

DNSSEC and Hardening Security in the Internet Infrastructure: The Public Policy Questions

Amy Friedlander
Stephen Crocker
Allison Mankin
W. Douglas Maughan
Douglas Montgomery

Revised August 30, 2005

This is a paper from the practitioner community. We are engaged in an effort to strengthen security in the Internet infrastructure. Our immediate task is to deploy a new Internet protocol, DNS Security Extensions (DNSSEC), which promises to harden features of the Domain Name System (DNS), a key element in the infrastructure of the Internet. In our work, we find ourselves at the intersection of the following questions:

1. How do we stimulate innovation in infrastructure services when those services are provided in a competitive, largely private commercial environment and the returns are likely to occur in the long term and will also be shared?
2. What is the appropriate role of government in fostering infrastructure development when we are committed to largely privately-owned and operated infrastructure facilities and services?
3. What is the balance among national and homeland security interests and global Internet management – or governance?

We are most interested in strategies to create conditions conducive to deployment, the first question, but recognize that the three questions are interdependent since one avenue for diffusing innovation might be a blunt-edged government mandate. That strategy is neither viable nor preferred, since the infrastructure in question -- the Internet -- is global, decentralized and locally operated but given coherence by the consensus-based protocols, and predominantly owned by private interests.

Privatization of infrastructure development and services excited substantial interest in the 1980s and 1990s, spurred in part by the divestiture of AT&T as well as by the emerging information infrastructure, typically associated with the expansion and commercialization of the Internet. Infrastructure possesses substantial national public importance as a range of actions demonstrate, from the earliest national highway systems, proposed as a defense measure in 1938, to the formation of the National Communications Systems, formed in the wake of the Cuban Missile crisis in 1962, to the Critical Infrastructure Protection program of today. There is always a fundamental tension between providing for the public good and doing so through competition. Moreover, the notion of infrastructure embodies several notions: cooperation among entities to provide consistent if not uniform service, economic advantage both to providers and to the consumers of these services, and competition among the providers themselves. In short, we want

infrastructure providers to share and to compete. And finally, infrastructure providers are well aware of the perils of both moral hazard and free riders.

So let's start with a simple question: What is the DNSSEC protocol? And then proceed with a discussion of how deployment is proceeding and the issues, questions, and policies the protocol and its deployment provoke.

1. What is DNSSEC? Why is it needed? And how does it work?

Sometimes compared to a telephone directory for the Internet, DNS maps easily remembered names, like `tprc.org`, to numeric Internet Protocol (IP) addresses. It is a distributed database of information (or records) in a hierarchical organization, typically described as either a “tree” or a series of “parent-child” relationships, starting with the root, through the top level domains (TLDs), domains, and sub-domains. The system is vulnerable to a particularly insidious form of attack in which an attacker tampers with a look-up, that is a request for DNS information from one system (the resolver) to a system where that information is stored (the name server), and the compromised information is returned to the user undetected. The attack can be directed at an intermediate point called a “cache”, where DNS information is stored, so both the owner of the name and the requester may be following good security practices, but the transaction is still compromised. As a result, users can be inadvertently diverted to a site that may look perfectly respectable or their e-mail messages can be observed or read en route.

DNSSEC attaches cryptographic signatures to records in the DNS by employing public and private key pairs. As the names suggest, a “public key” is made broadly available; the “private key” is kept private, typically stored offline in a secure location and used by the name server, the program where the DNS data are stored, to encrypt the records’ signature. The point of these encrypted signatures is to allow a querying/receiving system in look-up (the resolver) to use the public key to validate the information (i.e., to answer the question is the data received the same as the data sent?) and to authenticate the signature (i.e., to answer the questions, did the correct source -- the authoritative name server -- send me the data and is the information I received the same information that was sent?). Deploying DNSSEC allows DNSSEC-enabled systems to detect compromised information. It does not stop or prevent the attack per se, but allows systems to identify DNS data that may be untrustworthy. It also allows systems to determine that a name does not exist.¹

¹ The protocol is embodied in three documents: R. Arends, R. Austein, M. Larson, D. Massey, S. Rose, DNS Security Introduction and Requirements, RFC 4033, October 10, 2004 (<http://www.ietf.org/rfc/rfc4033.txt>, verified June 21, 2005); R. Arends, R. Austein, M. Larson, D. Massey, S. Rose, Resource Records for the DNS Security Extensions, RFC 4034, October 10, 2004 (<http://www.ietf.org/rfc/rfc4034.txt>, verified June 21, 2005); R. Arends, R. Austein, M. Larson, D. Massey, S. Rose, Protocol Modifications for the DNS Security Extensions, RFC 4045, October 10, 2004 (<http://www.ietf.org/rfc/rfc4045.txt>, verified June 21, 2005). A brief introduction will be available in Amy Friedlander et al., DNSSEC: A Piece of the Internet Infrastructure Security Puzzle. Accepted for publication, *Communications of the ACM*, forthcoming.

