

**RESEARCH IN PROGRESS: NOT FOR CITATION
WITHOUT AUTHOR PERMISSION**

DNSSEC and Government Interests in Internet Architecture

Brenden Kuerbis, Doctoral Student
Syracuse University, School of Information Studies

August 31, 2006

Abstract: This exploratory research examines the development and adoption of the domain name system security extensions (DNSSEC) standard. It identifies that the DNSSEC standardization process made a distinct shift beginning in 2000-2001, as U.S. government concerns about use of the Internet infrastructure intensified. U.S. government influenced changes in the standard included restricting the possibilities of public-key distribution within the DNS, and leveraging its oversight of the root zone file to affect how the standard would be deployed operationally. More generally, the research highlights how the U.S. government responded to the liberalization of the export control regime for encryption by mounting an effort to prevent the basic standards and architecture of the Internet from enabling encrypted communications. Additionally, it draws attention to the government's security and economic interests in DNSSEC and the evolving relationship between the ostensibly private sector standards setting organizations (i.e., the IETF), which was created by government, and state power in the era of Internet protocol communications.

1. Introduction

The DNS is a distributed, hierarchical database, providing the ability to lookup address information (and more) of host machines connected to the Internet. It does this by mapping domain names (e.g., ICARUS.SYR.EDU) with which humans are more familiar, to numeric internet protocol addresses (e.g., 128.230.1.49), which provide the network with required routing information. When a DNS server is queried by a host's resolver software, it resolves a domain name to a specific IP address and then returns the results to the querying host.¹ Once complete, a host can initiate communication and send packets to the destination host.

¹ Or IP address in the case of a reverse lookup.

As such, the DNS is a critical component of the Internet's infrastructure which is queried billions of times each day by Internet users. While obvious that the DNS has held up quite well given the rapid growth of the Internet over the last 10 years, there is general agreement that as operated today the system is vulnerable to data corruption which could impair its functionality (National Research Council 2005).² These vulnerabilities were first identified in 1987, elaborated on in 1990, and included identification of threats concerning forged data responses and cache poisoning (Mockapetris 1987; Bellovin 1995).

In response to these threats, the Internet Engineering Task Force (IETF) proposed the Domain Name System Security Extensions (DNSSEC) which introduced public-key cryptographic signatures into the DNS infrastructure to ensure the integrity and authenticity of information retrieved by DNS queries. The DNSSEC standard was first published as a RFC in 1997, around seven years after security vulnerabilities in the DNS initially became publicly known and about six years before it became an IETF topic of discussion.³ In 2001, operational issues were discovered in the standard as written related to its key handling capabilities. As a result, changes were proposed to the standard, which was re-written in three documents (i.e., 2535bis) between 2001 and 2005. In March 2005, the revised DNSSEC standard was approved by the Internet Engineering Steering Group (IESG) and published in three separate RFCs covering requirements, additional resources, and protocol modifications.⁴ In addition to significantly expanded technical dimensions, the RFCs now had numerous authors, including technical experts from registries, DNS management and software providers;

² It is also vulnerable to other types of attacks, notably, Denial of Service Attacks, one of which infamously disabled several root servers in 2002, see <http://d.root-servers.org/october21.txt>. DNSSEC does not address this vulnerability.

³ See Eastlake III, D. E., & Kaufman, C. (1997). *RFC 2065: Domain Name System Security Extensions*. Retrieved March 15, 2005, from <http://www.ietf.org/rfc/rfc2065.txt>. This document became obsolete with the publication of Eastlake, D. (1999). *RFC 2535: Domain Name System Security Extensions*, from <http://www.ietf.org/rfc/rfc2535.txt>

⁴ See Arends, R., Austein, R., Larson, M., Massey, D., & Rose, S. (2005). *RFC4033: DNS Security Introduction and Requirements*. Retrieved March 15, 2005, from <http://www.ietf.org/rfc/rfc4033.txt>; Arends, R., Austein, R., Larson, M., Massey, D., & Rose, S. (2005). *RFC4034: Resource Records for the DNS Security Extensions*. Retrieved March 15, 2005, from <http://www.ietf.org/rfc/rfc4033.txt>; and Arends, R., Austein, R., Larson, M., Massey, D., & Rose, S. (2005). *RFC4035: Protocol Modifications for the DNS Security Extensions*. Retrieved March, 2005, from <http://www.ietf.org/rfc/rfc4033.txt>.

government agencies; and individuals associated with government and industry supported research centers. While not unusual to see authorship of IETF standards fluctuate in response to evolving expertise requirements, this wholesale change of authorship signaled a distinct maturation of organizational interest in shaping the direction the standard took.

This document proceeds in the following manner: Section 2 reviews relevant theory applicable to the study of DNSSEC's development and adoption into the Internet's infrastructure. Section 3 provides an overview of the basic functionality provided by DNSSEC and how it introduces public key cryptographic signatures and a hierarchical model of trust to the DNS. Section 4 outlines government interest in DNSSEC and recalls some critical developments so far in the evolution of DNSSEC, including early government influence limiting the scope of public key storage in the DNS, and the leveraging of USG oversight of the root zone file to influence key management. These actions begin to weave a story of how the U.S. government responded to the potential widespread use of encrypted communications, enabled in part by the liberalization of the export controls regime for encryption software and network effects associated with an authenticable DNS, by mounting a concerted effort to prevent core standards and architecture of the Internet from enabling these technologies. Additionally, it draws attention to the evolving relationship between the ostensibly private sector standards setting organizations (i.e., the IETF), which was created by government, and state power in the era of Internet protocol communications.

2. Theorizing DNSSEC

Our knowledge of the economics and political economy of standardization and previous response of government to instances of technological change is what makes the development and adoption of DNSSEC an interesting area of inquiry. Economic theories of standardization inform us that there is no guarantee a standard will be adopted, even when everyone benefits from it. And even if a standard is agreed upon, it may not necessarily be the socially optimal one. This is because standardization activities play a core role in corporate strategy and public policy making (Besen and Johnson 1986). Standards can be used for competitive strategy, either by firms purposefully pursuing

incompatibility or by promoting their own “sponsored” technology to achieve competitive advantage (Besen & Saloner pg. 219). We also know that market power and financial resources are important to getting a particular standard adopted (Weiss and Sirbu 1990). Recent work has made great strides in clarifying the economic incentives surrounding DNSSEC that could inspire firms to act strategically. Fetzner and Jim (2004) identify potential conflicts of interest for the largest registry, VeriSign; due to implementation costs associated with DNS server software development, and signing and maintaining (i.e., periodic key-rollover) the largest zone file. Furthermore, DNSSEC potentially could cannibalize its existing Secure Sockets Layer (SSL) certificate authority business, which accounts for a substantial and growing portion of their revenues.⁵ Drawing largely from the economics literature on network effects and externalities, Ozment and Schechter (2006) develop a model for bootstrapping the deployment of security technologies into existing networks and suggest that a combination of software development subsidies, partial adoption mandates, and bundled technologies could increase the likelihood and rate of DNSSEC adoption. Huston (2006) notes that the additional operational and infrastructure costs associated with deployment are not trivial, particularly for registries which run larger zones, and manufacturers of client resolver software, and other end-user applications. In each case, it is not abundantly clear that the benefits outweigh the costs. Zone operators would inherit the additional workload of key management and signature generation, with no immediately apparent benefits. End users and application developers would possibly gain the most, with enhanced confidence in DNS query resolutions and the possibility of using a secure DNS for authenticable data

⁵ SSL provides for some degree of authentication, relying on built-in software in commonly used browsers to obtain digital certificate information that can be used to authenticate a server the user is connecting with and establish an encrypted session based on public-key cryptography. However, SSL is heavily reliant on the user manually establishing the veracity of the certificate. The cornerstone of VeriSign’s SSL business, digital trust services, was established in 1995, with the introduction of website digital certificates. This line of business was expanded to sell enterprise PKI services that allowed business to deploy customized digital certificate services for use by employees, customers and business partners (VeriSign 2002). In 1998, the company began promoting, selling and distributing PKI services under contract with global affiliates. After adjusting for revenues derived from the acquisition of U.S.-based Network Solutions in June 2000, revenues from international affiliates exceeded or approached 25% of total revenues between 1999 and 2001. VeriSign predicts that revenues from these security services will continue to grow, while its other line of business remain steady or dwindle (VeriSign 2004).

distribution.⁶ Overall, each actor is evaluating the tradeoffs associated with cost of deployment versus “practical probability of attack and the potential consequent costs.”

