Bootstrapping the Adoption of Internet Security Protocols

Andy Ozment
MIT Lincoln Laboratory
& University of Cambridge

Stuart E. Schechter
MIT Lincoln Laboratory

June 26–28, 2006

Abstract

The deployment of network-wide security enhancements to the Internet has proven more difficult than many had initially anticipated. We leverage existing models of networks' value to model the problem of bootstrapping the adoption of security technologies. We describe a variety of policy interventions and deployment strategies that can help to catalyze this adoption. Using this framework, we provide a series of short case studies for previous attempts to deploy security technologies to the Internet. We then provide a detailed study of strategies for deploying security-enhanced protocols into the Internet's Domain Name System (DNS). Finally, we show how the adoption of these DNS security enhancements can help to alleviate bootstrapping problems that have impeded the deployment of other security-enhanced protocols.

This work is sponsored by the I3P under Air Force Contract FA8721-05-0002. Opinions, interpretations, conclusions and recommendations are those of the author(s) and are not necessarily endorsed by the United States Government.

This work was produced under the auspices of the Institute for Information Infrastructure Protection (I3P) research program. The I3P is managed by Dartmouth College, and supported under Award number 2003-TK-TX-0003 from the U.S. Department of Homeland Security, Science and Technology Directorate. Points of view in this document are those of the authors and do not necessarily represent the official position of the U.S. Department of Homeland Security, the Science and Technology Directorate, the I3P, or Dartmouth College.

Presented at:
The Fifth Workshop on the Economics of Information Security (WEIS 2006). Cambridge, England: June 26–28, 2006.

1 Introduction

> "Our economy and national security are fully dependent upon information technology and the information infrastructure. At the core of the information infrastructure upon which we depend is the Internet, a system originally designed to share unclassified research among scientists who are assumed to be uninterested in abusing the network."
>
> —*The National Strategy to Secure Cyberspace* [40]

1.1 Motivation

The *National Strategy to Secure Cyberspace* highlights the need to secure three key protocols that underlie the Internet: the Internet Protocol (IP), the Domain Name System (DNS), and the Border Gateway Protocol (BGP) [40].

This national strategy was released in February of 2003. By then, the Internet Engineering Task Force (IETF) had already completed new standards for adding security to the internet protocol and was close to finishing work on standards for securing the DNS. However, adoption of these protocols by the Internet's users and other stakeholders is far from certain.

Like the Internet itself, the technologies to secure it exhibit positive network effects: their value to individual users increases as others adopt them. As such, the initial benefits obtained by early adopters might fall significantly below the costs of adoption.

The technologies underlying today's Internet were heavily subsidized by US government agencies, yet still required decades to become widely adopted. Those calling for the deployment of technologies to secure the Internet have much shorter time frames in mind.

The secure protocols highlighted in the national strategy face additional hurdles, as they compete with the widely deployed insecure protocols that they are intended to update or replace. Existing economic models of competition between network technologies show how superior technologies often lose out to inferior technologies that have a head start in the race for user adoption.

To overcome these hurdles, we are pursuing a better understanding of the bootstrapping problem faced by those of us who work to deploy security technologies into the Internet. We also seek to better understand the types of strategies and interventions available to accelerate the rate at which users adopt security technologies. Specifically, we seek approaches to ensure the expeditious adoption of DNSSEC: a new standard for security enhancements to the DNS. We focus on DNSSEC because its successful adoption can accelerate the adoption of other protocols.

1.2 Roadmap

To better understand the adoption of security technologies into existing networks, we develop a model using existing characterizations of network value. This model, introduced in Section 2, is then used to formalize the problem of bootstrapping security technologies in Section 3.

Strategies and interventions to stimulate the adoption of technologies are introduced and categorized in Section 4. Case studies of the adoption of existing security technologies are presented in Section 5.

We then turn to the problem of deploying security into the DNS. For those unfamiliar with DNS, we provide a brief introduction to the system and its protocols in Section 6.

We describe the security enhancements to the DNS in Section 7. We show how the deployment of security into the domain name system can help to stimulate the adoption of other security protocols. Next, we introduce the stakeholders of the DNS in Section 8.

With all of the background in place, we then analyze the strategies and interventions available to stimulate the adoption of security into the domain name system. This discussion, in Section 9, builds on the strategies and interventions first presented in Section 4.

We place our model and analysis in the context of previous work in Section 10, and we conclude in Section 11.

2 Modeling Adoption of Network Security Technologies

To model the problem of deploying new security technologies into an existing network infrastructure, we begin by assuming that the existing network has a set of users U.

Deploying the security technology comes at a cost. We will focus on the per-user cost of deployment as these costs are often the primary barriers to adoption. We define c_i to be the fixed cost for user $u_i \in U$ to deploy the technology. We assume that this cost-to-adopt remains constant; however, if the adoption cost depends on the adoption choices of other users, the cost could easily be represented as a function.

2.1 A general model of network value

The benefit a user derives from joining a network depends on who else has joined the network. The same is true for those who adopt a security technology that enables them to join a more secure region of a larger network: the benefit is a function of the members of the secure subnetwork. We define $b_i(A)$ to be the benefit

that a security technology provides user $u_i \in U$ when she is part of a larger set of adopters $A \subseteq U$.

Because we view security technologies as goods, we assume that the benefit an adopter derives from the technology does not decrease when others adopt. In other words, the benefit of the security technology can only increase (or stay the same) as the set of adopters grows.

$$\forall i \in U, A \subseteq B \subseteq U : b_i(B) \geq b_i(A)$$

2.2 Adoption and social welfare

Deployment of a security technology by a group of adopters $A \subseteq U$ improves social welfare when the net social benefit of this deployment is positive:

$$\sum_{u_i \in A} (b_i(A) - c_i) > 0$$

If the goal of deployment is to maximize social welfare, we can formalize this goal as the maximization of net social benefit:

$$\max_A \left(\sum_{u_i \in A} (b_i(A) - c_i) \right) \quad (1)$$

Depending on the properties of the benefit function b, finding the optimal adoption set (or even determining if any adoption set exists) may be computationally hard.

2.3 Winners and losers

When identifying an adoption set using only the constraint shown in Equation 1, there may be a subset of adopters who lose out as a result of deployment: their costs-to-adopt outweigh their benefits. In other words, an adoption set A may have losers (L) as well as winners (W).

$$\begin{aligned} \forall u_i \in L : b_i(A) &< c_i \\ \forall u_i \in W : b_i(A) &\geq c_i \\ W \cup L &= A \end{aligned}$$

Losers reduce their individual welfare by adopting the technology, but they improve the overall welfare because others benefit at no additional cost.

We can categorize solutions to technology adoption problems by the fate of the losers: solutions can either be constrained so that the adoption set contains only winners or they can include mechanisms to force losers to adopt.

To ensure that the adoption set contains only winners, we can place a pareto optimality constraint on deployment.

$$\forall u_i \in A : b_i(A) > c_i \quad (2)$$

One way to encourage losers to adopt a technology would be to use market mechanisms to internalize the positive network effect: transfer the winners' surplus to adopters who would otherwise suffer a loss by adopting. Alas, market solutions are difficult to achieve because the surplus is often distributed over a large number of users.

Alternatively, a social planner could mandate that users adopt. This approach requires that the planner either fund the costs of compliance or face resistance from those whose costs are greater than their benefit when adoption is mandated.

2.4 Simplified models of network value

Two common assumptions are made to simplify how networks are valued: adoption-set equivalence and adopter equivalence. These assumptions result in a simpler benefit function.

We say that the benefit function b_i^* is *adoption-set equivalent* if its value depends only on the fraction of users that have adopted the technology, regardless of who the adopters are. We can thus define the adoption-set equivalent benefit function b_i^* in terms of another function $\mathcal{B}_i(n)$:

$$b_i^*(A) = \mathcal{B}_i(x) \text{ where } x = \frac{|A|}{|U|}$$

We say that $\mathcal{B}$ is *adopter-equivalent* if all users derive the same benefit from being part of an adoption set.

$$\forall u_i, u_j \in U; x \in (0, 1] : \mathcal{B}_i(x) = \mathcal{B}_j(x)$$

A number of existing models of network effects are adoption-set and adopter equivalent. For a network of n users, Metcalfe [28] defined the value V of a network to a user as the value of each potential interconnection to the $n-1$ other users. We call such a network a Metcalfe's Law (ML) network.

$$V^{\text{ML}}(n) = n - 1$$

Metcalfe defined the value of the entire network as the aggregate of its value to individual users. Thus, Metcalfe's law states that the value of the network is approximately the square of the number of users [28, 46].

$$\sum_{i=1}^{n} V^{\text{ML}}(n) = n(n-1) \approx n^2$$

The value function V^{ML} for Metcalfe's law takes as its input the number of users n in a growing network. Our benefit function $\mathcal{B}$ takes as its input the fraction x

