most users assume that results returned by search are trustworthy. Manipulation of results is a concern.	There can be competing search services, so users are not dependent on only one. A search engine provides a powerful point of control. Could bypass DNS.	Typical services are designed with ease of use in mind.
URL in messages	URLs are often shortened, so branding may be obscured.	URLs in malicious messages may become another prime vector for abuse.	Question: Do users understand that the URL can be hidden?	Easy to click. But even more unrealistic to assume users will look at URL.
Special names in URLs	In general, no branding. Typical names are arbitrary strings.	The specialized resolution service could be target of insecurity	Users must trust that the operator of the resolution service will not abuse that control point.	Names used out of context can cause confusion.
Apps	Often URLs are not visible, so no option for branding.	App behavior is opaque. Techniques to enhance resilience as well as possible vulnerabilities are hidden.	The app controls the functioning. App need not use general name resolution scheme like DNS.	Purpose-specific design should enhance useablilty.
User-created content hosted in network services.	If URLs are used, the name is specified by the hosing site. No option for branding.	Systems under centralized control may be better able to respond to security issues.	Users must trust that hosting sites will preserve content and user identity.	System designers have focus on useability.

[1] We have removed the column labeled User Behavior, since to some extent this analysis is more forward looking, and replaced it with a new category called Trust and Control.

At the other end of the spectrum is a service like Rebrandly,[133] which will allow you to purchase any domain name you like, and then use their service to shorten your own links, thus preserving what they consider the valuable branding from your chosen domain name. And yourls.org provides open source software so a web site can implement its own link shortening.

Link shorteners raise security concerns, because it becomes very difficult for ordinary users to tell where the link actually points. The bit.ly service allows the user to append a "+" to the shortened link before following it, in which case bit.ly will show the user where the short link actually points. However, few users would know about this or undertake to use the feature.

As with other mechanisms that have been initially been devised for what might be called convenience, the option for control and intermediation seems to become a primary objective.

URLs in messages: The modern version of receiving a URL in an email is getting a URL in a text message or tweet. An original motivation for link shorteners was to fit a URL into a tweet (Twitter now provides a built-in link shortener for this purpose). The issues are more or less the same as with a URL in an email, except that the increasing use of shorteners makes it harder (or perhaps impossible for most users) to tell where the link actually points.

Specialized search: Generalized search engines like Google are being complemented by specialized search tools, such as reputation systems like Yelp or eCommerce sites like Amazon. Yelp provides a lot of information about the institution being reviewed, but also includes the URL of the website for the reviewed organization. Inspection of the actual link on a Yelp page will reveal that Yelp is acting as an indirection service—the link actually first goes to Yelp. Obviously, this arrangement allows for the collection of analytics, as well as forwarding. On the Yelp app (as opposed to the web interface) the URL of the reviewed enterprise is not even displayed until the user clicks on it. Most users may never see the URL.

With search in general, but especially with these specialized search tools, there is a fear that the operator of the service will use their control of the service to bias the results. The motivation might be economic or ideological, but whatever the reason, it is hard for a user to tell why results are returned in the order that they are.

Search services today return a URL to the user, who can click on that URL to get to the web site associated with the result. However, a search could exploit the control they have over the resolution process to completely bypass the DNS for the user. The search service could resolve the location of the desired site to an IP address (using whatever method it chooses) and return a URL that contains the IP address rather than a DNS name. This sort of action would represent a massive shift in control from the user to the search service. Whether it would be in the interest of the site associated with the search result, or adverse to that interest, would depend on the details. The high-level point is that once a new point of control is introduced into the ecosystem, the generality of the Internet mechanisms may provide a number of ways to exploit that control point.

[133] https://rebrandly.com/.

Special names in URLs: It is implicit that the name in a URL is to be resolved using the DNS. However, in some cases the creator of the URL may want to use a different system to resolve the names. There is no way to signal this explicitly. (In a different design of a URL this might be coded explicitly, but the option was not included in the original design.). What has happened in practice is that URLs have been created with names the DNS cannot resolve, so the user must understand the context in which to use the URL and only request resolution in that context.

A good example of this situation is the .onion name introduced by the TOR foundation to name hidden resources within their anonymity network. The string ".onion" looks like a Top Level Domain, but it was not issued by ICANN.[134] If a user tries to resolve this sort of URL using a standard browser, the query will fail. Only if the user knows to use the URL in the context of a browser that supports the TOR extensions can the name be resolved. The TOR Foundation defines and oversees the resolver for these names. Users of TOR generally do trust the TOR service, and assume the resolver will work correctly, but attempting to use the names in the wrong context can cause great confusion.

To the extent that other sorts of names turn up that look like DNS names but cannot be resolved by the DNS, the level of confusion may grow, and this sort of confusion may contribute to what is seen as the instability of the DNS.

The rise of apps: The development of the web led to the development of the browser, which is a general-purpose program capable in principle of displaying any web page and resolving any link. This generality implies that the method used to resolve a link has to be designed in a uniform way across all web pages. The trend we discuss here is migration away from the browser to apps that are specific to the service being offered. This shift has many implications, but from the perspective of the DNS and name resolution, it means that the app, in contrast to the browser, can use any naming scheme and any resolution mechanism it wants to name the components of the app that are downloaded across the net. Using DNS names for parts of the app is still an obvious approach, but the app could use other sorts of names, and rely on an app-specific name resolution service to map the names to content.