The economics literature also notes that in environments characterized by network effects⁷ seemingly insignificant events (e.g., performance of prototypes, political circumstances) can become consequential, and “tip” markets toward a particular standard (Arthur 1989, 116; David 1985). Similarly, historical path dependence in development can result in technologically inferior standards emerging to take over a market (e.g., QWERTY keyboards) (David, 1985). Together, incumbent standards can “lock-in” and be remarkably persistent due to their installed base and network effects, dominating new technologically superior standards that do not enjoy network effects (David 1985; David and Greenstein 1989). With few exceptions the standards economic literature examines the role of the firm, although some have examined the role of government in the development of standards. David (1987) argues that government’s influence is indirect, often as an important buyer or user of products which incorporate certain standards. Greenstein (1992) importantly identifies that standards may develop through market mechanisms and organizations that combine market participants and government guidance.

The political economy of standardization situates the standards adoption process in the global context. It clearly demonstrates that international actors, particularly governments, have strong economic interests in and preferences for specific technological standards. Perhaps most importantly, it begins to uncover the domestic sources of those preferences, and how government policy toward technical standards is usually connected to supporting domestic industry and protecting markets. Examples of domestic interests shaping government positions in standards negotiations include the

⁶ Huston is referring to the use of the CERT RR (see RFC 4398) in a DNSSEC environment for distributing certificates. However, the CERT record is reliant on a certification authority outside the DNS. See <http://www.cafax.se/dnssec/maillist/0000-00/msg00242.html> and <http://www.cafax.se/dnssec/maillist/0000-00/msg00244.html>.

⁷ Networked product environments are characterized by network effects, or the increasing utility associated with each additional network user. (Rohlf's 1974) These effects are an externality and contribute to a positive feedback mechanism, which drives an increasing installed user base and the creation of complementary products that encourage further adoption.

development of PAL/SECAM and NTSC in terrestrial broadcasting (Crane 1979), the contestation between GSM and CDMA in cellular communications (Pelkmans 2001), and the long progression and eventually divergent HDTV standards in the U.S., E.U. and Japan (Hart 2004; Galperin 2004). However, while domestic economic demands often dictate government positions in standards development, governments have institutional and political interests of their own, especially in the area of national security and surveillance. These interests often come into conflict with the economic interests of domestic industry groups and often occur during instances of technological change. Satellite communications, global positioning systems and encryption software provide examples.

Throughout the 1970s and 1980s, the USG confronted changes in satellite communications which forced tradeoffs between foreign and economic policy (Snow 1997; 1985). Policy decisions made were shaped largely by foreign relations, national security and intelligence interests within the government who were potential users and beneficiaries of the changes (Kavanaugh 1985). Mueller (1991) examines how the U.S. government responded to competitive pressures on Intelsat⁸ by issuing a Presidential Directive which set a “separate system policy” (SSP) that permitted alternative satellite carriers, but restricted their ability carry public switched network traffic (PSTN). This political bargain sought to remedy conflicting political and military interests expressed by the Department of State and Department of Defense (DoD), with commercial interests voiced by the Commerce Department and Federal Communications Commission. In creating the SSP, the U.S. simultaneously redirected competitive pressures and protected Intelsat’s market and carefully crafted governmental coalition, while publicly promoting liberalization in satellite communications. However, as Mueller points out, the success of alternative satellite carriers, increasing opposition to the PSTN restriction, and competition from lightly regulated undersea cable operators initiated the undoing of Intelsat’s monopoly on international PSTN traffic. Despite the attempt to maintain tight

⁸ Intelsat was a geopolitical initiative of the United States which created an international consortium for the commercial development of satellite communications created in 1965 at the height of the U.S.-Soviet space race, see Snow (1977) for a detailed history and the competition and foreign policy issues involved.

control over the regime, the demand for effective competition in the global marketplace usurped both foreign and security policy objectives.

Another well-known example is the evolution of the global position system (GPS). Much like the Internet, the GPS originated as a Defense Department research project, and quickly became a mission critical system for the U.S. military. Similarly, its global accessibility and widespread usefulness created strong demand for it to be opened to civilian use after initial development. Eventually, the DoD labeled the GPS as a “dual-use,” with global positioning technology turning into a multi-billion dollar industry, supporting scientific, engineering and commercial applications, and creating a bevy of interests in administration of the GPS. The recognition that the GPS could provide substantial economic benefit to users of navigation information and the conflict with security-oriented interests to maintain tight control over the technology was recognized early on in the development of the GPS (Lachow 1995). By 2000, the US government also recognized the need to incorporate broader opinions in shaping the development of the technology (Larsen 2001). However, no single forum existed for addressing the full range of concerns regarding GPS or for making agreements (Pace 1996).⁹

National security concerns revolved around the availability and accuracy of GPS signals. As originally designed, the GPS provided selective availability (SA) of two signal standards, one crude, “standard” positioning service signal (SPS) with accuracy to plus-or-minus fifty meters for civilian use, and one encrypted, precise positioning service signal (PPS) with accuracy to plus-or-minus ten meters for military use. In response, technological improvements, including alternative differential GPS (DGPS) reference stations emerged in the marketplace which augmented the GPS signal, offering better positioning data than PPS. Eventually SA was eliminated, providing civilian and military users the same level of positioning and navigation accuracy (Larsen 2001; Ashkenazi 2000). Nonetheless, worries that signal could be degraded has driven interest in

⁹ Interestingly, this has parallels with the current issue of Internet governance. Only recently was the Internet Governance Forum created with the express purpose of facilitating discourse among diverse stakeholders and “between bodies dealing with different cross-cutting international public policies regarding the Internet,” see the Tunis Agenda at <http://www.itu.int/wsis/docs2/tunis/off/6rev1.html>

developing a competitive infrastructure. Lembke (2002a; 2002b) describes competition surrounding global positioning systems, particularly the emergence of GALILEO and the intersection of technology and security policy within the European Union. The development of the GALILEO satellite navigation system has benefited from the desire of industry representatives and political actors to have a European alternative to the American GPS system.

An instance particularly relevant to the study of DNSSEC is the USG response to the growth of the Internet and use of encryption software. As Diffie and Landau (1999; Forthcoming) have covered in detail, the USG has a long history of trying to control the use of encryption. Attempts to control the spread of cryptographic software were largely based on export controls, which sought to constrain software by function and key length (9).¹⁰ Other techniques included unsuccessful attempts to weaken encryption standards (i.e., DES) by escrowing key data to allow for government access to communications, and successful attempts by law enforcement to gain statutory authority (i.e., CALEA) to require wiretapping capability in communications (and more recently information) carriers networks.¹¹ Intergovernmental agreements led by the US were also pursued, like the Wassenaar Arrangement and its COCOM predecessor, which sought to use countries domestic export authority to provide notification and curb the spread of encryption software.