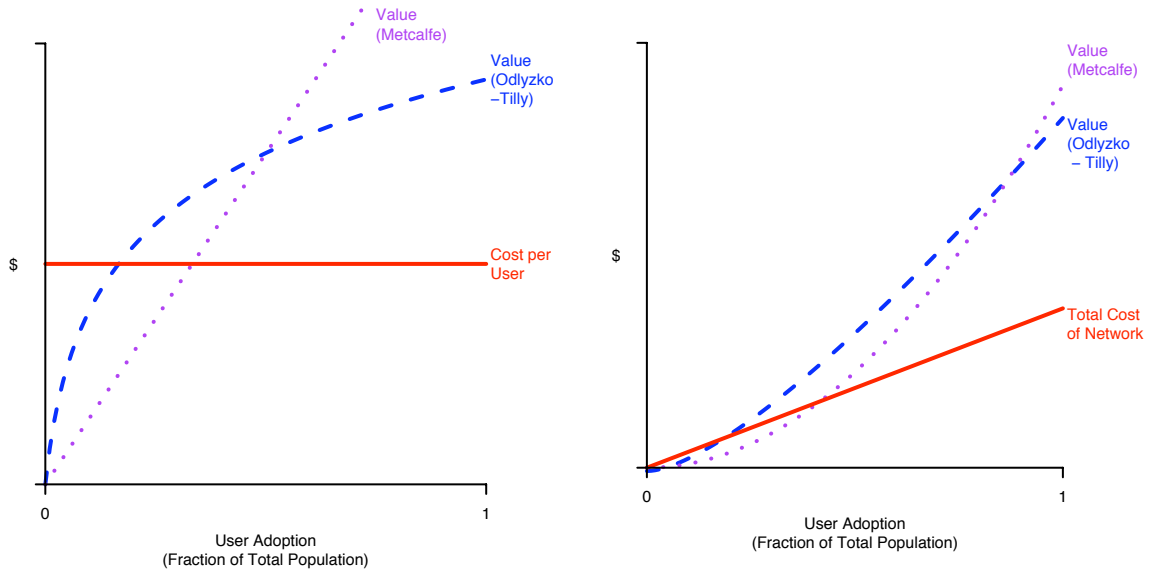

(a) Value of network to an individual

(b) Total value of a network

Figure 1: Models of the Value of a Network

of adopters in an existing network with N users ($N = |U|$). When modeled using Metcalfe's Law and assuming a maximum individual benefit M, our benefit function could now be written as follows:

$$\mathcal{B}_i^{\mathrm{ML}}(x) = K\left(x - \frac{1}{N}\right) \text{ where } K = \frac{M \cdot N}{N-1}$$
$$\sum_{i=1}^{N} \mathcal{B}_i^{\mathrm{ML}}(x) = N \cdot K\left(x - \frac{1}{N}\right)$$

Recognizing the influence of diminishing returns in the value of a growing network, Odlyzko and Tilly [30] define the aggregate value of a network to be $n \log n$. Assuming adopter-equivalence, the value of an OT network to individuals and to all users can be summarized as follows:

$$V_i^{\mathrm{OT}}(n) = \log n$$
$$\sum_{i=0}^{n} V_i^{\mathrm{OT}}(n) = n \log n$$

Converting their value function to represent the benefits when a fraction of users x deploys, and where M is the maximum benefit per user, yields the following two equations:

$$\mathcal{B}_i^{\mathrm{OT}}(x) = \log_e(1 + Q \cdot x) \text{ where } Q = e^M - 1$$
$$\sum_{i=0}^{n} \mathcal{B}_i^{\mathrm{OT}}(x) = N \log_e(1 + Q \cdot x)$$

Figures 1a and 1b illustrate the differences between Metcalfe's Law and the model of Odzlyko and Tilly. Figure 1a shows the benefit and cost of joining the network as viewed by an individual. Figure 1b shows the aggregated benefits and costs over all users.

As we approach the problem of bootstrapping security technologies into existing networks, we will use the Odlyzko-Tilly model of network value.

3 The Problem of Bootstrapping

We illustrate the bootstrapping problem in two cases: a general case under the simple adoption model and a case specific to the adoption of security protocols that add authentication to existing networks.

3.1 Bootstrapping in the simple adoption model

Deploying new network technologies can be challenging because users may act independently. Furthermore, they may only deploy when the network is in a state at which the *immediate* benefits of adopting the technology outweigh the costs.

Game theoretic models of adoption often assume that users must choose to adopt a technology in parallel, unable to observe the adoption decisions of others for some time. We will assume that users who decide to adopt a new technology can do so instantaneously. Such assumptions are realistic when the technology is readily

available from a service provider at the flip of a switch. Assuming instantaneous deployment allows us to factor out choices made in anticipation of decisions by others.

When instantaneous decisions are possible, users can wait to adopt a technology until the benefits of doing so outweigh their costs. We assume there is no penalty for late adoption. Though we assume fixed costs in our model, costs may actually go down as other users deploy.

Because we assume that users need not adopt until the benefits outweigh the costs, adoption by a set of n users ($n = |A|$) requires an ordering of these users $A_n = (u_1, u_2, \ldots, u_n)$ such that:

$$\forall i \in [1, n] : b_i(A_i) > c_i \tag{3}$$

Unfortunately, a minimal level of deployment is often required before *any* users can obtain benefits that will outweigh their costs. This condition creates a bootstrapping problem: no users adopt the technology, although a subset of the users would benefit if they all adopted.

Figure 1a shows a simplified model of a network technology that faces a bootstrapping problem. Both models' benefit lines lie below the cost line on the left side of the graph. The costs of being an early adopter outweigh the benefits.

3.2 Bootstrapping the adoption of authentication

Authentication is a security feature of communications protocols: it enables the receiver of a communication to verify that a communication came from the purported sender. Authentication also implies that a receiver can verify that the communication has not been changed in transit. If all senders are required to prove the authenticity of their communications, then receivers can classify all communications into one of two categories: (1) *provably authentic* or (2) not provably authentic and presumed *forged*. If a communication proves to be authentic, the recipient can accept and use the information with confidence. If the communication is not provably authentic, then the recipient can disregard it as a forgery.

However, if some legitimate senders (non-adopters) do not prove the authenticity of their communications, the messages they send will fall into a third category: *authentic but not provably so*. Senders who provide proofs of authenticity for all outgoing communications (adopters) do not generate messages in this category. Unfortunately, recipients may not be able to ascertain whether a message that is not provably authentic falls into the second or third category: messages that are authentic—but not provably so—cannot be distinguished from forged messages.

To know whether a message that cannot be proven authentic should be treated as forged, the receiver of a message must know whether the sender is an adopter who only sends messages that can be proven authentic. If the receiver cannot distinguish adopters from non-adopters, he must either accept all messages that cannot be proven authentic or reject them. Accepting messages that cannot be proven authentic eliminates the key benefit of adopting authentication. Discarding messages that cannot be proven authentic incurs the cost of cutting of all communications with non-adopters. This choice exacerbates the bootstrapping problems faced by protocols that add authentication to an existing network.

The need to differentiate adopters from non-adopters can be illustrated with an example from the world of email. Adopters of email authentication sign outgoing messages and have the option to verify these signatures when receiving email. In our example, a user *Bob* has adopted authentication and receives an unsigned message purportedly from *Alice*.

If Alice has adopted authentication, she signs all of her email. She thus expects Bob to reject unsigned messages that purport to be from her but cannot be authenticated. If Alice has not adopted authentication, she does not sign her messages. She thus expects Bob to accept messages from her even though they are not signed. To know whether to accept an unsigned message purportedly from Alice, Bob must know whether Alice has adopted authentication.

Solving this problem requires a secure mechanism through which Bob can determine if Alice has adopted authentication. For example, if Bob already knows Alice he might consider it safe to call and ask if she signs her messages. Unfortunately, the Internet has lacked a general mechanism with which to securely determine whether a system or its users has adopted an authentication technology.

Without such a mechanism, adopters cannot determine whether to accept or reject unauthenticated communications from users not yet known to be adopters. There is little benefit to deploying authentication if unsigned messages are treated as authentic. However, rejecting unsigned messages imposes a significant cost: the benefits of being a member of the unauthenticated network are lost. Figure 2 illustrates the costs and benefits of adopting authentication and rejecting unauthenticated communications.

Users will only begin authenticating incoming messages when the benefit of authentication outweighs the cost of ignoring messages from those who have not adopted authentication. This tipping point is illustrated in Figure 2.

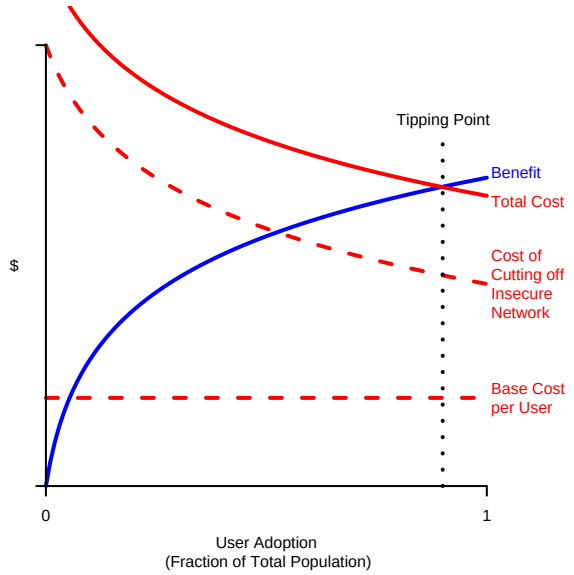

Figure 2: Bootstrapping authentication technologies

4 Approaches to Stimulating Technology Adoption

How can bootstrapping problems be overcome? There are a number ways to stimulate the adoption of a security technology, some of which require more intervention by policy makers than others.