A domain or a subset of the domain called a “zone” is considered DNSSEC enabled when it signs its records. In principle, when DNSSEC is fully deployed, .edu will vouch for anyuniversity.edu, which will, in turn, vouch for lawschool.anyuniversity.edu, and so on. At the very top of the DNS hierarchy, the root will vouch for the top level domains. Top level domains (TLDs) are classified into two major groups: the gTLDs, which are the familiar .com, .org, .edu, .info, and so on; and the ccTLDs, which are the country codes: .uk, .se., .jp, .au, and so on.

There are a lot of names and addresses out there, which are administered by a lot of institutions, using a lot of software. As of July 2005, there are roughly 53,328,4187 unique host names spread across 83,336,109 domains just in the first three levels of the DNS hierarchy.² Moreover, these systems exist under a variety of national jurisdictions and legal and regulatory regimes. And systems in the private sector are also operated under an array of organizational conditions with varying missions, incentive structures, and responsibilities. Thus, a university might be motivated by the prestige associated with early adoption while a private corporation must allocate its internal resources among a series of competing budgetary demands while answering to a board of directors and ultimately to stockholders.

So interesting questions arise concerning who decides and how adoption should be motivated? The software is not free. The human investment in terms of time and training may be substantial. And the effect on performance – a major question for any commercial entity – remains to be seen. Indeed, the technical community is working on this issue, among others.³ Additionally, there is the potential for disruption during the switchover. The migration to DNSSEC will, in fact, be incremental⁴ so new and legacy systems have to coexist, leading to more complexity and hence cost. And finally DNSSEC is not a silver bullet.⁵ It addresses an important set of threats to the Internet but not all of them. And because it is an innovation to the infrastructure, which is by definition shared, no single provider can be sure of recovering costs and the likelihood of the free rider problem is all too real.

² Internet Systems Consortium, ISC Internet Domain Names Survey (July 2005), <http://www.isc.org/index.pl?ops/ds/>.

³ Bernhard Ager, Holger Dreger, and Anja Feldman, Exploring the Overhead of DNSSEC, [May 17, 2005], <http://www.net.informatik.tu-muenchen.de/~anja/feldmann/papers/dnssec05.pdf>. Discussions among members of the DNSSEC Deployment Working Group on this and other technical topics are publicly available at <http://www.dnssec-deployment.org/wg.php>. Members of the working group organized a workshop on performance and metrics at the recent meeting of the IETF in Paris, August 3, 2005.

⁴ We note that the protocol has been designed to allow for incremental adoption and coexistence of DNSSEC-compliant and legacy systems.

⁵ Threats are described in D. Atkins and R. Austein, RFC 3833: Threat Analysis of the Domain Name System (DNS), August 2004. <http://www.ietf.org/rfc/rfc3833.txt>.

2. Tipping Point, Technology, and Public Policy Issues

Internet standards are proposed, debated and established within the Internet Engineering Task Force (IETF). From the perspective of the Internet's operation, adoption of a protocol is voluntary; no central authority mandates use but cooperation is induced by the nature of the network. That is to say, entities join the network to access other nodes on the network, presumably to maximize access to other nodes. Thus, while adoption of DNSSEC does not compromise legacy DNS systems, as more nodes become DNSSEC compliant, it becomes more advantageous for individual nodes to join, and at some point, the balance "tips" in favor of deploying DNSSEC.⁶ So deployment becomes a tipping point problem, and one role of government is to help create conditions that encourage the tipping point to occur sooner rather than later. How that might be done is discussed later in this article.