Most apps do not allow the user to type in URLs; they only provide access to the content they choose to make available. Thus, the Yelp app mentioned above chooses to provide a way for a user to reach the web sites of reviewed establishments, but the app does not display the URL, and when a user clicks on the "website" icon, what actually happens is under the control of the app. In contrast to a browser, where what happens has to be standardized so that any web page can use it, an app can behave as the designer chooses.

User-created content: While commercial enterprises may worry about their domain name and their brand, much content that is created today is made available through cloud services such as

[134] The status of that name is complicated. The IETF defines certain names as *special use* names. After the TOR Foundation started using the name, the IETF agreed to include .onion on the list of special use names. The IETF asserts that it is by their authority that certain names are withheld from allocation by ICANN. See https://tools.ietf.org/html/rfc6761. For an interesting discussion of the tensions here, see http://domainincite.com/19293-icann-just-gave-a-company-a-new-gtld-for-free.

YouTube and Facebook that make the content available inside their content management framework. YouTube videos do have URLs that can be extracted and pasted into other content, but the creator of that content has no control over the URL. The domain name is www.youtube.com. The user does have any control over branding, does not purchase a domain name, and so on. On Facebook, the branding is associated with the name of the account, and content is uploaded into that space. Facebook does not encourage users to link into individual pieces of content, but rather expects users to stay within the Facebook system, and explore the space of content using the Facebook tools. Some enterprises (perhaps more often small ones) have forsaken setting up a web site in favor of a Facebook presence. They get their search and branding using the capabilities of Facebook, not a general-purpose search engine or a domain name.

4.2.3. Human Behavior Online

We have discussed the various ways in which users encounter and use URLs, whether in search results or embedded in email. There are various presumptions about how users perceive URLs—for example the proposition that "just the right domain name" is critical as a component of branding. However, validating those assumptions requires that we understand how humans use, misuse, or flatly ignore the information conveyed by URLs. Fortunately, much is known about the forces that influence human behavior, online and off. In this section, we briefly review some of the relevant behavioral research.[135]

In the early days of the web, there were two ways to navigate to a web page: either click a link on another page or type a domain name into the address bar of a browser (e.g., "example.com"). If users are to remember and type a URL, a brief and memorable URL is valuable.

However, there are two potentially separate aspects to the way a user interacts with a URL. If the user is going to manually type in the URL, it needs to be short and memorable. But those attributes are separate from whether the URL implies anything about the quality or reliability of the web page. It became almost an article of faith among Internet marketers that having the right domain for a web page was critical. The implication was that the user would use the domain name as part of deciding to visit the web site.

We noted above that as the search industry emerged and flourished,[136] the results of a search query contained many indicators that the user might observe to select a result. The URL is only one of them. Building a search engine depends critically on URLs, since they facilitate the collection and indexing of Web sites.[137] However, this role is again distinct from any branding implied by the name.

---

[135] In preparing this section, we would like to acknowledge the contribution of Sara Wedeman of Behavioral Economics Consulting Group (BECG).

[136] Contemporaneously – 1994 saw the release of Lycos, Yahoo, Ask Jeeves; Altavista and Looksmart debuted in 1995; Baidu was launched in 1996; followed by Yandex and Google in 1997, and many others over the years to come.

[137] Evans (2009).

While we cannot know, precisely, whether early users were as concerned with short, memorable URLs as Web marketers believed them to be, we now have behavioral data. These data show us that either users are no longer focused on the domain name/its credibility, or are largely unaware of the issue. Inquiries into whether users pay attention to domain names suggest they do not. These experiments were not investigating issues of branding, but whether users looked at URLs to make judgments about quality, reliability and security. The evidence that users do not pay attention to URLs is seen as an increasing source of risk, due to all the security vulnerabilities inherent in visiting sites whose true provenance is difficult to ascertain, and in some cases may be purposely obscured. Even experienced Web users pay insufficient attention to URLs, and their assessments of which sites are fraudulent is rarely better than less knowledgeable users.[138,139]

Another noteworthy factor is that the domain name is a verbal cue, whereas attractiveness prompts greater trust on the part of users.[140,141] There are several eye tracking studies that examine how people *actually* search. It turns out that visual stimuli get much more attention than do verbal stimuli. Domain names are verbal: pictures and colors and multimedia stimuli are visual, and yes – they do get more attention than do words.

With the dramatic increase in use of social media, one frequently-used method for making interpersonal connections is through the sharing of links. Typically, these are embedded links, and the social media provider provides prompts that facilitate sharing, which means that the user does not even need to look at the link at all before sending it to another. This can be especially pernicious, since many social media users send things to friends and family: people they trust and who trust them.[142] This enacts the "social proof" principle of influence and persuasion,[143] wherein people feel more comfortable engaging in a behavior when they see members of their peer group behaving similarly. Users pay attention to who sends the link, not the contents of the link itself.

Security researchers are very concerned with the challenge of helping users avoid dangerous web sites. Data showing that domain names do not attract much user attention abound. Instead, a great deal of the research focuses on how to get them to pay attention, and to act accordingly. (This objective would be of great appeal to brand managers, of course.) Most researchers recommend more training to remedy user inattention to security risks, often combined with additional cues for users to perceive and understand. Evidence suggests that this approach will not be effective:[144]

> *To determine how well users are able to recognize and identify phishing web pages with anti-phishing tools, we designed and conducted usability tests for two types of phishing-detection applications: blacklist-based and whitelist-based anti-phishing toolbars. The*

[138] Albakry, Vaniea, & Wolters (2020).