Confidentiality aspects of public key encryption technology within DNSSEC standard have long been of concern to the USG.¹² In fact, early DNSSEC deployment efforts were tripped up by export controls, resulting in the temporary removal of BIND¹³ prototype source code from the Web. Originally classified in 1996 by the Bureau of

¹⁰ These controls were largely based on the key length, e.g., 40 bits, software used. Key length determines the difficulty of unauthorized decryption in a properly designed encryption algorithm. E.g., an increase of one bit doubles the cost to an intruder (Diffie and Landau, 2005 pg. 9)

¹¹ FCC order

¹² As noted earlier, DNSSEC uses public key encryption for authentication and determining integrity of DNS data. However, the same encryption algorithms used to provide that functionality can also, with modification, be used to provide data confidentiality – the main issue of concern for authorities conducting surveillance.

¹³ BIND, or the Berkeley Internet Name Daemon, is an implementation of DNS server software published by ISC and used on the vast majority of name serving machines on the Internet.

Export Administration (BXA) as authentication software and therefore exempt from Export Administration Regulations, the software was reclassified as a controlled item almost a year later by BXA, making it subject to export restrictions. According to Tien (1999) BXA concern centered over the inclusion of RSAREF, a collection of the various cryptographic routines source code which could be modified to provide data confidentiality. However, as Tien (1999) argued, BXA had questionable jurisdiction and its action was based upon an unwritten rule to reclassify the software. Eventually, with liberalization of export controls, the restriction was lifted.¹⁴

Over the 1990s the USG reluctantly liberalized its export control policy for encryption software. This policy choice was partly driven by the advances of public-key cryptography and its enablement of Internet commerce, which forced confrontation between public interest groups and the software and banking industry, reliant on the technology for secure communications and transactions in a global economy, and the intelligence and law enforcement communities engaged in surveillance activities.¹⁵ Additionally, the widespread shift of norms and practices in software development to open source modes decreased the controllability of software (Diffie and Landau, Forthcoming pg. 17).

By January 2000, the USG had lifted nearly all export restrictions on encryption software products save a few important differences (14). The liberalized rules were targeted largely at retail products, and distinguished between commercial and government customers.¹⁶ Crucially, sales to governments and products intended for protecting large commercial communications infrastructure could not be exported without licensing arrangements subject to review by the Department of Commerce's

¹⁴ See Gilmore (2000) at <http://www.chiark.greenend.org.uk/pipermail/ukcrypto/2000-February/047193.html>

¹⁵ See Diffie and Landau (2000), Giacomello (2005) for histories of how the “unlikely coalition” of public interest advocacy and private sector interests defeated government law enforcement interests, resulting in the almost complete liberalization of U.S. export controls on strong encryption software in January 2000. See also Levy (2001).

¹⁶ Eventually, remaining retail restrictions pertaining to “one-time reviews” were eliminated in response to European liberalization. This reaction was driven by the exposure of the USG Carnivore program, which engaged in commercial espionage of US allies.

Bureau of Industry and Security (15). This restriction has proved especially problematic as communications over the Internet often travel a diverse path of commercial, government (i.e., PTT), and private networks and was not necessarily “compatible with the U.S. desire to protect the critical infrastructure of the industrialized world (20).” As Diffie and Landau importantly conclude, the revised export control regime recognized the importance of international commerce and shifted the burden associated with surveillance to governments.

DNSSEC raises complex policy issues as governments struggle to balance different – and many times, contending national interests related to security, economic or foreign policy objectives. Enduring questions remain about whether security or rather economic interests and potentially anti-competitive behavior drive government policy response to technological change. In considering the case of DNSSEC, a central policy problem identified is who maintains the private key and signs the root zone (Mathiason, Mueller, Klein, Holitchser, & McKnight 2004), and more importantly, the motivations for why the root will or will not be signed. These questions become particularly intriguing when we consider how these government interests are pursued given the nature of Internet communications, a tapestry of open standards and largely privately-owned infrastructure which does not map perfectly to territorial states. More generally, this problem raises interesting questions for policymakers and researchers, including:

- As strategic actors themselves, how do states integrate domestic economic and national security objectives into Internet standards?
- Similarly, how do non-state actors respond to these objectives, are they complicit or merely acting as expected? How do global actors navigate strategic government behavior and integrate their own interests into Internet standards? More generally, how much autonomy do non-state actors actually have?

3. DNSSEC, function and form

How does DNSSEC work?

Using data contained in four additional DNS resource records (RRs), DNSSEC authenticates the origin of DNS data for a particular *zone*¹⁷, assures integrity of a response from the DNS, and authenticates denial of existence.¹⁸ Together, these measures provide general DNS data authentication, and protect against tampering in DNS caches and forged DNS responses. Huston (2006) among others provides a detailed explanation of each RR type and their interoperation, a quick overview which draws from this work is as follows:

DNSKEY – the public key portion from a cryptographic key pair associated with a specific secured DNS zone (e.g., .org). The complementary private key is maintained in a secure manner by the zone administrator and used to sign zone data. Together they are used to encipher/decipher a zone’s data.

RRSIG – a signature for a zone’s specific resource record set (RRset). E.g., the DS or A records in the .com zone have an associated RRSIG, which is generated by creating a hash of that RRset and signing it with the zone’s private key.

NSEC – is a “gap spanning record,” used when there is no authoritative data to return to resolver DNS queries. Its compliment, NSEC3, accomplishes the same but additionally prevents the accidental “discovery” of the entire contents of a zone file using a hash of ordered domain names.

DS – a record signed by and authoritatively served from the parent zone, which contains a hash of a child zone’s DNSKEY.¹⁹ Importantly, a resolver uses this RR

¹⁷ The DNS is delineated into a hierarchy of zones, e.g., the root (“.”) zone, the .edu or syr.edu zones. The root zone is the parent of the child .edu zone (and other TLDs), which in turn has a child zone(s), e.g., syr.edu. Each zone maintains a DNS server(s), which provides mapping data. Most of the data stored pertains to the immediate zone; however, DNS servers also maintain a cache of recent DNS query resolutions. If the DNS server file doesn’t contain specific address information for a received query, it returns the IP address of another DNS server in the domain name hierarchy which could possibly resolve the query.

¹⁸ This is the case when a query for a non-existent domain name is submitted by a resolver.

¹⁹ Each child zone maintains a key signing key (KSK) and zone signing key (ZSK). “The child zone has its ZSK signed by its KSK. The public parts of both the KSK and the ZSK are served from the apex of the child zone. The KSK can be exposed less frequently than the ZSK and thus be stored securely. The ZSK is needed to sign the other records in the zone, is used more often and thus may be changed rather frequently, with a new ZSK needing a signature by the KSK. A ZSK roll over does not need an interaction with the parent zone. A child zones ZSK is signed by its KSK. (Arends and Koch 2005)”

“delegation point” to authenticate a child zone’s public key. The DS record can additionally be used to implement trust anchors in cases where the complete delegation hierarchy is bypassed. (Arends and Koch 2005)

It is important to recognize that any system that uses public key encryption needs an infrastructure to validate those public keys. In this regard, DNSSEC introduces public-key infrastructure (PKI)²⁰ qualities to the DNS (Ateniese & Mangard 2001). By using chains of public-key signatures to authenticate DNS data, it essentially forms a certificate chain similar to that used by SSL or PGP (Weiler 2004). The DS RR is analogous to a certificate which binds a globally unique entity (e.g., a zone) to a specific public key and is issued, authenticated and distributed by an authoritative parent zone. The veracity of any child zones signed data (e.g., RRSIG(A)) is assured in part by querying the appropriate DS record from the parent zone. In this manner, the implementation of DNSSEC creates a trusted directory of naming data.