In this section, we will illustrate the different approaches to bootstrapping using graphs of costs and benefits. While the approaches we describe are general, the illustrations are generated using an Odlyzko-Tilly model of adoption benefits and fixed per-user adoption costs.

4.1 Global mandate

The most invasive approach to stimulating the adoption of a technology is to mandate that users adopt and to enforce this mandate by imposing costs (*e.g.* fines) on those who fail to comply. Users who adopt thus receive the conceptual 'benefit' of not being subjected to noncompliance costs.

Figure 3a illustrates the total benefit of adoption to an individual as the sum of the individual's existing adoption benefit and the compliance benefit. If the first user to adopt receives no benefit other than that of compliance, this compliance benefit must outweigh the cost of adoption.

4.2 Partial mandate

Once some minimum number of users have adopted a technology, the technology may reach a tipping point at which adoption will take place without further intervention. A partial mandate can be used to induce this minimum number of users to adopt.

Figure 3b illustrates this tipping point; it occurs when the fraction of users who have adopted, x, is large enough that $\mathcal{B}(x) > c$.

4.3 Bundling complements

Another way to increase a technology's net benefit is to bundle a complementary technology. In particular, complementary technologies that benefit early adopters are best able to spur adoption.

When all users benefit from a complementary technology, the complement's benefits will resemble the compliance benefits from a global mandate (Figure 3a). When only a subset of users benefit, the complement's benefits will more closely resemble the compliance benefits of a partial mandate (Figure 3b). Figure 3c shows a complement whose benefit exhibits positive network effects and is always larger than its cost.

4.4 Facilitating subnetwork adoption

Adoption may occur naturally in situations where a single group is large enough and sufficiently well coordinated. The group's size must enable it to increase the fraction of users who adopt, x, past the tipping point so that $\mathcal{B}(x) > c$. The group must then be coordinated well enough to ensure that those members do adopt.

For example, most of a technology's positive network effects for a large organization may occur within that organization. In Figure 3d, one large organization will obtain sufficient benefit from the technology for it to adopt within its own organization. The adoption by that subnetwork will then encourage adoption by the remainder of the network.

The deployment of fax machines occurred through this mechanism: companies initially bought fax machines to connect their own offices.

The designers of a technology may be able to take advantage of this effect; they can craft features that provide significant benefits within organizations and other subnetworks.

4.5 Coordination

A related approach is for individuals and groups to coordinate to adopt together. By doing so, they can form larger subnetworks than would be possible on their own.

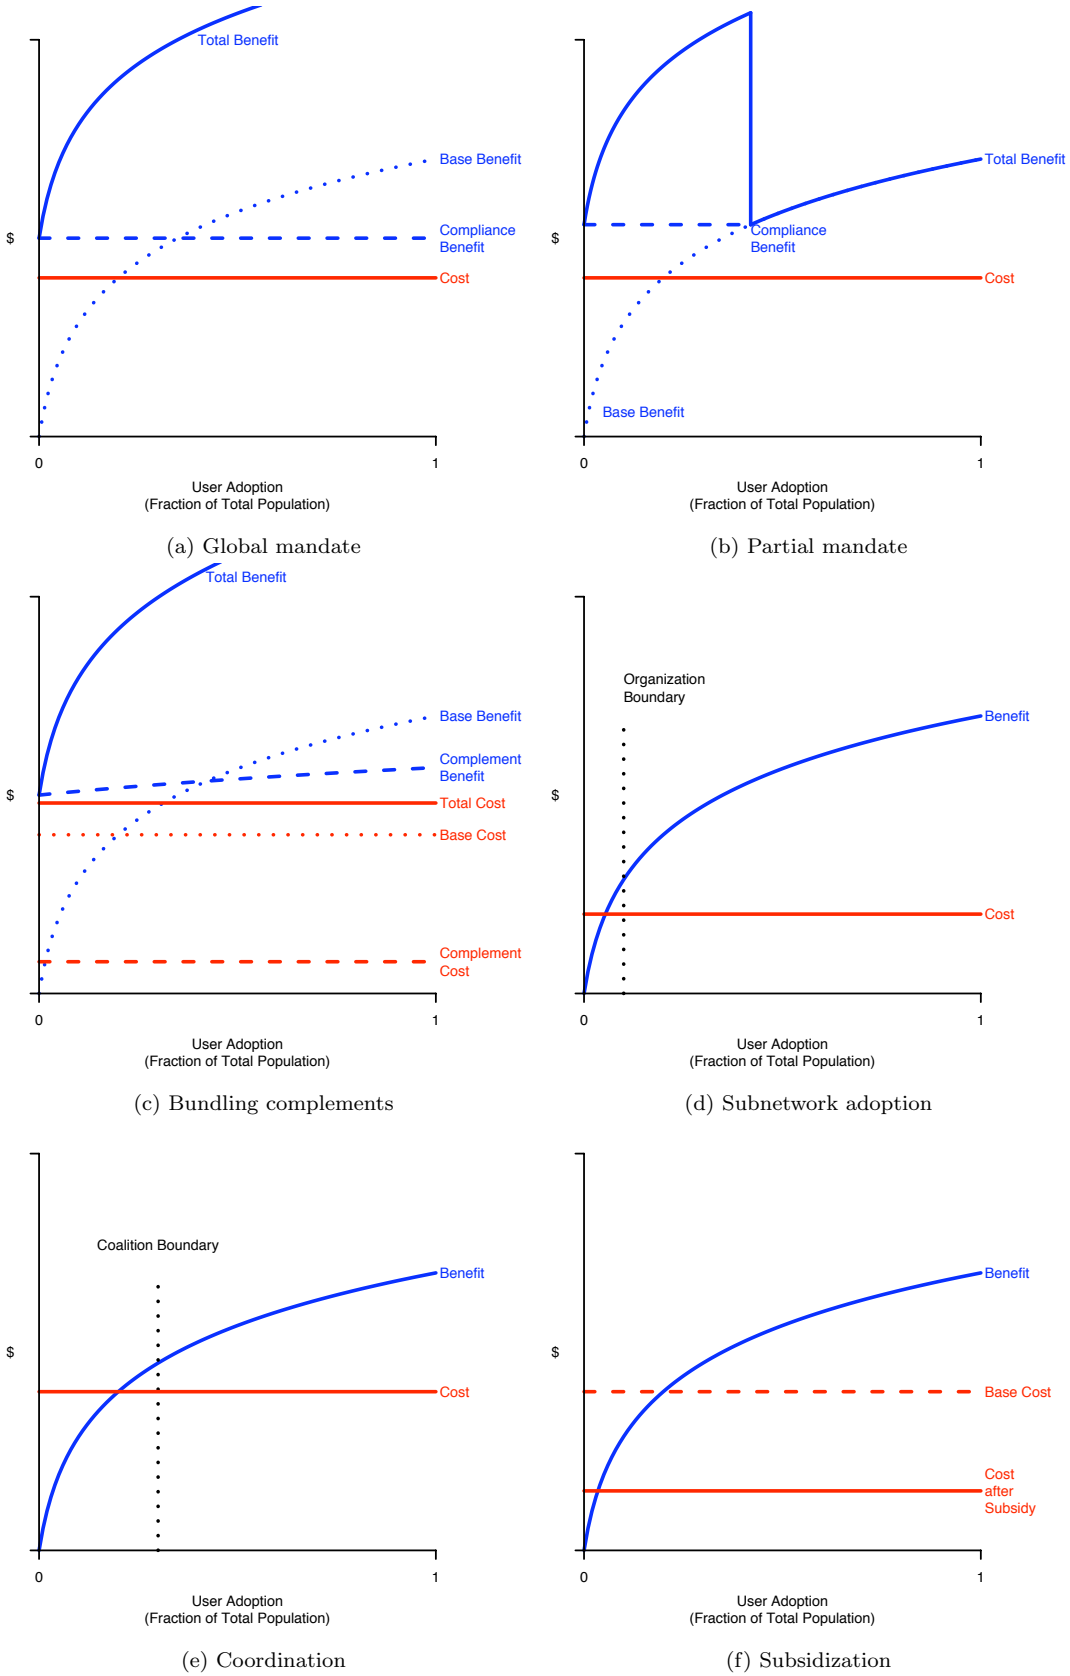

(a) Global mandate

(b) Partial mandate

(c) Bundling complements

(d) Subnetwork adoption

(e) Coordination

(f) Subsidization

Figure 3: Approaches to stimulating technology adoption

Such coordination efforts are not without costs: partners must be recruited, agreements must be drafted, and enforcement mechanisms may be required to ensure that all parties adopt. Much of the existing work on technology adoption addresses it as a problem of coordination [13]. Figure 3e illustrates a coalition that has coordinated adoption such that $\mathcal{B}(x) > c$ and the bootstrapping problem is overcome.

4.6 Subsidization

The approaches above seek to increase the benefit of adoption; another approach is to reduce the cost of adoption via subsidization. One form of subsidy is for a social planner to reward individuals who adopt (*e.g.* via tax breaks). Alternatively, a government or industry group may invest in technologies or processes that universally lower the costs faced by all adopters. Figure 3f shows the result of an investment that lowers the costs for all adopters.