[139] Thompson et al. (2019).

[140] Pengnate and Sarathy (2017).

[141] Barnes & Vidgen (2000).

[142] Abou-Warda (2016).

[143] Cialdini (1984).

[144] Li et al. (2014).

research results mainly indicate no significant performance differences between the application types. We also observed that, in many web browsing cases, a significant amount of useful and practical information for users is absent, such as information explaining professional web page security certificates.

One possible 'answer' is lying in plain sight. Computer, smart phone, and tablet screens are visual media. Visual cues carry more weight than verbal cues. Visual cues, however, deflect the attention of the user away from the textual URL.

One suggestion to improve security was that we should not try to train users to make security decisions based on URLs and similar passive cues, but rather that "the security community focus on triggering active warnings when a website's identity is suspicious." [145] The work reported in that paper provides strong evidence that users do not look at and cannot make sense of what is in a URL. This conclusion implies several things. First, it invokes an abstract group called "the security community" to solve the problem, thus throwing a major problem over the fence. But with respect to what behavioral studies teach us, it is further evidence that URLs are not a major component is the decision-making of most users. "Active warnings" are things like popup windows, which (depending on how they are designed) can either provide useful information or further confuse and frustrate the user. But the direction this advice takes us is to give the user a richer set of information on which to base decisions, ideally in a comprehensible form. The same considerations will apply in other cases, such as search results. Search results contain a variety of cues. Which of these will users focus on? Proposing alternatives to URL-based quality/credibility assessment will, of course, horrify those who think that selection of "just the right domain name" is essential to branding, but these studies suggest that this is the direction we must go to create a space in which ordinary users can make rational decisions. With respect to security, user data show us that it is within our power to increase security, but designers must work with, not against, human behavioral patterns.

The behavioral studies we report in this section are just the tip of the iceberg, and those who are concerned with how URLs are being users should avail themselves of that work.

4.2.4. Lessons about the possible future

The importance of a domain name for branding may change, and may become less important over time. URLs are hidden behind link text that does not match the actual URL, masked by link shorteners, supplemented by surrounding information in the results of web search, and so on. In the mobile space, a search may end up at an app store rather than a web page. There are still strong advocates for picking a name that captures the essence of a given enterprise, but the importance of this may fade. This trend would have important implications for competition in the DNS space.

There are a variety of ways to name various sorts of content. Content can be hosted in a traditional web site, but also hosted inside a service such as Facebook or YouTube. For content such as a

[145] Thompson et al. (2019).

publication, it can be hosted by a service that will provide a DOI for the object. Some content can be made available through an app, which means it need have no external name at all.

Many analyses of the DNS ecosystem look at it as a closed system. There are many classes of actors in the DNS system, including ICANN, the registries, the registrars, resellers and so on. But the DNS is embedded in a larger context of other actors, and that larger context may also shape and constrain how the DNS evolves. There are operators of alternative naming systems such as the IDF. There are services that host content that do not require the creator to obtain their own domain name, and so on. In trying to predict how the DNS ecosystem might evolve, it is important to speculate about how those other actors, as well as registrants, might react depending on how the DNS evolves.

4.2.5. Market power and switching costs

Our discussion of market power identified the issue of switching cost for registrants with an investment in their domain name. Although the existence of switching costs does not prove (by itself) the existence of market power, or if such market power exists, that it will be abused, reducing switching costs is generally beneficial for market efficiency, and in this case, for registrants. The challenge of how best to address this for legacy registrants is difficult, but for users now purchasing a domain name who want to avoid being captured by the registry in the future, the evolutionary path of the DNS suggests some things that a user can do.

For users that create videos, it is already a common practice to store these inside a service like YouTube rather than hosting them internally. For things like academic publications, there are lots of sites that will host content, and a site that will provide a DOI can give a permanent link that can be shared. For publishing a URL on offline (paper) material (e.g., advertising, business cards, menus and the like), link shorteners provide protection because of the indirection. Some link shortening service will generate a QR code that maps to the link, which means the user can just point their smart phone at the offline material and follow the link. This is the most extreme form of preventing the user from seeing the URL.

If a domain name owner decided to migrate to a new name, the most problematic challenge is to find the links that point to their web site and figure out if these can be changed. The most critical links would be those associated with search engines, but in the case of specialized services such as Yelp, finding these might be facilitated. Google today provides a service where the owner of a web site can get the URLs of web sites that point to the owner's site. It is not too difficult to imagine that if migration from a domain name becomes a more common requirement, firms would emerge to help manage the migration, just as there are firms today that "manage" placement in search results—a process called Search Engine Optimization or SEO. If the ecosystem moves in the direction where what we might call the "managed links" to a site (for example those returned by search engines) are orchestrated in a structured way, links from random sites are preferentially based on an indirection service, and so on, switching costs might become manageable for many domain owners.

Of course, this future need not happen. It will happen if domain name owners lose faith in the stability of the system and worry about what will happen at renewal time. If the industry can send a strong signal that users of domain names need not worry about abusive treatment, this will push

the industry in the direction of continued dependency on the DNS and economic stability of the actors in that ecosystem.

The other critical issue relates to security.