The technology is neither created, deployed, nor used in a cultural vacuum. Two technical issues have already provoked discussion related to Internet governance. The first concerns the root and the second the privacy implications of a possible practice known as zone walking or zone enumeration.⁷

First, the root. A DNSSEC-enabled zone, which is a subset of DNS information that is administered by a given entity, employs public-private key pairs to digitally sign its records. A DNSSEC-capable resolver, that is, the system that receives the DNS information in response to a query, uses the public key of the parent to authenticate the source of the information. When the system is fully DNSSEC enabled, all of the layers of the hierarchy will vouch for the next layer down, from the root on down. Thus a question arises: Who holds the root's keys? And indeed, who manages the root and the root's keys? In particular the *private* key belonging to the root?

⁶ Several studies have looked at the network effects on technology adoption in telephony, notably William Paul Barnett, *The Organizational Ecology of the Early Telephone Industry: A Study of the Technological Cases of Competition and Mutualism*, (Ph.D dissertation, University of California, Berkeley, 1988); Kenneth Lipartito, *The Bell System and Regional Business: The Telephone in the South, 1877-1920* (Baltimore, MD: The Johns Hopkins University Press, 1989); Milton Mueller, *The Telephone Connection: Interconnection, Competition and Monopoly in the Making of Universal Telephone Service, 1894-1920*, (Ph.D. dissertation, University of Pennsylvania, 1989). A summary that considers these studies from the perspective of natural monopoly theory, increasing returns to scale, and technology adoption is provided in Amy Friedlander, *Natural Monopoly and Universal Service: Telephones and Telegraphs in the U.S. Communications Infrastructure, 1837-1940* (Reston, VA: The Corporation for National Research Initiatives, 1995).

⁷ *Report of the Working Group on Internet Governance*, June 2005; <http://www.wgig.org/>. The background report to this document (*Background Report, The Working Group on Internet Governance*, <http://www.wgig.org/docs/BackgroundReport.doc>, verified August 2, 2005) notes the potential clash between DNSSEC and national privacy policies in a footnote, page 17, note 13. Paul Vixie, a well known expert on Internet engineering and security, has called attention to this point in his comments; see "Some Comments on Working Group on Internet Governance (WGIG)," 21 July 2005 (<http://fm.vix.com/internet/governance/wgig-report-july05.html>, verified August 2, 2005).

Moreover, on June 30, 2005, the Department of Commerce issued a statement of principles in which it acknowledged the role of other governments with respect to the ccTLDs but also announced it would take no actions that adversely affect the effective and efficient operation of the DNS and would therefore “maintain its historic role in authorizing changes or modifications to the authoritative root zone file.”⁸ Questions about managing the root, including the keys, go to major aspects of managing the Internet’s infrastructure, the extent to which management should be internationalized and under what rubric, and the role and control of ICANN. These issues have already been raised by the Working Group on Internet Governance and have begun to be reported in the technology press.⁹

Zone walking/zone enumeration. The second technical issue is equally charged, that is, zone walking or zone enumeration and privacy. In addition to the DNS data and the signatures, the resolver, the system that initiated the query and receives the information from the name server, from the name server, also receive information on the next signed entry in the DNS database. (This lets the resolver distinguish between the non-existence of a record and existence of a record for which no response arrives because its signed record currently is invalid, for instance, if an attacker has tampered with it.). Using the information on the next entry received in each response, successive queries can be launched enabling the query system to ascertain the contents of a given zone file (or subset of the DNS database under a single administration), even if the zone has a policy against such information being shared. This is called "zone walking" or "zone enumeration" and has been identified by a number of parties, more so in Europe than in the U.S., as a potential privacy threat.

Besides violating policy set by administrations, the specific threat in the harvest of zone information obtained from zone walking is not hard to envision. Aggregation of entries from zones often exposes internal information about organizations because of useful naming conventions that include locations and departments in the system names. Although this privacy risk is not viewed by all as preeminent, there is recognition in the technical community that the DNSSEC specification needs strengthening in this aspect and it is being addressed in the IETF working group framework with a technology called NSEC3.¹⁰

⁸ U.S. Principles on the Internet’s Domain Name and Addressing System, June 30, 2005, http://www.ntia.doc.gov/ntiahome/domainname/USDNSprinciples_06302005.htm.

⁹ Should we wait for CNN headlines about DNS insecurity before taking action? *News from .aero -the domain of aviation*, Newsletter No.10, June 2005, http://www.information.aero/index.php?id=183&no_cache=1.