5 Case Studies

Three technologies illustrate the challenges of bootstrapping the adoption of a security technology: SSH, HTTPS, and IPsec. Some of the policy techniques discussed above have been used to encourage deployment of these technologies; the extent to which these technologies have been adopted ranges from minor to near-universal.

5.1 The Secure Shell (SSH) protocol

At the 2003 Workshop on Economics and Information Security, Nicholas Rosasco and David Larochelle [33] presented a case study of the deployment of the secure shell (SSH) protocol: a protocol that enables a client to remotely issue commands to a server. The authors argued that network externalities were not "a significant factor impeding the adoption of SSH."

Rosasco and Larochelle attributed the success of SSH to the low cost to adopt the technology and the protocol's utility within organizations. The cost to transition to the `ssh` client command was low because it offered the same general interface as the commands it replaced. The command also retained the functionality that the user expected.

Another reason the deployment of SSH did not face a bootstrapping problem is that the protocol is used most heavily for intra-organizational communication: clients primarily access the servers of organizations to which they already belong. The benefits of adopting SSH were also clear to system administrators who could mandate its adoption within their organizations: many feared password sniffing long before SSH arrived to prevent it.

By transitioning from `telnet`/`rsh` to `ssh`, an organization can receive the full benefit of the transition almost immediately. As a result, SSH serves as an excellent example of deployment via subnetwork adoption.

In Section 3.2, we argued that technologies that introduce authentication into existing networks face special bootstrapping problems. SSH is such a technology: it introduces authentication of servers to the clients who wish to access them. Clients must know whether or not to use the secure and authenticated protocol (SSH) or insecure protocol (`rsh` or `telnet`) when initiating a connection to a server.

SSH clients keep track of whether servers authenticate themselves in order to determine which should authenticate themselves in the future. If the first connection to a server is not attacked, the client will detect that the server authenticates itself and future connections can be protected. This approach was effective because SSH clients usually connect to servers with whom they have previously communicated. Because SSH was deployed preliminarily within organizations, system administrators were able to use other organizational channels to inform their users when their organization transitioned to SSH.

5.2 HTTPS

Security was added to the HyperText Transfer Protocol (HTTP) after the world wide web was already a well established network.

Transitioning from HTTP to HTTPS is costly. A server administrator must generate a public key, purchase a certificate authenticating the key, and install the key and the certificate into her server. Certification authorities charge anywhere from $20 to almost $1,000 for these certificates. Furthermore, the cryptographic algorithms used by HTTPS consume significant computational resources.

Unlike SSH, web users often browse to countless servers run by myriad organizations. Browsers frequently load pages from organizations with which their users have no prior relationship. These browsers (clients) have no secure mechanism to determine which sites (servers) should authenticate themselves and which do not support HTTPS. Users only get the benefits of server authentication if they know which sites support it.

The deployment of HTTPS was driven, in part, by industry-imposed partial mandates. Credit card companies such as Visa [45] and Mastercard [27] required online merchants to adopt HTTPS.

HTTPS thus became widely deployed among mer-

chants and financial institutions. However, partial mandates did not push HTTPS beyond the tipping point at which more sites would choose to adopt. While there are 40 million registered domain names in `.com` alone, the total number of SSL server certificates (which are required for server authentication) is less than 300,000 [39].

To this day, browsers still default to insecure HTTP regardless of whether a server has deployed HTTPS. Browsers require HTTPS server authentication only when sent to an `https`-prefixed address. Almost all sites that support HTTPS still support HTTP, and use this insecure protocol to redirect users to the server-authenticated protocol. As this approach is itself insecure, users must verify whether or not security has been activated: a task which has proven intractable for most. Thus, the potential security benefits of HTTPS have not been fully realized.

5.3 IPsec

The Internet Protocol Security (IPsec) standards [22] were designed to provide encryption and authentication to any Internet communications. As with HTTPS, the costs of key management provide a significant barrier to the adoption of IPsec.

As a result, IPsec is primarily employed in situations where this cost can be overcome. Organizations use IPsec to implement virtual private networks, providing software to their users that includes the appropriate keys to talk to a single server. As this application has no network benefits outside the organization, IPsec's use in VPNs (subnetworks) has not helped to spur its adoption elsewhere.

IPsec has been bundled into the next version of the Internet Protocol, IPv6 [41, 35]. In turn, IPv6 is being driven in part by subnetwork deployment: the US Department of Defense can derive the majority of its benefit from IPsec and IPv6 internally, so it is working to deploy by 2008 .

6 A Brief Introduction to the Domain Name System

The deployment of the DNSSEC security technology to the existing Domain Name System (DNS) serves as our primary case study. Before we examine the deployment process, we will provide a brief overview of the DNS and of the DNSSEC standard.

6.1 Architecture of the DNS

The domain name system is a directory that maps the names of Internet hosts to essential information about these hosts, such as their IP addresses.

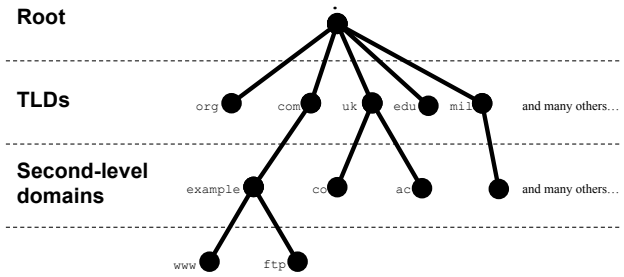

Figure 4: The DNS hierarchy is an inverted tree, with the root zone as the top of the hierarchy.

The DNS is a hierarchical database, which is reflected in the names of Internet hosts. These names are composed of the different levels of the hierarchy, written in bottom-up order. Figure 4 shows the DNS hierarchy.

For example, the domain name '`example.com.`' indicates three ascending levels of hierarchy: '`example`', '`com`', and '.' (*i.e.* root). The topmost domain, the root, is represented by the dot at the end of the domain name. All domain names are descendants of the root, so this dot is optional and is almost always omitted.

The DNS is implemented as a distributed database. The group of records stored by a server for a domain is called a *zone*.[1] A zone can be served by any number of servers. For example, the records for the root zone are available from servers at thirteen IP addresses. DNS clients come pre-configured with a list of these thirteen root server addresses.

The root zone database contains records that provide the IP addresses of its children in the hierarchy, the Top-Level Domains (TLDs). Examples of TLDs include `com`, `net`, `org`, `ca`, and `uk`. Like the root zone, a TLD's zone database stores records that provide the IP addresses for its child subdomains, the second-level domains. The `example` subdomain of the `com` TLD is an example of a second-level domain.

6.2 Vulnerabilities of the DNS

The domain name system is vulnerable to attacks on communications to the system as well as attacks on the system itself.

6.2.1 Denial of service

The goal of a denial of service attack is to make a system unavailable to its users. For example, an attacker could

[1] A domain's zone may also include any records for the descendants of that domain.

overload a service with so many artificial requests that it is unable to handle legitimate requests.

6.2.2 Man in the middle

Attackers who can capture communications on the path between clients and DNS servers are positioned to carry out man-in-the-middle attacks: intercepting DNS requests and forging responses to them. For example, when a client requests the address of a subdomain's server, an attacker could intercept this request and forge a response containing the address of his own server. Users of wireless networks are especially vulnerable to man-in-the-middle attacks by 'rogue' access points.

6.2.3 Server compromise

Another way to attack the integrity of the DNS records a client receives is to add, delete, or modify the records *before* they are requested. To access the records, an attacker could compromise one of the hosts that serves them. A recent analysis by Ramasubramanian and Sirer [32] found that 17% of the DNS servers on the Internet have known security vulnerabilities. An attacker who compromised those servers could control 30% of the names listed in the Yahoo and DMOZ.org directories.

A more complete threat analysis of the domain name system can be found in IETF RFC 3833 [7].

7 DNSSEC: Security Extensions for the DNS

The Domain Name System Security extensions (DNSSEC) standard [4, 5, 6] enables clients to verify both the integrity of DNS records and the authenticity of those records' origin.

7.1 Signing zones

DNSSEC specifies a set of extensions to the DNS that enable domain operators to cryptographically sign the data within their domains. Clients that request data from the DNS use these signatures to establish that the records returned by a server are authentic: that they have not been changed since they were signed by the domain's owner.

Clients can establish the authenticity of any domain's records by starting with a single pre-configured public key: the root zone's public signing key. The corresponding private key is used by the root zone's operator to sign each of the existing records in the root zone. Those signatures are themselves stored as records in the root zone.

7.2 The chain of trust

The root zone's key is the anchor of a chain of trust: this key is used to sign the public key of each top-level domain (TLD). Each TLD then uses its private key to sign the public key of each of its second-level domains.

Assume, for example, that the root zone, the `com` TLD, and the `example.com` second-level domain are all DNSSEC-compliant. A client who wishes to obtain the IP address for `www.example.com` first contacts a root zone DNS server. From that server, the client obtains records that contain the IP addresses of the `com` zone's DNS servers and the `com` public key.[2] The client uses its pre-configured root public key to verify the integrity and authenticity of those records.