4.3. Security

Many of the pernicious problems that afflict users on the Internet today have some intersection with the DNS. At a high level, the issue is that a DNS name may turn out to point to a malicious web site. That site might mimic a well-known site, download malware, offer fraudulent products, and so on. This situation raises the question of whether the DNS system should play any role in policing abusive behavior. At one end of the spectrum are web sites that host objectionable content such as child exploitation material or terrorist recruitment or incitement. One way to take this content down would be to remove or redirect the domain name of that content. Registries and registrars do not want to be tasked with determining which content is objectionable, and this resistance seems reasonable. But if blocking of a domain is mandated, this raises issues of collateral damage (what other content might be hosted under that name) and who should be making the determination. At the other end of the security spectrum are harms that arise because a malicious actor has penetrated a registrar or registry (or a site hosting a name server) and modified the information there. That is clearly a problem the operator needs to deal with. In the middle are what are sometimes called technical abuse,[146] where malicious actors openly purchase domain names for their purposes.

In this case there are questions about the responsibility of the various actors either to track their registrants or detect and block these names. There is a secondary industry of defenders who monitor activities on the Internet and classify domains as abusive so that various other actors in the ecosystem can, if they choose, block access to those names.

---

[146] During the Fall of 2019, ICANN hosted stakeholder meetings where the DNS community focused on the need to address continuing DNS abuses (see ICANN66 (October 2019) GAC Public Safety Working Group, Registries Stakeholder Group, and Registrars Stakeholder Group, see https://66.schedule.icann.org/meetings/1116759). There was discussion of a range of issues, including: (a) Definition of what constitutes DNS abuse; (b) Best practices that might be voluntarily adopted by stakeholders; (c) Recognition and concerns that ICANN, registrars, and registries lack effective tools to address abuses; (d) Recognition and concern that terms of contracts are not being adequately enforced; and (e) Better information is needed to track effectiveness of efforts to address abuses, including more information on incidents, efforts to address, and better tools to identify responsible parties. The call to define what constitutes DNS abuse was regarded as important to establish the scope of what types of abuses ICANN, registrars, and registries may be held accountable for. This is an industry-sponsored, voluntary initiative to define the scope for identifying technical DNS abuses, see "Framework to Address Abuse," December 2019, available at http://dnsabuseframework.org/media/files/2019-12-06_Abuse%20Framework.pdf. This effort was initiated by the registry and registrar community and was launched in October 2019 as the DNS Abuse Framework (see http://dnsabuseframework.org/) and includes a number of the largest registries and registrars (e.g., PIR, GoDaddy, Donuts, Tucows). Many of these same entities were also part of the Domain Name Association (DNA, https://thedna.org/what-is-the-domain-name-association/) which is trade group that was launched in 2017 to recommend voluntary DNS security "best practices" (https://domainnamewire.com/2017/02/08/udrp-copyrights-dna-proposes-healthy-domains-best-practices/). This is consistent with the industry's preference for self-regulation

It would take a paper at least of this length to discuss the ecosystem of malicious registration (which some registrars seem to facilitate) and blocking. What we want to focus on here is possible harms to the DNS itself. In 2019, ICANN held a workshop in which participants reaffirmed community consensus that abuse continues to pose a risk to the stability and security of the DNS (although there was little consensus on what should be done to address those risks).[147] We are concerned with the ways in which abuse might threaten the DNS itself.

Penetration of systems is an obvious challenge. But it is not clear that abusers registering domain names in the ordinary way would threaten the stability of the system. There is a general insight from the security analysis of systems: abusive behavior may disrupt the overall system less than the preventive measures put in place to block the abuse. While many will call for reducing abuse, the challenge is to design mitigations that do less harm than the abuse. Blocking access to TLDs would harm all the domains registered in that TLD; blocking of a domain blocks every URL in that domain. Depending on the range of use of a domain name, the collateral damage may be minimal or massive. It is these sorts of choices that could be seen as destabilizing the DNS.

The harm from abuse itself is not the instability of the DNS but the loss of trust by users and domain name owners, who may move away from aspects of the DNS out of fear. If a user has clicked on a shortened link and ended up at a malicious web site, that user may refuse to use any further shortened links, thus harming the utility of that scheme. If users repeatedly receive malicious span that mimics the email they might get from a legitimate service, they may refuse to trust any email from online services, thus breaking an important path by which users can get information. The flood of spam pretending to be from the IRS, the Social Security Administration, banks and other critical services has added a great deal of friction to online communication, and the overall online experience. If users have ended up at an "imposter" web site due to any number of manipulations, they may lose trust in online services, rejecting online banking and other services that can otherwise be provided with great convenience over the Internet.

The overarching question is which parts of the ecosystem (which again is larger than just the DNS) should share in the responsibility to curb this behavior. We see today individual actors taking steps to protect their users and keep them from defecting from the online experience. Communication is being implemented inside apps, as opposed to using email, and so on.

But the high-level conclusion is that if the DNS continues to be a vector for the implementation of abusive behavior, users will lose trust in the DNS and service providers will reshape their user experience to avoid the use of the DNS (or avoid using it in ways that allow for abuse). This trajectory could reduce the value of domain names, and potentially the registration of new names. If an organization decides that having a presence on Facebook is more appealing and more trustworthy than a web site, those firms may move there and abandon their web presence.

