March 22, 2012

Working Group 5
DNSSEC Implementation Practices for ISPs

Final Report

Table of Contents

Figures

Tables

1 Results in Brief

1.1 Executive Summary

Internet service providers (ISPs) provide the vast majority of U.S. consumers' and businesses' Internet connectivity, making them crucial to the wide deployment of security technologies such as DNSSEC. Working Group 5, "DNSSEC Implementation Practices for ISPs," examined the pros and cons of ISPs' adoption of DNSSEC as knowledge and acceptance of this security technology increases, and attempted to create a set of recommendations for ISPs that do want to adopt DNSSEC.

The Working Group recommends that:

- ISPs implement their DNS recursive nameservers so that they are at a minimum DNSSEC-aware, as soon as possible.
- Key industry segments, such as banking, credit cards, healthcare and others, sign their respective domain names with DNSSEC.
- Software developers, such as those creating operating-system, web-browser, and other Internet-focused applications, study how and when to incorporate DNSSEC validation functions into their software.

These recommendations are covered in greater detail in Section 6.3, "Recommendations."

2 Introduction

Working Group 5, "DNSSEC Implementation Practices for ISPs," was asked to examine "best practices for deploying and managing the Domain Name System Security Extensions (DNSSEC) by Internet service providers (ISPs). In addition, the Working Group shall recommend proper metrics and measurements that allow for evaluation of the effectiveness of DNSSEC deployment by ISPs."

This Working Group enjoyed input from a broad range of experts, from major ISPs and from non-ISP organizations, who were able to comment knowledgeably on DNSSEC's importance to the security of the DNS. The result is a final report that addresses the full range of ISP and other concerns about DNSSEC deployment, and helps clarify existing and potential obstacles to same along with potential solutions.

While some ISPs have deployed DNSSEC internally and for their customers, most have not, and this Working Group's task was to determine the pros and cons of ISP adoption of DNSSEC and recommend how ISPs might best achieve this task.

In brief, ISPs' desire for security is counterbalanced by concerns about:

- Effectiveness
- Ability to resolve customer issues in failure
- The expense and increased workload incurred by adoption
- Concerns that a DNSSEC-enabled system is less forgiving than the current largely non-DNSSEC enabled system, which may pass costs for others' inadequate operation or maintenance of their DNSSEC signatures to ISPs
- Other threats, including potential DNSSEC enablement of more intense DNS amplification attacks

The fourth concern listed above was highlighted by the recent episode in which nasa.gov allowed its DNSSEC signatures to expire, and customers were unable to reach nasa.gov. Rather than complaining to NASA, consumers typically contacted their ISPs, costing those companies customer-service dollars.

It was against this backdrop that the Working Group's deliberations took place.

2.1 CSRIC Structure

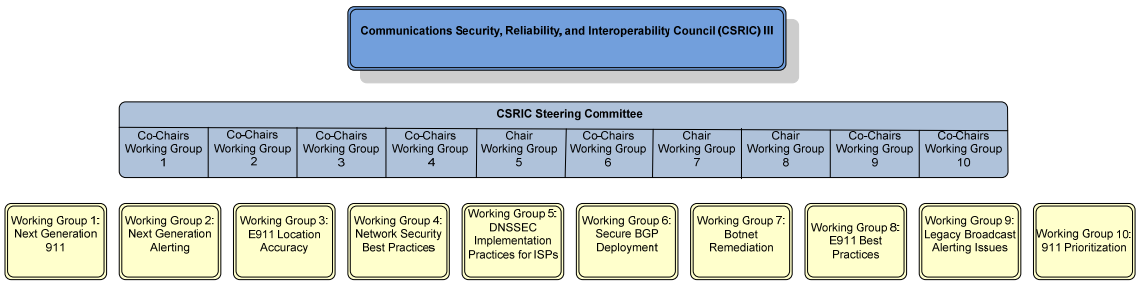

Figure 1 – CSRIC structure

2.2 Working Group 5 Team Members

Working Group 5 consists of the members listed below.

Name	Company
Chair: Steve Crocker	Shinkuro
Daniel Awduche	Verizon
Michael Burns	Cablevision
Warren Kumari	Google
Matt Larson	Verisign
Jason Livingood	Comcast
Daniel Mason	CenturyLink
Chris Mikkelson	CenturyLink
Doug Montgomery	NIST
Russ Mundy	Sparta
Rod Rasmussen	Internet Identity
Brian Rexroad	AT&T
Chris Roosenraad	Time Warner Cable
Todd Szymanski	Sprint
Matt Williams	Cox
Suzanne Woolf	ISC

Table 1 - List of working group members

3 Objective, Scope and Methodology

3.1 Objective

From the description of Working Group 5 on fcc.gov (available for download at http://www.fcc.gov/pshs/advisory/csric3/wg-descriptions_2-17-12.pdf; the announcement pertaining to all 10 working groups is at http://www.fcc.gov/encyclopedia/communications-security-reliability-and-interoperability-council-iii):

"The Domain Name System Security Extensions (DNSSEC) are widely recognized as the best hope for improving the long-term security of the Internet's critical domain name system. Standards for DNSSEC are now mature and implementation has begun in the government as well as the enterprise sector."