The process highlights a chosen top-down conceptual model of trust which maps to the existing hierarchical structure of the DNS (Perlman 1999). As Massey et al. (2001) outline, this hierarchical or “tree-based” validation model is not the only trust model possible for DNSSEC; alternatives include trust lists and mesh architectures.²¹

²⁰ The concept of public-key infrastructure systems dates to Diffie and Hellman’s (1976) proposal of public-key cryptography and suggestion of a Public File key directory that user’s would refer to find other user’s public keys. Noting the bottleneck limitations and security risks of using a single file to distribute public keys in a large network, Kohnfelder (1976) proposed the concept of a certificate, which binds a entity’s identity to a key using a digital signature and stores them in a repository. Importantly, this alleviated the need to trust the repository, and allowed it to be replicated and made fault tolerant (Gutman 2002). However, despite removing these burdens, a certificate authority is still required to attest to the authenticity of the digital certificate (i.e., the relationship between the public key and entity’s identity). The combination of a unique identity bound to a public key linked by a digital certificate and issued from a certificate authority, along with procedures for issuing and revoking those certificates, is known as a public-key infrastructure (PKI).

²¹ See also United States General Accounting Office (2001), pgs. 62-66, for a general description of alternative trust models in PKIs.




Figure 1: Hierarchical Model Certification Path (Source: GAO-01-277)

Any model of trust has both strengths and weaknesses. Straight forward in its approach, a top down model (see Figure 2 above) employs a single “root” certificate authority that is trusted by all users and which issues certificates to subordinate certificate authorities, who may subsequently issue their own certificates. By following certification or “trust” paths, an end resolver can always verify the validity of a certificate by tracing back to the known and trusted root. Obviously, this model emphasizes the existing hierarchical organizational structure of the Internet’s DNS infrastructure (i.e., ICANN → Registries → Registrars) and highlights the need for common agreement among all the actors on certain management and operational policies. It also highlights a potential bottleneck in the root zone file (RZF). In this model, the RZF represents a single point of failure, if DNSSEC related data is compromised there, all subordinate certificate authorities and certificates issued must be replaced, a potentially massive undertaking. Furthermore, longer the trust paths require more intensive processing by validating resolvers.

4. Shaping DNSSEC

US government interests in DNSSEC development and adoption originate amongst a variety of organizations, including administrative, agency, research and military. These interests have manifested themselves in activity on multiple fronts, both related to the standards development and policy choices related to operationalization.

DHS and DoC collaboration

Formally created in 2002, the Department of Homeland Security (DHS) became quickly active in the area of cyber security; its activity grounded in several high profile reports outlining the critical nature of the Internet's infrastructure, including the DNS.²² An indicator of DHS's particular interest in the development of the DNSSEC manifested itself in a 2003 Memorandum of Understanding (MoU) with the Department of Commerce's Technology Administration, which includes the National Institute for Standards and Technology. Early versions of the MoU specifically called for "technical collaboration, [and] review of each others work [in the area of] encryption standards (Carney 2003)." As Carney notes, given the law enforcement and national security focus of the DHS, collaboration in the area of standards, particularly those that utilize encryption techniques, continued to be an area of concern for policymakers.²³

By 2004, DHS's Information Analysis and Infrastructure Protection Directorate was,

Focused on securing the domain name infrastructure [by advancing] the diffusion and use of the Domain Name System Security Extensions (DNSSEC) protocol as a replacement for the traditional domain name infrastructure...[including working]...with Federal researchers and officials and the private sector to develop a roadmap to accelerate the development and deployment of a secure domain name infrastructure. [This work also included] the identification of technology requirements and development of models to aid in assessing the performance impact of utilizing DNSSEC in operational environments (U.S. House of Representatives Committee on Science 2005a, pg. 62).

This collaboration was manifested in the National Institute for Standards and Technology's Computer Security Division taking an active role in the standard's

²² See e.g., the *National Strategy to Secure Cyberspace* at <http://www.whitehouse.gov/pcicb>, and the *National Strategy for the Physical Protection of Critical Infrastructures and Key Assets* at <http://www.whitehouse.gov/homeland/book/>

²³ See Diffie and Landau (1998) for the history of NSA role in encryption standards and the Computer Security Act of 1987, which moved official development of federal encryption standards (for unclassified material) to NIST from NSA. The point is still not lost on lawmakers; the Bush administration proposal to move the NIST's Computer Security Division to the DHS (see Homeland Security Act of 2002 as introduced in House) during its formation was met with strong resistance by Senator Cantwell (WA) and others familiar with earlier battles over encryption, ultimately the proposal failed.

development by participating in the IETF’s DNSEXT Working Group. In terms of the IETF development process, Figure 1 (below) indicates how government interest in DNSSEC shifted significantly beginning in 2001. From the early 1990s to 2000, USG contributions on the DNSEXT Working Group listserv, *namedroppers*, were diffused across a variety of organizations, totaling almost 90 comments.²⁴ The bulk of USG comments to the list during that time came from Department of Defense related entities, including DARPA, DISA, military research labs and NASA. However, in 2001, the Department of Commerce’s National Institute of Standards and Technology (NIST) became heavily involved, and subsequently commented over 200 times between 2000 and 2005, with the majority coming in 2003-04 leading up to the publication of the revised DNSSEC standard documents.²⁵

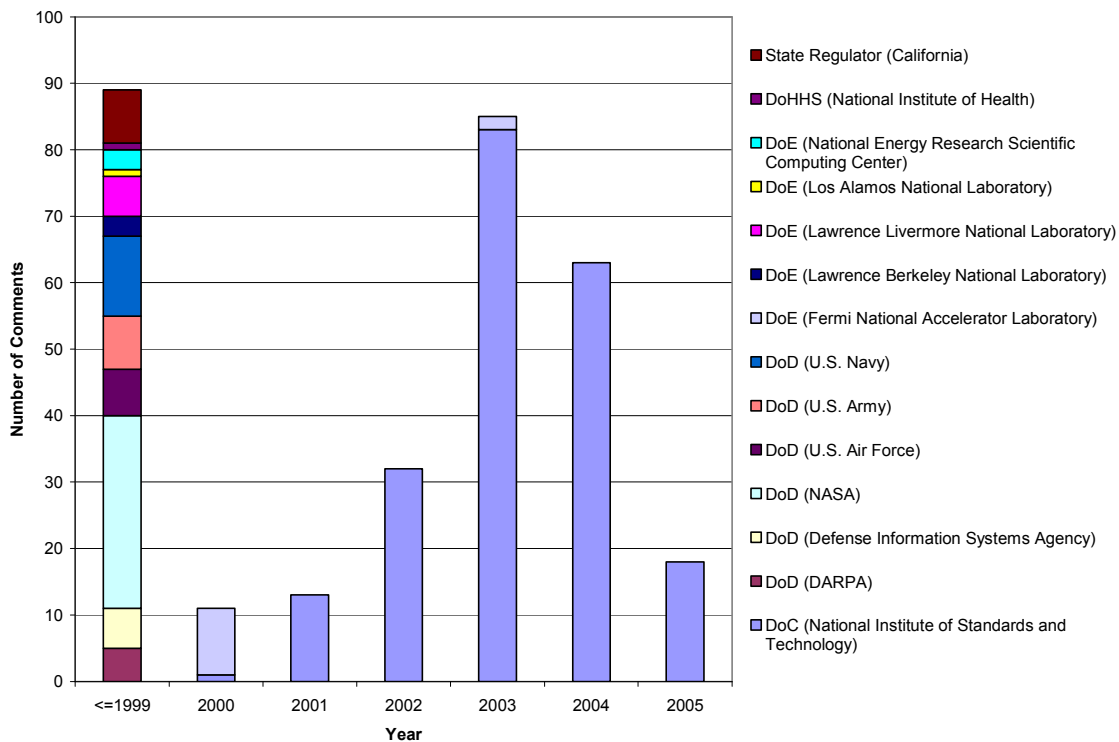


Figure 2: USG activity on *namedroppers*

²⁴ As an organization with a globally distributed membership, much of the IETF standards development process takes place via its Working Group listservs.