[147] See ICANN66 (October 2019) GAC Public Safety Working Group, Registries Stakeholder Group, and Registrars Stakeholder Group, see https://66.schedule.icann.org/meetings/1116759. One of the most extensive analyses of DNS abuse associated with the new gTLDs is Korcyznski, et al. (2017). See, also, Piscitello & Strutt (2019).

A recurring question is the extent to which we should attempt to train the users to detect and avoid dangerous experiences. A specific example of this question is whether we should expect users to be able to look at a URL and make a value judgement about whether it might be dangerous. As we have described above, there are many ways in which what the user sees may not reflect what will actually happen: links with text that does not reveal the actual link, link shorteners, and so on. One point of view is that users cannot be expected to do this, and need to be protected by the design of the system. An extreme view is that users should never see a URL. They should always be hidden inside research results, behind useful link text, and so on. This trajectory, of course, would totally eliminate any brand value associated with a name.

This trajectory would have a major economic impact on the registries and registrars to the extent that a significant percentage of domain names are purchased speculatively, in the hope that the name might be valuable to some future owner.[148] Some TLDs have been purchased at auction for millions of dollars.[149] If it turns out that names don't matter that much, speculators will stop investing, interest in launching new TLDs will fade, and the economic consequences will be substantial. A lot of web sites today offer to help the user pick just the right domain name and TLD, and make the argument that the right name is a powerful element of branding. This hypothesis needs further evaluation, and tracking over time.

While asking users to take primary responsibility for their own security is an unrealistic and burdensome approach, it may be reasonable to ask users to understand specific rules. For example, we may be able to train users to understand the context in which a given URL is probably trustworthy. Legitimate search engines should undertake to detect and reject bogus search options that lead to malicious sites, so it is probably reasonable to train users that they can trust URLs returned from search services. Specialized search services such as Yelp must deal with bogus reviews, but it is probably reasonable to expect that a link from a Yelp review will actually end up at the site for the enterprise. On the other hand, URLs in email are untrustworthy, and users should be aware of this. These sorts of rules are probably workable for a typical user, with some training. The riskiest part of the ecosystem from the security perspective is email, because of the ability of the malicious actor to forge most of the cues that help a user distinguish malicious messages. This risk may migrate to some extent into systems like Facebook if attackers can create an account that mimics a legitimate provider, but at least in that case it is clear where the responsibility lies to police the behavior. In the case of a distributed system like the DNS, or the highly distributed email system, it is less clear where the responsibility for policing abuse should lie.

---

[148] We do not presently have good data on the activity of speculation in domain names and secondary market transactions. There are numerous sites (see earlier discussion), but it is hard to validate what their activity really is since most transactions are not publicly disclosed or verifiable.

[149] For example, one source reported that the most expensive domain name sold was cars.com that was purchased by the Gannet Co. in 2014 for $875 million (based on their financial accounting report, see https://www.sec.gov/Archives/edgar/data/39899/000003989915000006/gci-20141228x10k.htm). Another source lists the following top-10 domain sales ($millions) as Voice.com ($30), Sex.com ($13), Tesla.com ($11), Fund.com ($10), Porn.com ($10), Porno.com ($9), HealthInsurance.com ($8), We.com ($8), Diamond.com ($8million), and Z.com ($7million) (see http://www.dnjournal.com/pdf/Marchex%20Historical%20Top%20500%20Domain%20Sales-FINAL.pdf).

None the less, one path that might greatly improve the overall trustworthy character of the system would be to work out better practices and conventions for including URLs in email. There are many options there, which we do not claim to be able to enumerate, but this area is worth a focus by the development community.

Of course, a migration away from the DNS does not free those who have an online presence from fear of abuse or instability. There is no guarantee that the link shortening services are any more stable or disciplined against hold-up abuse then the registries. There is no guarantee that the DOI system will continue to operate, although the creators of that system have gone to great effort to ensure its stability (and signal that commitment). The question is who will the firms that want an online presence choose to trust. Will they trust a highly distributed system like the DNS or a centralized system?

In the larger ecosystem, the DNS is just a sideshow. Its purpose is just to facilitate a connection. Content today is being hosted in the cloud—those who use the cloud must trust in the stability, longevity and security of the cloud. And so on. The feature that makes the DNS somewhat distinctive is the potential for URLs to propagate globally in ways that makes them hard to find and impossible to update. This makes the stability of the DNS (or the systems that might replace it) of particular importance. And the highly distributed character of the DNS makes it particularly hard to evolve. What we have seen in other domains is a move from a decentralized version of a system to a more centralized version. Email used to be highly decentralized, but much of today's email is handled by Google, Microsoft or Yahoo. One of the important drivers of this centralization is improved security. Social networking has been centralized into systems like Facebook. The DOI system is run by a single organization that holds responsibility for its stability. Without predicting the exact shape of the future, if the distributed version of the DNS cannot be made stable, usable and trustworthy, we may see a centralized variant of this system emerge, perhaps run by a large single provider such as Google, that takes on the responsibility for stability.

To the actors that make up the DNS, the system itself is precious and a source of revenue. But in the larger ecosystem, it could easily be a victim of a move to better overall security. The DNS is not the tail that wags the dog.

5. Conclusions and Directions for Future Research

Our focus on the DNS is motivated by its central role in the functioning of the Internet. While we speculate about trends that might reduce its importance, today it is necessary for essentially every interaction that happens on the Internet. To the extent that the DNS is unreliable or flawed, the user experience is impaired.