¹⁰ The current proposal is NSEC3, a solution which allows the same checking for the non-existence of a name, versus existence of a name whose record has a security problem, but which uses a cryptographic hash method (in which the next record’s information is converted with a one-way function to an abstract string, and any matching is performed after using the same one-way function on the candidate) to avoid the next record being identifiable. For more information see the IETF’s work in progress, B. Laurie (et al), DNSSEC Hash Authenticated Denial of Existence (<http://www.ietf.org/internet-drafts/>). [We note that the versioning of Internet drafts changes and the work evolves, so readers are encouraged to search the referenced page by title and author for the current state of work.]

An additional strong privacy risk linked to zone-walking occurs in the combination of the next record information with Whois information, a point Ram Mohan of Afiliias, a global registry services company, made at the ICANN meeting in Mar del Plata in early April 2005.¹¹ How might this scenario occur? The Whois database contains information supplied by entities that register domain names. The information, some of it personal, is usually publicly available. Someone, for whatever reason, walks a zone, and uses each zone entry as a key for requesting Whois records, thus obtaining large quantities of personal information. It can be argued that it was not the responsibility of DNSSEC's developers to worry about what some entity, whether malevolent or not, would do with zone data that might later violate privacy.¹² That this arguably should be handled by Whois and the policies that govern access to this information. From a different viewpoint, the DNSSEC NSEC3 work can be seen as enhancing the prospect that the DNS and the registration database have complementary privacy attributes. In a number of registry environments, there is discussion of adopting IRIS (Internet Registry Information Service),¹³ which has been developed within the IETF to provide security and privacy-motivated access controls, among other goals, as a replacement for Whois.¹⁴

But even if the technical community resolves these issues, would issues associated with root impede progress as has been claimed? In fact, not signing the root is not necessarily an impediment.¹⁵ The protocol allows for zones lower down in the DNS tree to self-sign, so DNSSEC can be effectively colonized at systems below the root before the root itself is DNSSEC enabled. Moreover, if only the root is signed, the system overall is not secure. So while the signing root will clearly have important symbolic as well as technical value, system wide-deployment is a more complex process that requires understanding motivations and incentives of many stakeholders at many levels in the DNS hierarchy.

¹¹ Comment in the DNSSEC Mini Workshop, ICANN Meetings, Mar del Plata, Argentina, April 5, 2005 (<http://www.icann.org/meetings/mardelplata/captioning-dnssec-05apr05.htm>, verified August 21, 2005).

¹² "DNS was originally designed with the assumptions that the DNS will return the same answer to any given query regardless of who may have issued the query, and that all data in the DNS is thus visible. Accordingly, DNSSEC is not designed to provide confidentiality, access control lists, or other means of differentiating between inquiries." R. Arends et al., *DNS Security Introduction and Requirements*, October 10, 2004.

¹³ A. Newton, M. Sanz, IRIS: The Internet Registry Information Service Core Protocol, RFC 3981, January 2005 (<http://www.ietf.org/rfc/rfc3981.txt>, verified August 21, 2005); A. Newton, M. Sanz, *IRIS: A Domain Registry (dreg) Type for the Internet Registry Information System (IRIS)*, RFC 3982, January 2005, (<http://www.ietf.org/rfc/rfc3982.txt>, verified August 21, 2005).

¹⁴ L. Daigle, *WHOIS Protocol Specification*, RFC 3912, September 2004 (<http://www.ietf.org/rfc/rfc3912.txt>; verified August 21, 2005).

¹⁵ The notion that the root must be signed as a condition for all other zones to sign is one of several misconceptions. For others, see DNSSEC Deployment Panel, "Myths or Facts," May 2005, <http://www.dnssec-deployment.org/calendar.php>.

3. Parallels with Y2K

The Y2K example¹⁶ offers a useful set of lessons as to how this process might unfold. Most of us will recall that the Y2K problem arose from legacy systems that allowed only two characters for dating the year rather than four. When 1999 rolled over into 2000, 1900 and 2000 would both be represented by “00”. And the results of this confusion would have potentially disastrous results. This did not occur; critical systems were updated and brought into compliance, and in some circles, it was thought that the Y2K problem had been over-hyped.

Admittedly, the Y2K technical issue and solution could be more tightly defined than the range of threats DNS potentially faces. However, the larger point, made by David Mussington in regard to Y2K, is that deployment in both cases is both a technical and organizational/social/economic challenge, requiring global cooperation among diverse stakeholder communities. Moreover, the information infrastructure proved critical to the operation of other infrastructures, notably finance and transportation, so security of the computer/networking infrastructure is intrinsic to protecting other critical infrastructures.