The client now contacts a `com` zone DNS server; from that server, it obtains records that contain the IP addresses of the `example.com` zone's DNS servers and the `example.com` public key. This time, the client uses its newly obtained `com` public key to verify the integrity and authenticity of those records.

From `example.com`, the client next obtains records containing the IP address of the host named `www.example.com`. The client uses its `example.com` public key, obtained from `com`, to verify this address.

Having securely obtained the address from the DNS directory, the client is now able to initiate a connection with `www.example.com`.

In this example, the root zone's key is used to verify the `com` zone's key. The `com` zone's key is then used to verify the `example.com` zone's key. This chain of trust allowed the client to ensure the authenticity and integrity of each record it used to obtain the IP address of its final destination, `www.example.com`. This example chain of trust is shown in Figure 5.

7.3 DNSSEC-compliance

DNSSEC-compliant domains sign all of their records and certify the keys of their DNSSEC-compliant children. The client can follow the chain of trust down the DNS hierarchy until there are no more child zones or until an unsigned (noncompliant) zone is reached. The integrity and authenticity of all the records in a domain can be verified, so long as those records are stored in a DNSSEC-compliant zone and all of that zone's ancestors are DNSSEC-compliant.

In our example above, each domain in which the client was interested was DNSSEC-compliant. If the `example.com` domain is not DNSSEC-compliant, then the client can only verify the integrity and authenticity

[2] The client actually obtains a signed record containing the *fingerprint* of the `com` zone's public key. The actual key is obtained from the `com` zone itself. For clarity of exposition, our description of this process omits this complication.

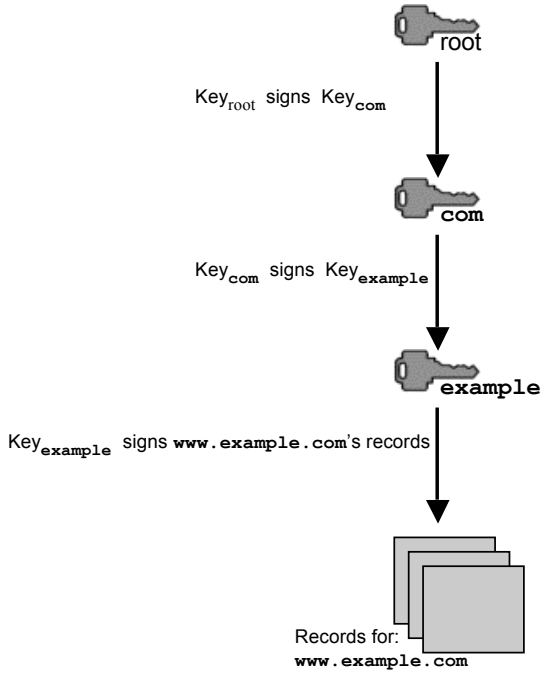

Figure 5: An example of the DNSSEC chain of trust. The client is pre-configured with the root zone's public key. It is able to use that key to verify the next key in the chain. This recursive process continues until the client is able to verify all of the information it is seeking (or until it encounters a DNS zone that is not DNSSEC-compliant).

of the information it received from the root zone and the `com` zone. The client receives no such assurances for the contents of `example.com`.

To ensure that clients know which records should be present (and signed) and what data is absent, each DNSSEC-compliant zone is required to store a signed table of contents: a record that lists all record types available for the zone. If a record type is present in the table of contents, then the set of records of that type must be present and signed as a group. A client that always checks this table of contents and verifies its signature cannot be deceived into believing a record is absent when it should be present.

7.4 Benefits to the client

Clients that obtain records from DNSSEC-compliant domains are protected against two of the attacks described in Section 6.2: man-in-the-middle attacks and compromised servers.

A man-in-the-middle could attempt to modify, remove, or add records sent from a DNS server to the client. By verifying DNSSEC signatures, clients can detect modified or added records. By checking the signed table of contents, clients can detect when records have been removed.

When traveling down the hierarchy, the client learns from a parent domain whether the requested child domain is DNSSEC-compliant. The client then knows that the child domain has a table-of-contents record, so it can detect an attack in which the transmission of that record is blocked.

If a DNSSEC-compliant server is compromised by an attacker, clients can detect whether any of the records they request have been added, deleted, or modified—so long as the zone's key has not been compromised. Thus, zone operators should not store their keys on their DNS servers: keys should be stored and zones signed on hosts that are less accessible to outsiders (*e.g.* on offline hosts). Without the key, the attacker cannot sign new or modified records. If the attacker cannot sign modified tables of contents, he will be unable delete records without the change being detectable.

7.5 Effects on efforts to deploy authentication

DNSSEC also provides a solution for the special problem of bootstrapping the adoption of protocols that provide authentication, which was discussed in Section 3.2.

Using DNSSEC, a domain can publish in its signed zone a statement that it supports authentication. This signed statement thus signals potential communications partners that the domain is a member of this authentication technology's set of adopters. Those partners who wish to communicate with the adopter—or who receive communications purporting to come from the adopter—can use DNSSEC to securely identify that the domain has indeed adopted the authenticated protocol. If the domain has adopted the authentication technology, its communications partners can safely reject unauthenticated communications purporting to be from that domain.

DNSSEC thus enables adopters to benefit from authentication without giving up their ability to communicate with those who have not adopted authentication. By eliminating the cost of abandoning the unauthenticated network, DNSSEC mitigates the special bootstrapping problems that would otherwise impede authentication protocols.

8 DNS Stakeholders

The domain name system has many stakeholders, each of which must play a role in the adoption of DNSSEC. We will start at the root and work our way down to registrants and their end users. For the sake of brevity,

we will leave a comprehensive listing of stakeholders for further treatment elsewhere.

8.1 The root

The Internet Corporation for Assigned Names and Numbers (ICANN) has managed the root zone under a memorandum of understanding with the US Department of Commerce since 1998 [16]. The servers at the thirteen root addresses are operated by ICANN and eleven other organizations.[3]

One of ICANN's four primary goals is to preserve "the operational stability of the Internet" [17], so its Security and Stability Advisory Committee is exploring the adoption of DNSSEC. However, ICANN would only mandate adoption if the security benefits of DNSSEC outweigh any potential instability that could result from the deployment process. At the moment, that committee has not yet made a recommendation with respect to DNSSEC.

8.2 Registries

A *registry* is a domain that issues subdomains to customers (called *registrants*). Registries have two components to their business: they issue subdomains to registrants, and they operate the DNS servers for the registry's own domain. These DNS servers direct requests for the registrant's subdomains to the registrant's zone. ICANN awards contracts to run the generic TLD registries like `com`, but individual governments exercise control of the registries that run their country-code TLDs (*e.g.* `.ca` or `.uk`).

Registries' DNS servers must be openly accessible if they are to serve the Internet. This accessibility makes them vulnerable to attack. If a registry signs its zone data offline and then moves the DNSSEC-signed zone to the server, it can reduce the damage that would result from a server compromise. If an attacker did compromise the server and modify records, DNSSEC-aware clients could detect the modifications.

Still, the net benefit of deployment may not be sufficiently compelling. A study by Olaf Kolkman showed that full deployment of DNSSEC could increase the bandwidth required to serve a registry's zone by a factor of 2 to 3 [23]. Kolkman, who is one of the chairs of the IETF's DNS Extensions Working Group, has also asserted that deploying DNSSEC provides "no obvious immediate benefits" to registries [24].

One registry that faces especially high costs is VeriSign. The `.com` zone operated by VeriSign contains child records for the 40 million domain names registered in that zone [44]. These records are currently kept in memory, and adding signatures to all of them will likely require more memory than is available on VeriSign's servers. To reduce the transition cost, a modification to the DNSSEC standard is under development and nearing completion. This modification would make it possible to sign records only for those children that are themselves DNSSEC-compliant.

If registries are not concerned about the potential for server compromise, it may be necessary for ICANN—or the relevant national government—to either mandate deployment or to allow registries to charge a higher price for DNSSEC-secured domains. If registries charge a higher price for DNSSEC-secured domains than for insecure domains, then some customers will choose the cheaper, insecure domains. This price differentiation would thus result in lower levels of adoption.

8.3 Registrars

Registrars acts as the middle-men between registries and the registrants who want to obtain a domain name. Registrars were created to eliminate the monopoly power over TLDs that was held by registries that sold domain names directly.[4]

To support DNSSEC, registrars need only change some of their software and management processes. Registrars are in a commodity business, so they will only add this support if customers demand it. However, the reason that registrars are in a commodity business is that existing registrants can switch registrars with relatively low cost. As a result, only one registrar need offer DNSSEC to enable any registrant to obtain it, and at least one registrar already supports DNSSEC [8].

8.4 DNS service providers

DNS service providers maintain DNS zone servers for customers who do not want to maintain their own or who want a backup server external to their own organization. DNS service providers also provide management and operational support to customers who do want to maintain their own DNS servers. Many registrars also act as DNS service providers.

DNS service is a commodity, so it is priced by the market. Service providers could choose to charge a fee for DNSSEC. Alternatively, they could adopt DNSSEC universally in order to attract security-conscious customers. The latter might be attractive as, once the ser-

[3] 11 organizations + ICANN = 13 addresses because VeriSign operates two of these addresses.