²⁵ Somewhat confusingly, while *namedroppers* is the official IETF DNSEXT WG list, other important discussions took place elsewhere including *cafax.se* (<http://www.cafax.se/dnssec/maillist/>) and *keydist* (<http://www.cafax.se/keydist/maillist/>) lists. Nonetheless, if one accounts for government actor participation on those lists the trend described does not change.

While NIST's participation on *namedroppers* dramatically increased, its impact was more likely felt on the IETF DNSSEC editors' team, where NIST took the self-acknowledged lead "in the completion and progression of all core DNSSEC specifications (NIST 2004 pg. 36)."²⁶ By 2004, NIST had also drafted guidelines sponsored by the Homeland Security Advanced Research Projects Agency (the funding arm within the Science and Technology Directorate of the DHS), which covered deployment of DNSSEC including key management policies.²⁷ NIST's support extended to deployment aspects of DNSSEC as well, providing support along with ICANN to form the DNSSEC deployment working group in 2004 (Jones, et al. 2005 pg. 2) and hosting DNSSEC implementation monitoring tools. As Friedlander et al. (Friedlander, A., Crocker, S., Mankin, A., Maughan, W. D., & Montgomery, D. 2005 pg. 8-9) suggests, at one level the activity of NIST was a direct response to the 2002 Federal Information Management Act (FISMA) which required federal agencies to provide risk-based information security protections for internal information systems, as well as ones used on behalf of an agency (US General Accounting Office 2004 pg. 21). However, at another level it shows a determined effort by NIST to influence the development of the "new" DNSSEC standard to reflect the interests of the government.

U.S. Department of Defense

The Department of Defense (DoD) has a long history of activity concerning DNSSEC in both research and supporting operations. Much like the early days of the Internet, perhaps the most important portion of DoD's activity in DNSSEC was the sponsorship of research through its Defense Advanced Research Projects Agency (DARPA). DARPA provided research funding in the mid to late 1990s to Trusted Information Systems (TIS) and ISC to provide expertise in the specification of the standard and prototype implementation.²⁸ Between 2001 and 2005, DARPA's heavily supported Computing Systems & Communications Technology Program budgeted over

²⁶ The DNSEXT Working Group is one of three areas where NIST has been active; the other two are BGP and IPSEC.

²⁷ The final document was released in 2005, see Chandramouli & Rose (2005), *Special Publication 800-81: Secure Domain Name System (DNS) Deployment Guide*

²⁸ See *NETWORK ASSOCIATES SELECTED TO DEVELOP NEW INTERNET SECURITY STANDARD*. (1998). Retrieved September 10, 2006, from <http://www.landfield.com/isn/mail-archive/1998/Aug/0117.html>

\$350 million in funding for its Information Assurance and Survivability Project (ST-24).²⁹ As part of ST-24, the Fault Tolerant Networks (FTN) program developed secure enhancements to the DNS, including operational use of keys, incremental deployment of secure protocols and strategies for coping with faulty or malicious secured zones. Additionally, the FTN planned to “transition these technologies to critical information and telecommunications systems...essential for minimum network operations, through its Critical Infrastructure Program. (DARPA 2002, pg. 81-83)” In outlining plans for 2002, FTN perhaps foreshadowed the coming controversy about signing the root, as it planned to demonstrate the technology necessary to construct “a dynamic, topologically sensitive root context for any network topology, thus, removing the dependence of a single, fixed root content for the domain name server (DNS). (84)”

DoD has also clearly supported DNSSEC implementation early on in the .mil and .sml domains and sub-domains to support its operational activities (Defense Science Board Task Force on Defensive Information Operations 2001 pg. 35). Not surprisingly, defense related industry argued that the DoD’s reliance on private networks made it necessary to “secure the basic protocols that support the Internet,” including adopting DNSSEC in the .mil zone (U.S. House of Representatives Committee on Science 2005b). However, in addition to securing its own communication networks, it has long been interested in developing capabilities and a national policy on offensive cyber operations (e.g., psyops, spoofing of “enemy” websites) as information operations on the Internet grow in importance (Department of Defense 1998; 2003 pg. 50-51).

National Science Foundation

Not all government research support for DNSSEC was funneled through DARPA. As one of the nation’s premiere supporters of basic research, the National Science Foundation (NSF) has maintained a long relationship with institutions engaged in the Internet standards process. This includes building institutional capacity to support standards development and supporting work which anticipated influencing standards developed through IETF processes. As Figure 3 (below) indicates, much of the funding

²⁹ See <http://www.darpa.mil/body/pdf/FY03BudEst.pdf>

during the mid 1990s went toward work investigating the impact of IPv6 addressing, and transport protocols to deal with congestion and related quality of service problems.




Figure 3: NSF Funding for Research Related to IETF Standardization Activities³⁰

Funding since that time has dramatically increased and focused additionally on underlying routing and naming standards including Border Gateway Protocol (BGP) and DNSSEC. The connections between research and the IETF standards process also became more explicit. For instance, a 2004 grant to USC-ISI (the largest organizational recipient of funding overall) would provide “systematic assessment of the gap between the DNSSEC specification and the deployment constraints” in order to achieve “unified design improvement” of the standard.³¹ In 2004, NSF funded over \$3 million towards a collaborative effort between UCSD’s CAIDA project and ISC, makers of the BIND software, to examine the DNS infrastructure generally including the impact of DNSSEC

³⁰ See Appendix A for a list of NSF awards and classifications. For clarity, grants related to platform interoperability (e.g., wireless, middleware) were removed as they do not address specific IETF Areas or standards.

³¹ See <http://www.nsf.gov/awardsearch/showAward.do?AwardNumber=0524172>

on the Internet traffic, resolver behavior and root server performance with the goal of contributing to “the hardening and securing DNS over time.”³²

The KEY debate

It was recognized early on that DNS Security Extensions were relevant to more than just securing the DNS, “because they provide a possible method for distributing public keys for use in authenticating the key components that are essential in [communication exchanges using public-key cryptography] (Orman 1995).” Within the Working Group, these discussions began in 1993, and the earlier versions of the DNSSEC standard specified this capability for a generalized KEY RR, noting that “storage of keys can support general public key distribution services as well as DNS security.”³³ As defined at the time, the KEY RR contained a protocol type field which allowed differentiating between public key data for applications and protocols like PGP, SSL, IPSEC, and DNSSEC.

However, RFC 3445, introduced as a draft in November 2001 and published as an RFC in December 2002, proposed to eliminate this functionality by removing the storage of arbitrary public keys from the DNS.³⁴ The proposed “key-restrict” changes ignited a heated debate within the Working Group discussions about the purpose of the DNS. It pitted DNS infrastructure traditionalists, who viewed the purpose of the KEY RR as restricted to naming data, against the “appdev” crowd who sought to leverage a secured DNS to provide an authenticable distribution of other application and protocol data, i.e., public keys or certificates.