In his study for the RAND Corporation, Mussington found that enterprises can be motivated to invest in infrastructure when there is a clear liability and when they can see the economic consequences of not making the investment.¹⁷ So adoption of a technology within an institution can have cascading effects within a network of suppliers and vendors. The Marconi Company, then a British corporation, famously signed an exclusive contract with the marine insurer Lloyd’s of London in 1901 that effectively required international shippers to acquire Marconi equipment to qualify for insurance.¹⁸

In the Y2K instance, cooperation among transnational multinational corporations was essential. In addition to undertaking government-to-government discussions, government involvement assisted in developing priorities, facilitated information sharing and remediation, galvanized public awareness, and overcame asymmetries and localized market failures. In some instances, economic interests of multinational corporations aligned with the national security interests of developing nations to create powerful incentives. However, individual enterprises undertook assessment and remediation efforts, and their initiative was critical.

¹⁶ The lessons learned summarized here are taken from David Mussington, *Concepts for Enhancing Critical Infrastructure Protection* (Washington, DC: RAND, 2002). Mussington asked the question, was Y2K over-hyped? We do not endorse or challenge his findings relative to that question here but take only the cogent synopsis of the motivations among the many stakeholder communities involved.

¹⁷ *Ibid.*, p. ix-x.

¹⁸ Susan L. Douglas, *Inventing American Broadcasting, 1899-1922* (Baltimore, MD: The Johns Hopkins University Press, 1987), p. 70.

We note that government also played a direct role. The Federal Acquisition Regulation (FAR) was amended to mandate that all contracts for information technology services be ready for the Y2K date change.¹⁹ In addition, companies licensed by the Securities and Exchange Commission (SEC) were required to outline their plans for dealing with the Y2K problem. Without a plan, the SEC would not renew their business licenses. The SEC did not really start enforcing the requirement until 1995 but by 1997, companies were taking it seriously.

Thus, Y2K upgrades took place as both bottom up and top down processes with government undertaking key government-to-government activities as well as providing overall coordination and in some cases credibility. In addition, requirements can effectively create a market by providing a demand for these services and creating a need to develop necessary capacities. This suggests that efforts to resolve the management of the root, which is intrinsic to a top down strategy, must occur in parallel with broad bottom up strategies. One would consist of programs to encourage (or in some situations require) key stakeholders in the public and private sectors to deploy the protocol. The second would be programs to foster a climate for in which end users demand greater security from their infrastructure service providers, much as contemporary consumers participate in environmental recycling programs or demand air bags in their automobiles.

4. Hearts, Minds, and FISMA

The Y2K experience suggests three major government roles: to lead by example, to create a market for DNSSEC-enabled services, and to promote education and awareness. Although private industry is leading the effort to ensure that the core functions of the Internet develop in a secure manner, the U.S. federal government has an important role in articulating the issues. Policy documents such the *National Strategy to Secure Cyberspace* (<http://www.whitehouse.gov/pcipb/>) explain the relevance of the mechanisms of the Internet, including DNS. Released by the White House in February 2003, the document outlines a broad strategy to engage and empower Americans to secure the portions of cyberspace that they own, operate, or control, or with which they interact as well as to support “the development of secure and robust mechanisms that will enable the Internet to support the Nation’s needs now and in the future.” Together with the *National Strategy for the Physical Protection of Critical Infrastructures and Key Assets* (<http://www.whitehouse.gov/pcipb/physical.html>), it is part of the *National Strategy for Homeland Security* ([http://www.whitehouse.gov/homeland/ book/](http://www.whitehouse.gov/homeland/book/)). Precisely because they do reflect the public policies of the U.S. government, these documents also play an educational and public awareness role as does the report by the Computer Science and Telecommunications Board of the National Research Council, *Signposts in Cyberspace: The Domain Name System and Internet Navigation* (2005), which was partially funded by the US Department of Commerce and also explains DNSSEC and its role in enhancing the security of the Internet.

¹⁹ 48 CFR 39.106. See also U.S. Government Accountability Office, *Internet Protocol Version 6; Federal Agencies Need to Plan for Transition and Manage Security Risks* (May 2005), p. 20.