We have tried to provide some insights about the economics of the DNS, focusing on the important classes of actors. While we acknowledge the considerable uncertainty that surrounds our estimates of market size, revenues and so on, our rough estimates are sufficient to make an important high-level point: while the revenues in the DNS ecosystem are important to the firms that live there, the DNS is economically a very small part of the overall Internet ecosystem.

The original designers of the DNS did not anticipate that there would be so many tussles around the DNS—they saw it as implementing a simple but essential service. These tussles have triggered

calls for one or another sort of intervention in the DNS markets: is there a need for price caps on domain names, do registries that hold many long-standing domain names have market power, are there valid objections to the creation of more TLDs? These are important questions, but the larger questions that may trigger intervention in the DNS are beyond the scope of the DNS itself. Since domain names are essential both to legitimate and malicious purposes, there are calls to change the way that the DNS operates to reduce abusive behavior on the Internet. If regulators try to impose rules on the DNS operators to improve security, or other actors move to protect themselves from the Internet abuse that propagates using the DNS, this may raise the ire of some of actors in the DNS ecosystem that are disadvantaged. In the long run, the concerns of the actors internal to the DNS ecosystem will be overridden by the larger concerns about the stability of the Internet.

The DNS is a global system. Names for an IP address anywhere in the Internet can be registered in TLDs hosted across the globe. Malicious behavior can come from anywhere to targets anywhere.[150] This reality raises issues for regulation, since domestic regulation can only be of partial utility given the current character of the DNS. Those firms that make up the DNS ecosystem should be attentive to the impact that regulation, in particular security-related regulation, may have on the overall character of the Internet and the firms that populate it. For example, it might come to pass that the owner of an IP address that wants that address to be resolvable in a certain part of the world needs to register a domain name for that address in a subset of TLDs, because only names in those TLDs can be resolved in that region. A failure to deal with abuse by the actors that make up the DNS ecosystem may lead to the fracturing of the DNS itself—it will become a regional system rather than a global system.

We do not pretend to predict the future; there are too many uncertainties. But we do identify some of the forces that might shape the future, depending on their strength. A number of these potential forces push in the direction of minimizing the value of a URL for branding purposes, which in turn would reduce the incentive for speculators to purchase and attempt to resell names. It may turn out that the DNS evolves back to its core purpose of mapping strings of letters into network addresses. However, while disputes over the creation of new TLDs and the like get visibility, another aspect of the DNS has quietly, under the covers, evolved to be a critical part of cloud-based application design. The original design specification for the DNS was that it would always give the same answer, no matter from where the query came. Today, the function is the exact opposite. As cloud-based apps are connected to the Internet in multiple locations, the DNS has become the sophisticated mapping service that tries to connect the user to the closest copy of the service. This function is critical to the efficient and resilient operation of the Internet today, and this is the next generation of the essential service that the DNS provides.

Ultimately, the challenge for DNS management, is a challenge for Internet governance. The multi-stakeholder framework that was put in place in recognition of the Internet's international character has confronted significant criticism from all sides – from those calling for an institutional framework with greater regulatory authority to act internationally and from those arguing for even

[150] Moreover, although the DNS is part of the infrastructure that enables traffic to be routed in the Internet, there are multiple other ways to threaten routing in the Internet that are not directed at the DNS. For example, it is possible to attack the inter-AS routing protocol BGP that route traffic between autonomous systems. For recent work on such attacks, see Testart & Clark (2021).

less regulatory authority, sometimes because it threatens sovereignty and sometimes because it threatens undue interference in the marketplace. We are not surprised by this conflict and recognize it as a demonstration of the success of the Internet in growing into the global, multi-stakeholder entity that it has become. And, we are hopeful that multidisciplinary efforts to better understand the Internet from a technical, economic, and policy perspectives will continue. This will require continuing efforts to collect empirical data on the performance of the Internet, the business practices of key stakeholders, and the behavioral responses of users and service providers to changing technical, market, and policy conditions. Whether the number of new gTLDs is significantly expanded or not, there will be no shortage of on-going questions as to how best to govern the Internet.

6. References

Abou-Warda, S. H. (2016) "The Impact of Social Relationships on Online Word-of-Mouth and Knowledge Sharing in Social Network Sites: An Empirical Study" *International Journal of Online Marketing* Volume 6, Issue 1

Albakry, S, Vaniea, K & Wolters, M 2020, "What is this URL's Destination? Empirical Evaluation of Web Users' URL Reading." in ACM CHI '20: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* Honolulu, United States, 25/04/20. https://doi.org/10.1145/3313831.3376168

Allemann, A. (2021), "Breaking: Ethos Capital acquires Donuts," Domain Name Wire, January 22, 2021, available at https://domainnamewire.com/2021/01/22/breaking-ethos-capital-acquires-donuts/.

Analysys Mason (2017), "Unleashing the power of all domains: the social, cultural, and economic benefits of universal acceptance," paper prepared by consultancy Analysys Mason for Universal Acceptance Steering Group (USAG), 11 April 2017, available at https://uasg.tech/wp-content/uploads/2017/04/Unleashing-the-Power-of-All-Domains-White-Paper.pdf.

Barnes, S., and Vidgen. R. "WebQual: an exploration of website quality." *ECIS 2000 proceedings* (2000): 74. *ECIS 2000 proceedings, 2000* - aisel.aisnet.org

BCG (2021), "What's in a (Domain) Name? The $2 Billion Secondary Market for Dot.Com Domains," Boston Consulting Group (BCG), January 2021, available at https://media-publications.bcg.com/pdf/DotCom_Domain_Market_Report.pdf (visited February 11, 2021).