[4] A registry is the sole supplier of names within its namespace. Registrars resell these names to registrants. The wholesale price that registries charge to registrars is fixed by ICANN (for generic TLDs) or by their respective governments (for country-code TLDs).

vice provider has adopted DNSSEC, the marginal cost to support an additional customer is low.

Both registrars and DNS service providers can be expected to support and provide DNSSEC if their customers demand it.

8.5 Registrants

Registrants are the customers of the DNS community: they are the consumers of domain names. For example, the e-commerce company Amazon is a registrant with the `com` registry for its `amazon.com` domain.

Registrants can use DNSSEC to ensure that their DNS directory data is published securely. DNSSEC-aware clients will be able to detect—and render ineffective—attacks that attempt to modify this data. However, users of clients that are not DNSSEC-aware will not be protected.

The cost to registrants of using DNSSEC is primarily one of operational management: keys must be created and managed and zones must be signed. These costs are likely to be lower for clients whose zones are managed by DNS service providers, as these service providers can amortize the costs of automating these operational tasks.

8.6 End users

End users who adopt DNSSEC-aware clients benefit from the enhanced security of the results they retrieve from signed DNS zones. The primary cost of adoption is that of obtaining DNSSEC-aware client software[5] and DNSSEC aware applications.

We cannot expect individuals to demand DNSSEC-aware client software: few understand the benefits of the security technologies already installed in their clients (*e.g.* HTTPS). It is more likely that client software will become DNSSSEC-aware if this feature is bundled into an upgrade that is part of users' normal upgrade cycle.

Furthermore, not all end users adopt software as an individual choice. Many adopt whatever is recommended or required by their organization. Firms may recognize the benefits of DNSSEC for end users who need to securely access the firm's network services. These firms could thus increase the rate of the deployment of DNSSEC-aware clients.

Summary

Registrars and DNS service providers are likely to wait to deploy until they can detect demand from their customers: registrants.

As a result, three key stakeholder groups stand out from those listed above: (1) the root and registries, (2) the registrants, and (3) the end users and their client software. Widespread DNSSEC deployment is dependent upon members of these three groups seeing a sufficient net benefit from adoption.

The first group, the root and registries, do not exhibit network effects. Deploying security and stability to the domain name system is part of ICANN's mission, so a mandate is possible. Indeed, registries may adopt in order to avoid the regulation that would come with a mandate. Regardless, network effects do not play a significant role in their adoption, so these stakeholders will not be the focus of our analysis.

Rather, we will utilize our model of network effects and the bootstrapping problem to discuss strategies and interventions that can stimulate adoption by registrants and end users. The benefit that DNSSEC provides to registrants and end users does exhibit positive network effects. The greater the number of registrants that utilize DNSSEC, the greater the trust that end users can place in the Internet. The greater the number of end users who utilize updated client software, the greater the value to registrants of utilizing DNSSEC to protect DNSSEC-aware clients.

9 Stimulating the Adoption of DNSSEC

Any new network technology that seeks to displace an entrenched network faces significant challenges. Previous Internet security technologies have also faced significant adoption challenges; as a result, many individuals in the security community are skeptical that DNSSEC will be adopted by more than a handful of users. For DNSSEC, the challenges are exacerbated by the difficulty of quantifying the benefits of security.

Adoption of DNSSEC might occur naturally if the DNS was the target of the majority of today's attacks. However, today's Internet offers the attacker a number of easier targets. Ensuring that DNSSEC is widely deployed *before* the DNS becomes an attractive target is a much more difficult task.

The role of end users and registrants in the deployment of DNSSEC is analogous to that of VCR owners and video rental stores. As the fraction of users with DNSSEC-aware clients increases, so does the benefit for registrants who sign their DNS zones. As the fraction of registrants who sign their zones increases, so does the benefit for users who adopt a DNSSEC-aware client.

In the terms of the model we described in Section 2, the registrants and end users comprise the set of possi-

[5] The necessary client software is called a resolver. Clients like web browsers and email programs will need to update either their included resolver or the operating system resolver upon which they rely.

ble adopters of DNSSEC with whom we are concerned. This set of adopters will receive positive network effects based upon the size of the adopting population. At the moment, however, these potential adopters suffer from a bootstrapping problem.

In this section we will discuss possible interventions that could be used to stimulate DNSSEC deployment, many of which are already underway. The different intervention approaches will be classified using the categories introduced in Section 4. We will then illustrate the intended effect of these interventions using our model.

9.1 Global mandate

The National Strategy to Secure Cyberspace demands that solutions other than global mandates be found: "federal regulation will not become a primary means of securing cyberspace...the market itself is expected to provide the major impetus" [40, page 30]. Even if the administration were to consider such mandates, the US government does not have the power to deploy DNSSEC outside of the United States. On today's Internet, global mandates are not a viable option.

9.2 Partial mandate

On the other hand, governments can create partial mandates for a technology by requiring its adoption by their own agencies and by those who do business with them.

The proposed mechanism to mandate adoption within the US Government is a Federal Information Processing Standard (FIPS). FIPS are issued by the National Institute of Standards and Technology (NIST) pursuant to the Federal Information Security Management Act of 2002 (FISMA) [14]. The US federal government is exploring the use of these standards to mandate the adoption of DNSSEC for agencies with certain security requirements [15].[6]

Another potential target for a partial mandate would be banks and other financial services agencies. These entities both have high security requirements and are usually highly regulated. As a result, their regulatory agencies may have the political will, the desire, and the capability to mandate that these entities adopt DNSSEC.

[6]FIPS Publication 200 [29] mandates that government agencies select the "appropriate security controls" from NIST Special Publication 800-53. The latter publication includes requirements for a "secure name lookup service" [34, page 102] under further guidelines published in NIST Special Publication 800-81 [9].

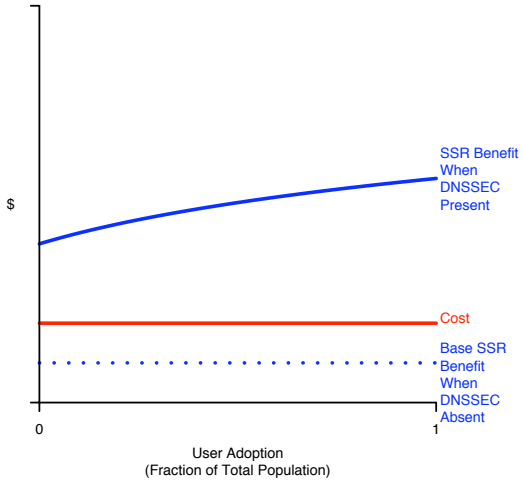

Figure 6: The value of the Service Security Requirement (SSR) record to providers of authenticated Internet services is almost entirely dependent on the use of DNSSEC.

9.3 Bundling complements

The potential for DNSSEC to address the special bootstrapping problems faced by authentication technologies (see Sections 3.2 and 7.5) makes it an ideal candidate for bundling.

9.3.1 Server authentication with SSR

One such complementary technology is the proposed Service Security Requirement (SSR) record [37]. When present in a DNS zone, this record indicates that the domain's web site or other service should always authenticate itself to clients who contact it. If the client can trust these records to indicate which domains have adopted authentication, then we can eliminate the special bootstrapping problems usually faced when adding authentication to a network (see Section 3.2). For example, the SSR record enables operators of web sites to indicate that browsers should only connect to the sites using the authenticated HTTPS protocol.

However, this SSR record is only valuable if clients can retrieve it securely and reliably. Storing an SSR record in the DNS ensures that it can can be retrieved as reliably as the the site's IP address. Using DNSSEC ensures that they can be retrieved securely. While attackers can still prevent SSR records from reaching clients, clients can detect these attacks and take an appropriate action (*i.e.* halting communication).[7] Without the

[7]Unfortunately, neither DNSSEC nor the SSR record can prevent an attacker from stopping the client from communicating

security provided by DNSSEC, an attacker can easily spoof or hide an SSR record. Thus, the SSR is valuable when it is bundled with DNSSEC, but has little value otherwise (Figure 6).

9.3.2 Authenticating email from organizations

Unlike the man in the middle attacks addressed by SSR, email forgery can be performed from anywhere on the network. New DNS records have also been proposed to address forged email. The presence of one of these records indicates that a message should be authenticated, and the contents of the record provide the information needed to authenticate the sender. A Sender Policy Framework (SPF) record indicates that an email can be authenticated by verifying the IP address of the sender [38]: valid IP addresses are listed in the record. Another proposed record, DKIM [3], indicates that the sender supports stronger authentication using public key cryptography: the public key is stored in the record. Unlike SPF, DKIM can protect against forgeries by network intermediaries (men in the middle.)

Both SPF and DKIM records make it more difficult to send spam and phishing emails forged to look as if they came from another domain. Even without DNSSEC, these records can increase the difficulty of forging emails from a participating domain. They are thus unlike the SSR record discussed above, because they provide significant value for both registrants and clients that have not adopted DNSSEC.

DNSSEC and DKIM are complements: working together they can prevent network intermediaries (men in the middle) from forging email from a sender. DNSSEC and SPF are not necessarily complements as they do not prevent forgery by network intermediaries. While DNSSEC can protect the SPF records, a network intermediary can still forge the IP address of a sender that has adopted SPF.