The authors, the Principal Investigator and IETF lead for the DNSSEC Project at the National Institute of Standards (NIST) and a researcher formerly at a government

³² See <http://www.nsf.gov/awardsearch/showAward.do?AwardNumber=0427144>

³³ See Eastlake, D. (1999). *RFC 2535: Domain Name System Security Extensions*, from <http://www.ietf.org/rfc/rfc2535.txt>

³⁴ See Massey, D., & Rose, S. (2002). *RFC 3445: Limiting the Scope of the KEY Resource Record (RR)*. Available at <http://ietf.org/rfc/rfc3445.txt> and the documents evolution here <http://tools.ietf.org/wg/dnsext/draft-ietf-dnsext-restrict-key-for-dnssec/> It should be noted that despite the authors strong recommendation “prohibiting storing application keys in the KEY resource record, it did not endorse or restrict the storing application keys in other record types. (6)”

funded research center (i.e., USC/ISI), reasoned that “application keys, such as PGP/email, IPSEC, TLS, and SSH keys, are not a mandatory part of any zone and the purpose and proper use of application keys is outside the scope of DNS.”³⁵ For them, DNSSEC and application keys differed in purpose and use and therefore the later added “unnecessary complexity and [increase] the potential for implementation and deployment errors...[and that]...limited experimental deployment has shown that application keys stored in KEY RRs are problematic.” Their main arguments for exclusion of other keys from the RR set centered on two issues, sub-typing and blind-signing of keys unrelated to securing the DNS.

“Sub-typing,” described the problem associated with the use of the RR header to indicate the presence of keys used for things other than DNS authentication. It was argued that additional overhead for resolver software and DNS hardware was created, as the resolver would have to sort through a variety of key “types” stored in a RR. Furthermore, it was suggested that each application which desired a public key in the DNS should require an additional RR set, so as to clearly delineate its purpose to resolver software and systems administrators.³⁶ It was also argued that under the original specifications the parent zone would also be “blind signing” varying application keys stored in the KEY RR. As outlined by an engineer for TIS Labs, “placing application keys in DNS and then relying on the DNS Security Extensions to provide security of the keys places an extra burden on the DNS administrator.”³⁷ In effect, the administrator would be responsible to a higher degree for the contents of a zone, vouching for KEY RR authenticity, and thereby potentially exposing itself to liability claims (Arends and Koch 2005).

The level of intensity for restricting the KEY RR varied among other participants in the Working Group list discussions. The co-chair of the DNSOP Working Group, employed by ISC, agreed with the technical arguments made about restricting the KEY

³⁵ See <http://tools.ietf.org/wg/dnsext/draft-ietf-dnsext-restrict-key-for-dnssec/draft-ietf-dnsext-restrict-key-for-dnssec-00.txt>

³⁶ See Rose, S. (2002). *Message 01259: Re: [Design] Re: the KEY debate*, from <http://ops.ietf.org/lists/namedroppers/namedroppers.2002/msg01259.html>

³⁷ See <http://www.cafax.se/keydist/maillist/2001-12/msg00038.html>

RR, but admitted he was “agnostic” about public key data being stored in the DNS under a different RR type.³⁸ To the contrary, a principle scientist for VeriSign argued that the “problem is that DNS is designed operated and deployed as a name infrastructure. The majority of the DNS servers are deployed without any form of security analysis.” Therefore it was a “very bad idea to make a change in a deployed protocol that changes its status from a non-security infrastructure to a critical security infrastructure [for distributing public keys].” Not surprisingly, it was suggested that another incumbent standard (XKMS) provided a key centric PKI and was already supported by vendors and companies (including VeriSign) interested in providing Web Services Security.³⁹

In many ways the KEY debate was like other battles over technical resources when subjected to economic and policy pressures – proposed changes are often met with arguments claiming the technical impossibility, or costliness and undesirable nature of a proposed change (Mueller 2006). Often this is enough to stymie changes in the resource, for example, the persistence of the US spectrum regime prior to introduction of market driven transferable rights and auctions, or the artificial limitation of the TLD name space (see Mueller & McKnight 2004). While the KEY debate raised valid concerns about data appropriateness, resource constraints and legal responsibility, the proposed change also reflected a re-alignment of the DNSSEC standard with the changing composition and interests of the Working Group members.

Importantly, the change shifted elsewhere the costs of defining and stewarding thru the IETF process proposed RR sets for public key storage by other applications and protocols. For instance, efforts were undertaken outside of the Working Group to define the CERT and APPKEY RRs. The burden moved away from the well coordinated and highly influential organizations involved in the DNSSEC Working Group (WG), including USG agencies, registries, DNS managers and BIND software developers to loosely organized and less powerful developers of secure public-key cryptography applications like PGP and S/MIME, which have been hampered by the lack of a

³⁸ See <http://ops.ietf.org/lists/namedroppers/namedroppers.2002/msg01225.html>

³⁹ See <http://ops.ietf.org/lists/namedroppers/namedroppers.2002/msg01286.html> and <http://ops.ietf.org/lists/namedroppers/namedroppers.2002/msg01349.html>

widespread, convenient, and authenticable distribution infrastructure for public keys (Bellovin 1998 pg. 47; Jones et al. 2005 pg. 1).⁴⁰ One non-US developer lamented that absent a WG decided approach for application keying material in the DNS, each application would have to design it separately,⁴¹ and another more stridently regarded the choice of the WG “as a denial of service attack on the effort to secure the Internet.”⁴² Arguably, by not permitting additional applications to leverage an easily accessible, authenticable DNS and associated network effects for public key distribution, widespread encryption use in Internet applications was inhibited and the costs associated with surveillance activities reduced.

Leveraging the root

Given the hierarchical model of trust specified in the standard, one could reasonably assume that successful deployment of DNSSEC was contingent on the signing and maintenance of DNSSEC relevant information in the root zone file. However, from the perspective of the DNSEXT Working Group, the likelihood of the root remaining unsigned has been known for some time. For example, in debating how resolvers should be configured to handle DNSSEC, representatives of the largest registry and DNS management companies (also a former DARPA Program Manager) argued in 2003 that “there will not be one true [DNSSEC enabled] root for a long time to come and a resolver needs to be able to configure multiple trust anchors at different points in the tree.”⁴³ This change introduced additional complexity to the deployment of DNSSEC, namely the inability of resolver software to validate TLD zone files absent the inclusion of 1) a preconfigured, trusted public key(s) in resolver software, or 2) some other mechanism for establishing trust outside of the complete delegation hierarchy. The first scenario is

⁴⁰ Public-key cryptography vastly simplified the process of encrypted communications, making it far more accessible and affordable. This was evident from the rapid dissemination of cryptographic software following the June 2000 lifting of most U.S. restrictions on the export of encryption technology whether in source or object code (Diffie and Landau 2002).

⁴¹ See <http://ops.ietf.org/lists/namedroppers/namedroppers.2002/msg01258.html> In fact, this has occurred with the publication of Josefsson, S. (2006). *RFC 4398: Storing Certificates in the Domain Name System (DNS)*. Available at <http://tools.ietf.org/wg/dnsext/draft-ietf-dnsext-rfc2538bis/rfc4398.txt>. This work was introduced into the IETF standards process in 2005.