BusinessWire (2019), "The 2019 IDN World Report Shows 20% IDN Growth," November 27, 2019, available at https://www.businesswire.com/news/home/20191127005538/en/The-2019-IDN-World-Report-Shows-20-IDN-Growth.

Carlton, D. (2009), "Report of Dennis Carlton Regarding ICANN's Proposed Mechanism for Introducing New gTLDs," prepared for ICANN, June 2009, available at http://archive.icann.org/en/topics/new-gtlds/carlton-re-proposed-mechanism-05jun09-en.pdf

CCT (2018), "Competition, Consumer Trust, and Consumer Choice Review Team (CCT) Final Report," 8 September 2018, available at https://www.icann.org/en/system/files/files/cct-rt-final-08sep18-en.pdf.

Cialdini, R. (1984), *Influence: the psychology of persuasion*, New York, NY: William Morrow & Company.

Clark, D. and kc Claffy (2015), "An Inventory of Aspirations for the Internet's Future." Technical report, Center for Applied Internet Data Analysis (CAIDA), 2015. Accessed March 21, 2019. https://www.caida.org/publications/papers/2015/inventory_aspirations_internets_future/inventory_aspirations_internets_future.pdf.

CRAI (2008), "Revisiting Vertical Separation of Registries and Registrars," report prepared for ICANN by CRA International, October 2008, available at https://archive.icann.org/en/topics/new-gtlds/crai-report-24oct08-en.pdf

Evans, D. (2009) "The Online Advertising Industry: Economics, Evolution, and Privacy" *Journal of Economic Perspectives* Volume 23, Number 3, Pages 37–60

Farhat, K. (2017), "Digital object architecture and the Internet of Things: Getting a 'Handle' on techno-political competition," Internet Governance Project, Georgia Institute of Technology, available at https://www.internetgovernance.org/wp-content/uploads/Karim_Farhat_IoT_IGP.pdf.

GoDaddy (2020), "GoDaddy Form 10-K for FY ended December 21, 2019," https://aboutus.godaddy.net/investor-relations/financials/default.aspx.

Huston, G. (2019), "DNS Wars," Blog Post, October 2019, available at https://www.potaroo.net/ispcol/2019-11/dnswars.pdf.

Huston, G. (2020), "Where is the DNS Headed?" APNIC Blog Post, 18 June 2020, available at https://blog.apnic.net/2020/06/18/where-is-the-dns-heading/.

ICANN (2013), "Who runs the Internet?", ICANN infographic, February 2013, available at https://www.icann.org/en/system/files/files/governance-06feb13-en.pdf

ICANN (2019), "ICANN Annual Report for FY2019," June 2019, available at https://www.icann.org/resources/pages/governance/financials-en

ICANN (2020), "Staff Report of Public Comments on Proposed Amendment 3 to the .COM Registry Agreement," 26 March 2020, available at https://www.icann.org/en/system/files/files/report-comments-com-amendment-3-26mar20-en.pdf.

ICC (2006), "Issues Paper on Internationalized Domain Names," International Chamber of Commerce, white paper, July 7, 2006, available at

http://www.intgovforum.org/Substantive_2nd_IGF/Issues_Paper_on_Internationalized_Domain_Names.pdf

Katz, M., G. Rosston, and T. Sullivan (2010), "An Economic Framework for the Analysis of the Expansion of Generic Top-Level Domain Names," report prepared for ICANN, June 10, 2010, available at https://archive.icann.org/en/topics/new-gtlds/economic-analysis-of-new-gtlds-16jun10-en.pdf

Kende, M. (2009), "Assessment of ICANN Preliminary Reports on Competition and Pricing," April 17, 2009, available at https://forum.icann.org/lists/competition-pricing-prelim/pdf2m9kAd0xph.pdf

Korczynski, M., M. Wullink, S. Tajalizadehkhoob, G. Moura, and C. Hesselman, C. (2017), "Statistical Analysis of DNS Abuse in gTLDs Final Report," available at https://bit.ly/2TVvg5j.

KPMG (2010), "New gTLD Program: Benchmarking of Registry Operations," available at http://icann.org/en/topics/new-gtlds/benchmarking-report-15feb10-en.pdf

Kreuger, F. and A. Van Couvering (2010), "A Quantitative Analysis of Trademark Infringement and Cost to Trademark Holders in New gTLDs," Minds + Machines Working Paper 2010-1, February 10, 2010

Lehr, W., D. Clark, S. Bauer (2019), "Regulation when Platforms are Layered," TPRC47: Research Conference on Communications, Information and Internet Policy, Washington DC, September 2019, available at http://ssrn.com/abstract=3427499

Lehr, W., D. Clark, S. Bauer, A. Berger, and P. Richter (2019), "Whither the Public Internet?" Journal of Information Policy 9 (2019): 1-42. doi:10.5325/jinfopoli.9.2019.0001

Li, L., Berki, E., Helenius, M. & Ovaska, S. (2014) "Towards a contingency approach with whitelist- and blacklist-based anti-phishing applications: What do usability tests indicate?" Behaviour & Information Technology, Vol 33(11) 1136-1147

Liu, B., Lu, C., Li, Z., Liu, Y., Duan, H.X., Hao, S. and Zhang, Z. (2018) "A Reexamination of Internationalized Domain Names: The Good, the Bad and the Ugly," In *DSN* (pp. 654-665).