9.4 Facilitating subnetwork adoption

In Section 4.4, we described how to stimulate the adoption of a technology by maximizing the benefits that an organization can derive internally. One way to increase the intra-organizational benefits of DNSSEC is to bundle it with technologies for authenticating the identities of individuals within the organization. DNSSEC signatures can certify the authenticity of keys used to identify individual users. The DNS can then be used to distribute these keys.

For example, many organizations would benefit from securing email communications amongst employees, even if external communications remained insecure. Deploying email encryption and authentication into organizations requires an infrastructure for certifying the authenticity of individuals' keys. Deploying encryption may also require an infrastructure for distributing recipients' keys to senders.[8]

Servers could use DNS and DNSSEC to retrieve and authenticate the keys of users logging into a system. Identity management solutions could use such an approach to provide single sign-on. Because both DNSSEC and identity management require key management, these management costs can be amortized over both technologies.

9.5 Coordination

Another approach to facilitating the adoption of DNSSEC is to coordinate groups of likely early adopters. For example, the Financial Services Technology Consortium (FSTC) helps research and coordinate technology adoption for the financial services industry. Standards bodies like the IETF and governance bodies like ICANN also provide a venue for entities to signal and coordinate their support for DNSSEC.

9.6 Subsidies

Among the fixed costs of deploying DNSSEC are those of software development, which can be grouped into three categories: DNS servers, resolvers, and key management.

To be DNSSEC-compliant, DNS server software must know when and how to return signature records for signed zones. DNS clients (known as resolvers) need code to verify these signatures. Those who operate DNS servers will also need new code to create and manage signing keys and to sign their zones.

The US Department of Homeland Security is working to reduce these development costs by subsidizing open source efforts to create software for all of these tasks. The lower placement of the cost line in Figure 7 illustrates the effect of these subsidies on those choosing whether or not to adopt DNSSEC.

Summary

Of the approaches described above, only global mandates are inapplicable to the effort to deploy DNSSEC. Formal efforts to use partial mandates and to subsidize the development of software are already occurring. Some groups, including the authors themselves, are exploring the value of bundling complements, facilitat-

with the server at all: these technologies do not prevent denial-of-service attacks.

[8] This is not a strict requirement if identity-based encryption is available, such as in the schemes proposed by Adida *et al.* [1, 2].

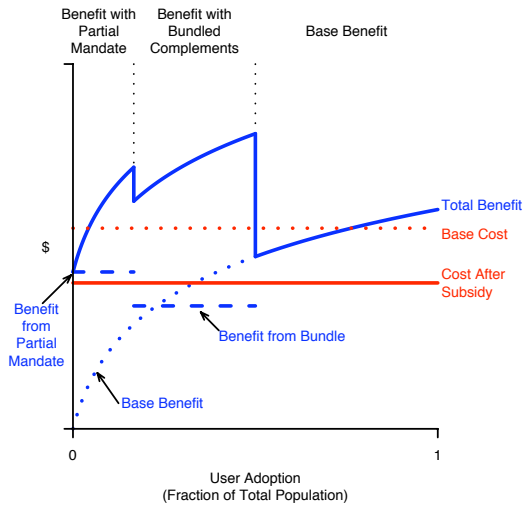

Figure 7: DNSSEC faces a bootstrapping problem. The following approaches are being used to encourage adoption: a partial mandate, bundling complementary standards, and subsidies.

ing subnetwork adoption, and coordinating likely early adopters.

Figure 7 shows the intended effect of these interventions. The US government's partial mandate will incentivize a set of early adopters. Another set of adopters will find value in bundled complements like SSR and DKIM/SPF. Both the partial mandate and the bundled complements are designed to raise the early benefit function over the cost function. The cost function itself is lowered by the software development subsidies provided by the US government.

We anticipate that the combination of these approaches will be sufficient to ensure the widespread adoption of DNSSEC by those who would benefit from it, but the rapidity of this adoption is still in question.

10 Related Work

Two areas of research relate to the economic issues surrounding the adoption of DNSSEC. The first is the body of economics literature on topics like network benefits, network externalities, and adoption externalities. The second includes the less numerous works that discuss the adoption issues specifically surrounding DNSSEC.

10.1 Economics of network adoption

Advocates of replacing insecure protocols with superior ones may find the early literature of network externalities discouraging. The seminal work on network externalities, by Michael L. Katz and Carl Shapiro, uses models to argue that "profit-maximizing firms may fail to achieve complete compatibility in cases where complete compatibility is socially optimal" [19]. In a later paper on the adoption of technologies in the presence of network externalities, Katz and Shapiro conclude that "in the absence of sponsors, the technology superior *today* has a strategic advantage and is likely to dominate the market" [20] (emphasis in original).

More optimistically, our model most closely resembles that of the early work of Joseph Farrell and Garth Saloner [13]. Like our model, that of Farrell and Saloner uses the adoption-set and adopter-equivalence assumptions that lead to symmetry among firms. Whereas our model assumes that individuals can adopt instantaneously, Farrell and Saloner assume that firms make decisions over n periods with parallel decisions and discrete payoffs. Thus, they assume that firms make decisions on the expectation of the outcomes of what others will do, rather than what others have already done. When combined with a common-knowledge assumption, they predict that symmetric firms will always adopt a superior standard—even if an existing standard is already widely deployed. Their model attributes inertia against adoption as a result of incomplete information.

Farrell and Saloner also show that adoption may occur even when it is not socially beneficial; this result was also later illustrated by Dixit [11].

In our study of DNSSEC, we assert that network effects result from the interaction between registrants who provide DNSSEC-compliant directory information and users who adopt DNSSEC-aware clients. Church *et al.* [10] classify network effects that result from such interactions as *indirect* network effects, and they describe how these effects can lead to what they call "adoption externalities".

Stan Liebowitz and Stephen Margolis note that network *effects* are pervasive in the economy. However, they argue that there is scant evidence of network *externalities*, which they define as network effects that lead to market failure [25]. They argue that "property rights, private negotiations, or government interventions...allow the externalities to be internalized". These arguments included a refutation of a number of case studies used to justify the existence of network externalities (such as those presented by Katz and Shapiro [21]).

Saloner and Shepard [36] later provided empirical evidence for the existence of such network externalities by showing that larger banks adopted ATMs more quickly than the smaller ones.

10.2 Deployment of DNSSEC

At least three prior works have discussed adoption issues faced by DNSSEC.

A 2005 report from the National Academy of Sciences [31] concludes that the security of the DNS would be "significantly improved if DNSSEC were widely deployed". The report notes that deployment costs include increased network bandwidth, computation, and key management

A paper by J. Scott Marcus [26] discusses both DNS security in specific and government roles in stimulating technology adoption in general. He illustrates general public policy stimuli that include mandates, seed money, consensus building (coordination), and government leadership. Marcus does not, however, apply these general stimuli directly to the problem of deploying DNSSEC.

In a later paper on DNSSEC deployment, Friedlander *et al.* [15] also explore mechanisms for stimulating innovation. They use historical case studies outside of security to provide context for their discussion: the examples range from the early 20th century deployment of telephones to the end-of-century Y2K investments. The authors suggest the use of NIST mandates for deployment of DNSSEC in government networks—Douglas Montgomery of NIST is among the paper's authors.

Friedlander *et al.* discuss two issues the affect the deployment tipping point: the signing of the root zone and the prevention of zone walking/enumeration. The first issue surrounds the question of who will control the keys to sign the root zone and perform the actual signing. The second issue is that the current standard enables attackers to enumerate every subdomain within a DNSSEC-enabled zone. This shortcoming will be resolved by a modification to the DNSSEC standard that is currently under development.

Christof Fetzer and Trevor Jim argue that VeriSign, the operator of the `com` and `net` registries, has a conflict of interest with respect to DNSSEC. VeriSign is also a Certificate Authority (CA): it sells public key certificates. Fetzer and Jim argue that DNSSEC and public key certificates are substitutes: both can be used to associate a public key with a domain name. One factor in support of this argument is that a standard has been drafted for authenticating certificates using DNSSEC [18], which obviates the need for some CA services. As a result of this tension, Fetzer and Jim believe that VeriSign will be reluctant to adopt DNSSEC for its `com` and `net` registries because doing so could cannibalize its CA business. For VeriSign, both the certificate and registry businesses appear to be of approximately equal size: both generate hundreds of millions of dollars of revenue each year [42, 43].

However, certificates that map domain names to public keys have become commodities, costing as little as $20. This competition has forced certificate issuers to cut costs: they have reduced the amount of inspection they perform before issuing a certificate to an entity. These market forces have also reduced industry confidence in traditional certificates.

By its very name, the Veri*Sign* brand is synonymous with the business of providing cryptographically signed certificates. Rather than allow its signature product to become a commodity, representatives of Verisign have asserted that VeriSign is in the businesses of certifying that public keys belong to *legitimate* businesses (as opposed to simply certifying that an entity owns a domain name). If these assertions are accurate, then deploying DNSSEC would not cannibalize VeriSign's certificate business: it would only kill the portion of the market that Verisign has already ceded to its competitors [12].