The items identified in the previous paragraph target the policy community and those who follow that community. Resources for the technical cognescenti are more plentiful. The US National Institute of Standards has published the *Secure Domain Name System (DNS) Deployment Guide* (Special Publication 800-81) to assist system administrators (as well as inform policy analysts) and the NIST web site offers a number of test and measurement tools for DNS administrators. The web site, www.dnssec-deployment.org, supports the community of the developers and early adopters.²⁰

However, an aggressive campaign by government to change hearts and minds can stimulate skepticism since government attempts to sway public opinion can also engender mistrust. Still, a September 2004 survey of technology leaders by the Pew Internet & American Life Project found that respondents were worried about the “vulnerability of the internet and the likelihood of an attack on the underlying infrastructure within the next ten years”.²¹ Yet the vast majority of end-users cite threats to their privacy and their data confidentiality as their principal concerns, seemingly unaware of the infrastructure threats.²² And to date, coverage of such threats outside of the technical press is quite thin.

So there is room for improvement and a role for government in concert with others.

Government can take more explicit actions as to foster technology transfer by deploying DNSSEC on its own systems and those of its contractors. This strategy effectively creates a market for enhanced infrastructure services while offering a relatively low-risk environment in which to develop the capacity as well as the tools. Historically, the US government has supported both infrastructure capacity and a market for infrastructure services, rural electrification projects in the 1930s being a good example. In the early 20th century, the cost of extending electrical systems was believed prohibitive in thinly populated areas, and rural Americans were under-served. New Deal programs took several forms from public administration of power generation to loans to enable construction of private facilities and consumers to buy electrical equipment and devices.²³ The space program is a more recent example where a sophisticated system proposes to enable prospective vendors to identify advanced technologies developed

²⁰ The resources cited in this paragraph are all in some way supported by the US government. Other important resources, not supported by the USG, include Olaf Kolkman’s *DNSSEC HOW-TO* (April 2005), which provides a step-by-step manual for deployment, based on RIPE’s experience, and the web site, www.dnssec.net, which has extensive coverage of educational and background material as well as pointers to tools and pilot projects.

²¹ Susannah Fox, Janna Quitney Anderson, Lee Rainie, *The Future of the Internet*, (Washington, DC: Pew Internet & American Life Project, January 9, 2005), p. i.
http://www.pewinternet.org/pdfs/PIP_Future_of_Internet.pdf.

²² Jeffrey Cole et al., *The UCLA Internet Report: Surveying the Digital Future, Year 3*. (Los Angeles: UCLA Center for Communication Policy, February 2003, pp. 48-50, 54.
<http://www.digitalcenter.org/pdf/InternetReportYearThree.pdf>.

²³ D. Clayton Brown, *Electricity for Rural America: The Fight for the REA*, Contributions in Economics and Economic History No. 29, (Westport, CT: Greenwood Press, 1980).

under federal contracts that can be commercialized, provides support through workshops and small business funding, and has resulted in applications in robotics, optics, materials science, and medical devices, among others.²⁴

In the examples of Y2K and more recently IPv6, the U.S. government is not building infrastructure directly as it did through the Tennessee Valley Authority but has created conditions that compel certain segments of the producer/manufacturing community to provide new technology and in so doing offers a market and precipitates building capacity that can then be sold elsewhere. The Y2K experience has been previously described. IPv6 (Internet Protocol Version 6) has been developed to replace the current version, IPv4 (Internet Protocol, IP Version 4).²⁵ The new protocol addresses a number of problems in IPv4, such as the limited number of available IPv4 addresses, and adds some improvements. In August 2005, the Office of Management and Budget (OMB) issued a formal memorandum, setting June 28, 2006 as the date by which agencies' infrastructure must be using IPv6 and "agency networks must interface with this infrastructure," thus inserting the new technology into the Internet infrastructure and into the production pipeline of commercial industry.²⁶

The US government is supporting exploratory work in deploying DNSSEC in .MIL and in using FISMA (Federal Information Security Management Act) as a route for deployment in government systems. It is a strategy of both carrots and sticks.