Manheim, K.M. and Solum, L.B., 2002. An Economic Analysis of Domain Name Policy. *Hastings Comm. & Ent. LJ*, *25*, p.359

Mueller, M. (1998), "The battle over Internet domain names: Global or national TLDs?" *Telecommunications Policy*, *22*(2), pp.89-107

Mueller, M. (2002a), "Competing DNS Roots: Creative Destruction or Just Plain Destruction?" *Journal of Network Industries*, (3), pp.313-334.

Mueller, M. (2002b), Ruling the root: Internet governance and the taming of cyberspace, Cambridge, MA: MIT Press.

Mueller, M., J. Mathiason and L. McKnight (2004), "Making sense of 'Internet Governance': defining principles and norms in a policy context," Working Paper Internet Governance Project, Syracuse University, April 26, 2004.

Mueller, M. L., & McKnight, L. W. (2004). The post-.COM internet: toward regular and objective procedures for internet governance. *Telecommunications Policy, 28*(7), 487-502. doi:https://doi.org/10.1016/j.telpol.2004.05.005

Mueller, M. (2006), "Toward an economics of the domain name system," in *Handbook of Telecommunications Economics: Technology Evolution and the Internet*, *2*, pp.443-487.

Mueller, M. and Badiei, F. (May 2017) *Governing Internet Territory: ICANN, Sovereignty Claims, Property Rights and Country Code Top Level Domain Names*, Columbia Science & Technology Law Rev., 18, pp.435-515.

Nicholls, T. (2013) "An empirical analysis of Internet top-level domain policy," Journal of Information Policy 3 (2013): 464-484.

NRC (2005), Signposts in Cyberspace: The Domain Name System and Internet Navigation, National Research Council (NRC): National Academies Press, 2005, available at https://www.nap.edu/download/11258.

OECD (2004), "Generic Top Level Domain Names: Market Development and Allocation Issues," DSTI/ICCP/TISP(2004), 13 July 2004, available at http://www.oecd.org/internet/ieconomy/32996948.pdf.

Pengnate, S. and Sarathy, R. (2017), "An experimental investigation of the influence of website emotional design features on trust in unfamiliar online vendors," Computers in Human Behavior, Vol. 67, 49-60.

PIR (2020), "Public Interest Registry 2019 Annual Report," available at https://thenew.org/app/uploads/2020/06/PIR-2019-Annual-Report-PAGES-1.pdf.

Piscitello, D. and C. Strutt (2019), "Criminal Abuse of Domain Names: Bulk Registration and Contact Information Access," report prepared by Interisle Consulting Group, LLC, available at http://www.interisle.net/criminaldomainabuse.html

RAPG (2010), "Registration Abuse Policies Working Group (RAPG) Final Report," prepared for submission to ICANN, 29 May, 2010, available at https://gnso.icann.org/sites/default/files/filefield_12530/rap-wg-final-report-29may10-en.pdf.

Schomp, K., M. Allman, and M. Rabinovich (2014), "DNS Resolvers Considered Harmful." In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks - HotNets-XIII*, 1–7. Los Angeles, CA, USA: ACM Press, 2014, available at https://doi.org/10.1145/2670518.2673881.

Sollins, K. and W. Lehr (2021), "Exploring the Intersection of Technology and Policy in the Future Internet Architecture Effort," 48th Research Conference on Communications, Information and Internet Policy (TPRC48, www.tprcweb.com), February 2021.

Testart, C. and D. Clark (2021), "A Data-driven approach to Understanding the State of Internet Routing Security," 48th Research Conference on Communications, Information and Internet Policy (TPRC48, www.tprcweb.com), February 2021.

Thompson, C., Shelton, M., Stark, E., Walker, M., Schechter, E. and Felt, A.P. (2019) "The web's identity crisis: understanding the effectiveness of website identity indicators," In *28th USENIX Security Symposium (USENIX Security 19),* pp. 1715-1732.

Tucker, C. and G. Rafert (2015), "Phase I Assessment of the Competitive Effects Associated with the New gTLD Program," available at http://newgtlds.icann.org/en/reviews/cct/competitive-effects-phase-one-assessment-28sep15-en.pdf

Tucker, C. and G. Rafert (2016), "Phase II Assessment of the Competitive Effects Associated with the New gTLD Program," available at https://newgtlds.icann.org/en/reviews/cct/competitive-effects-phase-two-assessment-11oct16-en.pdf

Verisign (2020a), "Domain Name Industry Brief," Vol. 17-Issue 3, August 2020, available at https://www.verisign.com/assets/domain-name-report-Q22020.pdf.

Verisign (2020b), "Verisign 2019 Annual Report and 2020 Proxy Statement," April 2020, available at https://investor.verisign.com/financial-information/annual-reports.

Verisign (2020c), "Troubling Efforts to Distort and Undermine the Multistakeholder Process," Letter to ICANN, February 4, 2020, available at https://mm.icann.org/pipermail/comments-com-amendment-3-03jan20/attachments/20200214/00fd9745/VerisignPublicComment_.pdf.

Verisign (2021), "Domain Name Industry Brief," Vol. 18-Issue 1, March 2021, available at https://www.verisign.com/assets/domain-name-report-Q42020.pdf.