⁴² See Richardson, M. (2002). *Message 01224: the KEY debate*. Available at <http://ops.ietf.org/lists/namedroppers/namedroppers.2002/msg01224.html>

⁴³ See <http://ops.ietf.org/lists/namedroppers/namedroppers.2003/msg01492.html>

considered unreasonable due to the large number of DNS zone keys required, which is only complicated by the scenario of having to periodically change keys (Massey, et al. 2001). Solutions to establish alternative trust anchors have been suggested. For example, ISC's sponsored DNSSEC Look-aside Validation (DLV) proposes that zone owners generate and sign their zones normally, and register their public keys with cooperating third party registries, which would allow other participating zones configured to participate in DLV to validate the key (Vixie 2004 pg. 1330). Of course, this solution is also dependent (and leverages) on the widespread deployment of newer versions of BIND which incorporates the logic.

One can only speculate how the Working Group within the ostensibly private-sector led IETF was so prescient in understanding the political and economic difficulties of signing the root zone file. However, the insight seems to be proving true, as there has been little indication to date that ICANN intends to build a trust anchor into the root of the Internet. While ICANN's recently negotiated VeriSign settlement dealt with root server management and DNSSEC deployment, it in fact, eliminated a previously specified completion date for enabling ICANN to edit, sign and publish the root zone.⁴⁴ Furthermore, if ICANN is willing to sign the root, there are other suggestions that it is willing to use its position to affect how DNSSEC is deployed.

First, in Marrakech the .de ccTLD operator raised its frustration during the DNSSEC workshops with the complete lack of process definition and transparency in regard to key management policies at the root.⁴⁵ In particular, concerns revolved around which party (i.e., IANA, DoC, or VeriSign) actually controls the root zone signing key, which in effect grants them the ability to refuse to sign data sent (i.e., DS records) sent for inclusion in the root. While speculative, one can hypothesize that refusal to sign a DS RR sent up from a zone would arise from the nature of the public key content, e.g., the algorithm used to create a zone's DNSKEY. And from observations of other decisions about controversial content in the root, we know that the process of content entry into the

⁴⁴ See <http://icann.org/topics/vrsn-settlement/revised-root-transition-agreement-redline-29jan06.pdf>

⁴⁵ See <http://www.icann.org/meetings/marrakech/captioning-dnssec-28jun06.htm>

root can be entirely subject to influence from politically driven interests.⁴⁶ Arguably the level of influence interests internal to the government (e.g., law enforcement and surveillance) could be much greater, and less transparent.

Second, although debate over the recently released proposed registry contracts for .info, .biz and .org has been consumed with variable pricing issues,⁴⁷ they contain other interesting provisions related to DNSSEC. Section 3.1 calls for the escrowing of public and private key signing and zone signing keys with mutually agreed upon third parties. Apparently introduced as language to ensure business continuity in the event of the registry collapse, it caught even the DNSSEC Deployment group by surprise, which questioned its purpose and associated tradeoffs.⁴⁸ This episode highlights the need for root authorities (whether ICANN or the TLDs) to be transparent in developing adequate operational procedures and key management policies. While a goal of providing business continuity in a critical communications system is perfectly reasonable, any trusted directory (which is what the DNS becomes with DNSSEC) that is required to make its encryption keys available to another party invites immediate comparison to early 1990s efforts by the USG to design “back doors” into infrastructure.⁴⁹ Again speculative, but key escrow would certainly make it possible for parties with access to private keys to forge DNS data. As Denning (1996) argues, large scale deployment of key escrow in globally available infrastructure introduces risk, and is not likely to be accepted “unless users are convinced that these risks have been made negligible through technical, legal, and procedural safeguards. At this point, these details are unknown, as is who or what is driving ICANN’s key escrowing requirements stated in the proposed contracts. But with the known long history of USG involvement in encryption standards to facilitate surveillance and the growing importance of information operations in communication

⁴⁶ Most recently, ICANN’s .xxx decision, where Department of Commerce officials instructed ICANN officials to not add the TLD as a result of political pressure from domestic interest groups, see Internet Governance Project. (2006). *Review of Documents Released under the Freedom of Information Act in the .XXX Case*. Retrieved September 7, 2006, from <http://internetgovernance.org/pdf/xxx-foia.pdf>

⁴⁷ See <http://icann.org/announcements/comments-summary-07sep06.htm>

⁴⁸ See <http://mail.shinkuro.com:8100/Lists/dnssec-deployment/Message/451.html?Language=>

⁴⁹ The infamous Clipper Chip embedded tamper resistant integrated circuits with a crackable algorithm into commercial communications devices, allowing USG authorities to exploit intercepted traffic when necessary (Diffie and Landau 1999 pg. 211; see also Denning 1996).

intensive world, understanding the motivations of the USG in DNSSEC development and adoption is worth exploring.

While it is possible to speculate as to internal government interests in shaping how DNSSEC is deployed, one should not lose sight of the substantial economic benefit to ICANN in refusing to sign the root. DNSSEC is a complex technology, and by no means mature and stable. All involved parties are familiar with the liability issues imposed on root certificate authorities in PKIs and many of these same liabilities are applicable within the model proposed by DNSSEC.⁵⁰ For those zones considered trust anchors, there are risks associated with misrepresentation of entities in RRs and the cost of verifying them before being added to the zone files; there are risks of data loss and private key compromise, which in turn requires development of adequate procedures particularly in regard to key rollovers. By not signing the root, and instead using contractually bound registries to serve as “islands of trust,” ICANN’s mitigates these liabilities. Overall, not signing the RZF itself pushes many of the costs of running DNSSEC successfully toward TLD registries. The strategy of shifting costs became even more apparent with the recent launch of the Registry Request Service and formation of the Registry Services Technical Evaluation Panel (RSTEP).⁵¹ In one regard, this can simply be viewed as an (over)response to the SiteFinder fiasco.⁵² However, as pointed out by a DNSEXT Working Group participant, the creation of this gate by ICANN places an additional burden on gTLDs wishing to implement DNSSEC.⁵³ The documentation requirements for gTLDs are onerous and somewhat ambiguous not knowing the status of whether the root will be signed, and the review process potentially open-ended.

Finally, the previous three years of global debate at the World Summit on the Information Society and the creation of the Internet Governance Forum highlight the widespread discontent with USG unilateral control over the root zone file. The

⁵⁰ See e.g., Baum, M. S. (1994). *Federal Certification Authority Liability and Policy*. Washington DC: National Institute of Standards and Technology available at www.verisign.com/repository/pubs/fca_liability.pdf

⁵¹ See <https://rrs.icann.org/>

⁵² See <http://www.icann.org/announcements/advisory-19sep03.htm>

⁵³ See http://www.circleid.com/posts/dnssec_to_gtld_administrations_by_icann/

introduction of DNSSEC to the DNS will only magnify this dilemma. The USG is in a tricky position concerning DNSSEC, as it struggles to balance contending security, economic, and foreign policy interests, and it remains to be seen what political bargains both internally and globally will be crafted.