The Federal Information Security Act (FISMA) of 2002 (Public Law 107-347-Title III) mandates that "each federal agency shall develop, document, and implement an agency-wide information security program" to protect the information systems that support its operations. This law assigns NIST to develop the Federal Information Processing Standards (FIPS), technical guidance, and testing infrastructure necessary to define and implement FISMA regulations for the civilian Federal Government. NIST's FISMA specifications include FIPS-199 that specifies a FISMA classification framework for federal information assets and FIPS-200 that specifies the minimal security requirements for various points within this classification framework. In addition to these base standards, NIST has developed a series of FISMA special publications that further define specific recommended security controls (SP 800-53), and the means by which the implementation of these controls can be verified (SP 800-37) and measured for effectiveness (SP 800-53A). The FISMA act of 2002, makes NIST FISMA FIPS

²⁴ See Commercial Technology Network and Affiliations, http://www.techbriefs.com/spinoff/spinoff2002/ctn_1.html. The site clearly seeks to promote NASA's program and we identify it not to claim that the program is successful but to make the point that the government does develop technologies that are, in turn, commercialized and distributed.

²⁵ S. Deering, R. Hinden, Internet Protocol, Version 6 (IPv6) Specification, RFC 2460, December 1998 (<http://www.ietf.org/rfc/rfc2460>, verified, August 21, 2005).

²⁶ Karen S. Evans, Memorandum for Chief Information Officers, August 2, 2005. Executive Office of the President, Office of Management and Budget, Washington, DC. We note that OMB's scope covers federal executive agencies.

mandatory and non-wavorable for federal information systems. A second piece of federal policy defined in OMB Circular A-130 Appendix III (Security of Federal Automated Information Resources) further mandates that agencies adopt NIST's supporting FISMA guidance documents and requires that agencies report yearly on their level of implementation. *Not* reporting can have serious consequences; annual funding can be held up and public rating of agencies' performance can be adversely affected.

NIST's Special Publication 800-81 *Secure Domain Name System (DNS) Deployment Guide*, which includes DNSSEC, is positioned to be incorporated in the next release of the FISMA guidance on recommended security controls (SP 800-53). This will effectively mandate that these recommendations be addressed in future rounds of OMB FISMA reporting.

5. How long?

We posed three public policy questions in this paper:

1. How do we stimulate innovation in infrastructure services when those services are provided in a competitive, largely private commercial environment and the return is likely to occur in the long term and will also be shared?
2. What is the appropriate role of government in fostering infrastructure development when we are committed to largely privately owned and operated infrastructure facilities and services?
3. What is the balance among national and homeland security interests and global Internet management – or governance?

Let's sum up our answers to the first two questions: The appropriate role of government is threefold: supporting education and public awareness projects, supporting development efforts, and providing a market for infrastructure services, which enables prospective providers to develop capacity with relatively low risk while establishing the government as a leader by example. In this last role, deploying DNSSEC is both a technology transfer and an infrastructure development challenge

This three-fold role incents participation by key private stakeholders, who may be motivated by potential commercial applications or by prestige. For financial institutions, DNSSEC potentially offers increased security for the internal and external transactions that are intrinsic to their business. For universities who have invested substantially in high performance computing and high speed networking and where adopting advanced technologies is a source of prestige, deploying DNSSEC has been characterized by one staff member at a recent Internet2 workshop as a "no brainer." This strategy will result, it is anticipated, in precipitating cascading effects throughout the system through a combination of limited government mandates (for its agencies, regulated entities, and vendors) and voluntary adoption. This strategy should be accompanied by a systematic

effort to cultivate enhanced awareness of Internet security, both what individuals can do and what they can demand of their providers.

We have said little about the third question, the balance among security interests at home and abroad. The decentralized nature of protocol adoption means that individual countries that administer domains – the ccTLDs as they are known – can make decisions independently as discussions about Internet governance continue. Indeed, clearly both Sweden and the Netherlands have been on the forefront of DNSSEC deployment efforts. We acknowledge that there is an inherent tension among interests, but hardening the Internet infrastructure whether country by country or zone by zone in the end will only benefit us all.

And how long will it take? The flippant answer would be as long as it takes and no more. More seriously, we think we will begin to see major steps within the next 1-to-2 years. For DNSSEC, one critical early milestone is deployment of the signed root and the root key. Once this is in place, some of the top-level domains will be signed relatively rapidly and a small number of end-user systems will begin checking signatures on a regular basis. We expect all of these to be complete by the end of 2006, with some possibility of achieving this a bit earlier. Simultaneously, we are seeing progress at other levels of the DNS hierarchy. As of this writing in late August 2005, we believe .se will publish its signed zone file within two months, and that .nl will sign its records in 2006. There are already signed zones within the .gov top level domain, and we hope to have major portions of .gov signed by the end of 2007. Finally, and very importantly, RIPE NCC, a European Internet infrastructure services provider, has begun to deploy DNSSEC in its operations.