11 Conclusion

Deployment of DNSSEC would both enhance the security of the domain name system and ameliorate the bootstrapping problems that impede the adoption of authenticated communications protocols. To achieve these benefits, we must first address the bootstrapping problems that impede the adoption of DNSSEC itself.

We have provided a model for the deployment of security technologies into existing networks. We have used the model to show the effect of strategies and interventions to catalyze the adoption of DNSSEC. Specifically, we have shown how a combination of software development subsidies, partial adoption mandates, and bundled technologies can increase the likelihood and rate of adoption. As these catalysts are put into place, we will soon learn whether it is possible to introduce infrastructure security enhancements into a widely deployed network.

References

[1] Ben Adida, David Chau, Susan Hohenberger, and Ronald L. Rivest. Lightweight signatures for email. In *Proceedings of the DIMACS Workshop on Theft in E-Commerce*, April 2005.

[2] Ben Adida, Susan Hohenberger, and Ronald L. Rivest. Lightweight encryption for email. In *Proceedings of the USENIX Workshop on Steps to Reducing Unwanted Traffic on the Internet (SRUTI 2005)*, pages 93–99, July 2005.

[3] Eric Allman, Jon Callas, Mark Delany, Miles Libbey, Jim Fenton, and Michael Thomas. IETF internet

draft draft-allman-dkim-base-00: DomainKeys Identified Mail (DKIM). http://www.ietf.org/internet-drafts/draft-allman-dkim-base-00.txt, July 9, 2005.

[4] Roy Arends, Rob Austein, Matt Larson, Dan Massey, and Scott Rose. RFC 4033 – DNS security introduction and requirements. http://www.ietf.org/rfc/rfc4033.txt, March 2005.

[5] Roy Arends, Rob Austein, Matt Larson, Dan Massey, and Scott Rose. RFC 4034 – Resource records for the DNS security extensions. http://www.ietf.org/rfc/rfc4034.txt, March 2005.

[6] Roy Arends, Rob Austein, Matt Larson, Dan Massey, and Scott Rose. RFC 4035 - Protocol modifications for the DNS security extensions. http://www.ietf.org/rfc/rfc4035.txt, March 2005.

[7] Derek Atkins and Rob Austein. RFC 3833 – Threat analysis of the Domain Name System (DNS). http://www.ietf.org/rfc/rfc3833.txt, August 2004.

[8] CBR Staff Writer. Secure DNS faces resistance. *Computer Business Review*, December 1, 2005.

[9] Ramaswamy Chandramouli and Scott Rose. Secure domain name system (DNS) deployment guide. Draft National Insitute of Standards and Technology (NIST) Special Publication 800-81, April 11, 2005.

[10] Jeffrey Church, Neil Gandal, and David Krause. Indirect network effects and adoption externalities. Foerder Institute for Economic Research Working Paper No. 02-30, December 19, 2002.

[11] Avinash Dixit. Clubs with entrapment. *The American Economic Review*, 93(5):1824–1829, December 2003.

[12] Joris Evers. Browsers to get sturdier. *ZDNet News*, December 12, 2005.

[13] Joseph Farrell and Garth Saloner. Standardization, compatibility, and innovation. *The RAND Journal of Economics*, 16(1):70–83, Spring 1985.

[14] The federal information security management act of 2002. Title III of Public Law Number 107-347: The E-Government Act of 2002, Signed into law December 17, 2002.

[15] Amy Friedlander, Stephen Crocker, Allison Mankin, W. Douglas Maughan, and Douglas Montgomery. DNSSEC and hardening security in the internet infrastructure: The public policy questions. In *The 33rd Research Conference on Communication, Information and Internet Policy*, Arlington, Virigina, September 23–25, 2005.

[16] Memorandum of understanding between the U.S. Department of Commerce and Internet Corporation for Assigned Names and Numbers. http://www.icann.org/general/icann-mou-25nov98.htm, November 25, 1998.

[17] Internet Corporation for Assigned Names and Numbers. ICANN | fact sheet. http://www.icann.org/general/fact-sheet.html.

[18] Simon Josefsson. RFC 4398 - Storing certificates in the domain name system (DNS). http://www.ietf.org/rfc/rfc4398.txt, March 2006.

[19] Michael L. Katz and Carl Shapiro. Network externalities, competition, and compatibility. *The American Economic Review*, 75(3):424–440, June 1985.

[20] Michael L. Katz and Carl Shapiro. Technology adoption in the presence of network externalities. *The Journal of Political Economy*, 94(4):822–841, August 1986.

[21] Michael L. Katz and Carl Shapiro. Systems competition and network effects. *The Journal of Economic Perspectives*, 8(2):93–115, Spring 1994.

[22] Stephen Kent and Randall Atkinson. Security architecture for the internet protocol. http://www.ietf.org/rfc/rfc2401.txt, November 1998.

[23] Olaf Kolkman. Measuring the resource requirements of DNSSEC. Technical Report RIPE-352, RIPE Network Coordination Centre, September 2005. http://www.ripe.net/ripe/docs/ripe-352.html.

[24] Olaf M. Kolkman. DNSSEC: Raising the barriers against DNS exploits. http://www.nlnetlabs.nl/downloads/RaisingTheBarrier.pdf, November 22, 2005.

[25] Stan J. Liebowitz and Stephen E. Margolis. Network externality: An uncommon tragedy. *The Journal of Economic Perspectives*, 8(2):133–150, Spring 1994.

[26] J. Scott Marcus. Evolving core capabilities of the internet. *Journal on Telecommunications and High Technology Law*, 3:123–164, October 2004.

[27] Mastercard International. Payment card industry data security standard. https://sdp.mastercardintl.com/pdf/pcd_manual.pdf, January 2005.

[28] Bob Metcalfe. Metcalfe's law: A network becomes more valuable as it reaches more users. *InfoWorld*, page 53, October 2, 1995.

[29] National Institute of Standards and Technology, U.S. Department of Commerce. Minimum security requirements for federal information and information systems. Federal Information Processing Standards Publication 200 (FIPS PUB 200), March 2006.

[30] Andrew Odlyzko and Benjamin Tilly. A refutation of metcalfe's law and a better estimate for the value of networks and network interconnections. http://www.dtc.umn.edu/~odlyzko/doc/metcalfe.pdf, March 2, 2005.

[31] Committee on Internet Navigation, the Domain Name System: Technical Alternatives, and National Research Council Policy Implications. *Signposts in Cyberspace: The Domain Name System and Internet Navigation*. National Academies Press, 2005.

[32] Venugopalan Ramasubramanian and Emin Gün Sirer. Perils of transitive trust in the domain name system. In *Internet Measurement Conference 2005*, Berkeley, California, October 19–21, 2005.

[33] Nicholas Rosasco and David Larochelle. How and why more secure technologies succeed in legacy markets: Lessons from the success of SSH. In *Proceedings of the Second Annual Workshop on the Economics of Information Security (WEIS2003)*, College Park, Maryland, USA, May 29–30, 2003.

[34] Ron Ross, Stu Katzke, Arnold Johnson, Marianne Swanson, and George Rogers. Recommended security controls for federal information systems. National Insitute of Standards and Technology (NIST) Special Publication 800-53 Revision 1, March 2006.

[35] Brent Rowe and Mike Gallaher. Could IPv6 improve network security? And, if so, at what cost? *Presentation to the Fourth Annual Workshop on the Economics of Information Security (WEIS2005)* http://infosecon.net/workshop/slides/weis_rump_RG.ppt, June 2, 2005.

[36] Garth Saloner and Andrea Shepard. Adoption of technologies with network effects: an empirical examination of the adoption of automated teller machines. *The RAND Journal of Economics*, 26(3):479–501, Autumn 1995.

[37] Stuart Schechter. IETF internet draft draft-schechter-ssr-01: Storing Service Security Requirements in the domain name system, January 6, 2006.

[38] Sender policy framework. http://spf.pobox.com/.

[39] Security Space. Secure server survey. http://www.securityspace.com/s_survey/sdata/200602/certca.html, March 1, 2006.

[40] The President's Critical Infrastructure Protection Board (PCIPB). The national strategy to secure cyberspace. http://www.whitehouse.gov/pcipb/cyberspace_strategy.pdf, February 2003. The President's Critical Infrastructure Protection Board is now known as the National Infrastructure Advisory Council (NIAC).

[41] U.S. Department of Commerce IPv6 Task Force. Technical and economic assessment of Internet protocol, version 6 (IPv6). http://www.ntia.doc.gov/ntiahome/ntiageneral/ipv6/final/IPv6final.pdf, January 2006.

[42] VeriSign, Inc. Verisign reports fourth quarter and full year 2005 results. http://www.verisign.com/static/036574.pdf.

[43] VeriSign, Inc. Company overview. http://www.verisign.com/stellent/groups/public/documents/presentations/031005.ppt, June 2005.

[44] VeriSign Naming and Directory Services. Registry operator's monthly report for May 2005. http://www.icann.org/tlds/monthly-reports/com-net/verisign-200505.pdf, June 2005.

[45] Visa International Services Association. Products and services: Online merchants. http://www.corporate.visa.com/pd/online_merchants/main.jsp.

[46] Wikipedia. Metcalfe's law. http://wikipedia.org/wiki/Metcalfe's_law.