References

- Arends, R., & Koch, P. (2005). *DNS for Fun and Profit*. Paper presented at the DFN-CERT: Workshop on "Sicherheit in vernetzten Systemen".
- Ateniese, G., & Mangard, S. (2001). *A new approach to DNS security (DNSSEC)*. Paper presented at the Proceedings of the 8th ACM conference on Computer and Communications Security, Philadelphia, PA, USA.
- Berger, D. F. (2004). *A Scalable Architecture for Public Key Distribution Acting in Concert with Secure DNS*. Unpublished Masters thesis, University of California-Riverside.
- Carney, D. (2003). *DHS and NIST to Collaborate*. Retrieved November 30, 2005, from <http://www.techlawjournal.com/topstories/2003/20030522.asp>
- Chandramouli, R., & Rose, S. (2005). *Special Publication 800-81: Secure Domain Name System (DNS) Deployment Guide*. Washington DC: National Institute of Standards and Technology.
- Crane, R. J. (1979). *The Politics of International Standards: France and the Color TV War*. Norwood, NJ: Ablex Publishing Corporation.
- David, P., & Greenstein, S. M. (1990). The Economics of Compatibility Standards: An Introduction to Recent Research. *Economics of Innovation and New Technology*, 1(2), 3-41.
- Defense Science Board Task Force on Defensive Information Operations. (2001). *Protecting the Homeland: Defensive Information Operations Volume II*. Washington D.C.: Office of the Undersecretary of Defense for Acquisition, Technology, and Logistics.
- Denning, D. E. (1996). Encrypting the Global Information Infrastructure. In *Computer Fraud & Security*: Elsevier Science Ltd.
- Department of Defense. (1998). *Joint Publication 3-13, Joint Doctrine for Information Operations*. Washington D.C.: Government Printing Office.
- Department of Defense. (2003). *Information Operations Roadmap*. Retrieved September 2, 2006, from <http://www.gwu.edu/~nsarchiv/NSAEBB/NSAEBB177/>
- Diffie, W., & Hellman, M. (1976). New Directions in Cryptography. *IEEE Transactions on Information Theory*, IT-22(6), 644-654.
- Diffie, W., & Landau, S. (1998). *Privacy on the Line*. Cambridge, Mass.: MIT Press.

Diffie, W., & Landau, S. (Forthcoming). The Export of Cryptography in the 20th Century and the 21st. In K. D. Leeuw (Ed.), *Handbook of the History of Information Security*. Amsterdam: Elsevier.

Eastlake, D. (1999). *RFC 2535: Domain Name System Security Extensions*, from <http://www.ietf.org/rfc/rfc2535.txt>

Faulhaber, G. R. (2005). Bottlenecks and Bandwagons: Access Policy in the New Telecommunications. In S. K. Majumdar, I. Vogelsang & M. Cave (Eds.), *Handbook of Telecommunications Economics* (Vol. 2, pp. 488-517). Amsterdam: Elsevier.

Froomkin, A. M., & Lemley, M. A. (2003). ICANN and Antitrust. *University of Illinois Law Review*(1).

Galperin, H. (2004). *New TV, Old Politics: The Transition to Digital TV in the U.S. and Britain*. New York: Cambridge University Press.

Giacomello, G. (2005). Sometimes security just does not prevail: The case of the 'cryptowars'. In *National Governments and Control of the Internet: A Digital Challenge* (pp. 26-55). New York: Routledge.

Gilmore, J. (2000). *DNSSEC history, and government's role in designing infrastructure*. Retrieved December 19, 2005, from <http://www.chiark.greenend.org.uk/pipermail/ukcrypto/2000-February/047193.html>

Greenstein, S. M. (1992). Invisible Hands and Visible Advisors: An Economic Interpretation of Standardization. *Journal of the American Society for Information Science*, 43(8), 538.

Gutman, P. (2002). PKI: It's Not Dead, Just Resting. *IEEE Computer Magazine*, 35(8), 41-49.

Hart, J. A. (2004). *Technology, Television, and Competition: The Politics of Digital TV*. Cambridge: Cambridge University Press.

Huston, G. (2006). *DNSSEC - The Theory*. Retrieved August 22, 2006, from <http://ispcolumn.isoc.org/2006-08/dnssec.html>

Josefsson, S. (2006). *RFC 4398: Storing Certificates in the Domain Name System (DNS)*. Retrieved September 15, 2006, from <http://tools.ietf.org/wg/dnsext/draft-ietf-dnsext-rfc2538bis/rfc4398.txt>

Jones, J. P., Berger, D. F., & Ravishankar, C. V. (2005). *Layering Public Key Distribution Over Secure DNS using Authenticated Delegation*. Paper presented at the Annual Computer Security Applications Conference, Tucson, AZ.

- Kavanaugh, A. (1985). Who Determined U.S. Satellite Policy on INTELSAT? *Journal of Communication*, 35(3), 70-79.
- Kohnfelder, L. (1978). *Toward a Practical Public-Cryptosystem*. Unpublished Bachelor's thesis, MIT, Cambridge, MA.
- Larsen, P. B. (2001). Issues relating to civilian and military dual uses of GNSS. *Space Policy*, 17(2), 111-119.
- Lembke, J. (2002a). EU Critical Infrastructure and Security Policy: Capabilities, Strategies and Vulnerabilities. *Current Politics and Economics in Europe*, 11(2), 99-129.
- Lembke, J. (2002b). *Competition for Technological Leadership: EU Policy for High Technology*: Edward Elgar Pub.
- Massey, D., Lewis, E., Gudmundsson, O., Mundy, R., & Mankin, A. (2001). *Public Key Validation for the DNS Security Extensions*. Paper presented at the DARPA Information Survivability Conference and Exposition (DISCEX II).
- Mathiason, J., Mueller, M., Klein, H., Holitchser, M., & McKnight, L. (2004). *Internet Governance: The State of Play*. Syracuse, NY: The Internet Governance Project.
- Mueller, M. (1991). *Intelsat and the Separate System Policy: Toward Competitive International Telecommunications* (No. 150). Washington DC: The Cato Institute.
- Mueller, M. (2006). IP addressing: the next frontier of internet governance debate. *info*, 8(5), 3-12.
- Mueller, M., & McKnight, L. W. (2003). *The Post-.Com Internet: Five Steps To an Open DNS*, from <http://dcc.syr.edu/miscarticles/NewTLDs2-MM-LM.pdf>
- National Institute of Standards and Technology. (2004). *Annual Report*. Washington DC: Computer Security Division.
- National Research Council. (2005). *Signposts in Cyberspace: The Domain Name System and Internet Navigation*. Washington D.C.: National Academies Press.
- Pace, S. (1996). The global positioning system: policy issues for an information technology. *Space Policy*, 12(4), 265-275.
- Pelkmans, J. (2001). The GSM standard: explaining a success story. *Journal of European Public Policy*, 8(3), 432-453.
- Perlman, R. (1999). An overview of PKI trust models. *Network, IEEE*, 13(6), 38-43.

Snow, M. S. (1976). *International Commercial Satellite Communications: Economic and Political Issues of the First Decade of INTELSAT*: Praeger.

Snow, M. S. (1985). Arguments For and Against Competition in International Satellite Facilities and Services: A U.S. Perspective. *Journal of Communication*, 35(3), 51-79.

Tien, L. (1999). *WRITTEN SUBMISSION OF MR. HUGH DANIEL IN THE MATTER OF APL9800007/Z066051/G006298*. Retrieved September 18, 2006, from <http://www.toad.com/dnssec/19990223-written-submission>

United States General Accounting Office. (2001). *Information Security: Advances and Remaining Challenges to Adoption of Public Key Infrastructure Technology* (No. GAO-01-277). Washington D.C.: United States General Accounting Office.

United States General Accounting Office. (2004). *Technology Assessment: Cybersecurity for Critical Infrastructure Protection* (No. GAO-04-321). Washington D.C.: United States General Accounting Office.

VeriSign. (2002). *Annual Report on Form 10-K*. Retrieved February 14, 2005, from http://verisign.com/stellent/groups/public/documents/annual_report/verisign2002annualreportview.pdf

VeriSign. (2004). *Annual Report on Form 10-K*. Retrieved February 14, 2005, from http://verisign.com/stellent/groups/public/documents/annual_report/verisign2002annualreportview.pdf

Weiss, M. B., & Sirbu, M. (1990). Technological Choice in Voluntary Standards Committees: An Empirical Analysis. *Economics of Innovation and New Technology*, 1, 111-