"This Working Group shall recommend the best practices for deploying and managing the Domain Name System Security Extensions (DNSSEC) by Internet service providers (ISPs). In addition, the Working Group shall recommend proper metrics and measurements that allow for evaluation of the effectiveness of DNSSEC deployment by ISPs. In addition to any other metrics, the Working Group shall address the following: availability of a zone, verification of received data, and validation of verified data. Finally, the Working Group shall recommend ways for the ISP community to demonstrate their intent to deploy DNSSEC, possibly by way of a voluntary opt-in framework."

3.2 Scope

While this Working Group focused solely on the use of DNSSEC as a means to protect against domain-name fraud, it should be noted that many other substantial threats regularly compromise the integrity of the DNS and figure in domain-name-fraud attacks. The following considerations were considered outside of scope for this working group:

a) Alternative approaches and countermeasures to protect against domain-name fraud were discussed but are not considered in this report.
b) Only the role of ISPs was considered in implementation of DNSSEC (particularly work toward ISP validation in recursive resolvers provided to ISP subscribers). The potential roles of alternative DNS resolver providers—e.g. DNS registrars, authoritative DNS providers—were discussed but are not considered in this report.

This Working Group's scope of research was also limited by time (a mid-March deadline to deliver its recommendations) and geography (members dispersed throughout the continental U.S.). However, since meetings could be held telephonically and via an easily used e-mail mailing list, these limitations are not thought to have had great effect on the Working Group's research.

3.3 Methodology

The Working Group proceeded along three stages, each of which consisted of one or more steps, in conducting its research and analysis:

Methodology

- Form working group with expertise
- Query working group regarding hurdles, challenges, etc.

Analysis

- List specific issues, formulate approach for each
- Note details of issue resolution

Findings

- Collate results of analysis
- Consensus and recommendations

The remainder of this section will focus on the discussion of Methodology; Analysis and Findings are covered in Sections 5.1 and 6.2, respectively, while Recommendations are in Section 6.3.

This Working Group enjoyed a broad range of participants among both ISPs (D. Awduche, M. Burns, J. Livingood, D. Mason, C. Mikkelson, B. Rexroad, C. Roosenraad, T. Szymanski, M. Williams) and non-ISP experts who have been part of the DNSSEC deployment effort (S. Crocker, W. Kumari, M. Larson, D. Montgomery, R. Mundy, R. Rasmussen, S. Woolf).

The Working Group was queried via a series of teleconferences and e-mail exchanges designed to elucidate the issues confronting ISPs as they decide whether and when to implement DNSSEC as part of their service offerings.

This Working Group focused specifically on the hurdles and challenges for ISPs to adopt DNSSEC and focused further on a particular aspect of ISP deployment: validation. This contrasts with providing signed DNS service to domain names that ISPs host or their own domains. Out of that process came a number of key drivers and challenges for ISPs that seek to adopt DNSSEC:

Drivers

- Protection against cache poisoning
- Security increasingly resonates with customers
- DNSSEC can be a market differentiator for early adopters
- DNSSEC may help ISPs avoid some costs if a cache poisoning attack occurs
- ISP DNSSEC awareness in DNS recursive nameservers is necessary for end-user validation (e.g., DANE[1])

Challenges

- Unclear U.S. government policy regarding use of DNS redirection to block botnets and advanced persistent threats (APTs) as well as other malicious or illicit activity
- DNSSEC efforts may create an inaccurate impression of DNS infrastructure security
- Loss of nonexistent domain (NXDOMAIN) revenues
- Perceived impact to Internet service reliability
- Poor WHOIS contact information complicates troubleshooting

[1] https://datatracker.ietf.org/wg/dane/charter/

- Effectiveness and impact of DNS amplification attacks may be exacerbated by DNSSEC deployment
- Possible unanticipated abuses of DNSSEC-enabled services for attacks
- Lack of direct financial benefit from DNSSEC adoption
- More signed domains needed
- End-system validation obviates the need for ISP validation
- DNSSEC may increase operating costs for ISPs and other DNS service providers
- Content distribution network service providers such as Akamai may face additional challenges in implementing and managing DNSSEC (although they may gain advantages as well)
- Alternate providers of DNS services may compete with services provided by ISPs, and those providers may not have any intention of implementing DNSSEC

Finding potential solutions to the above challenges will require follow-on activity.

4 Background

4.1 Brief Overview of the DNS

The Domain Name System (DNS) is a distributed hierarchical database which contains a listing of Internet resources and various types of information associated with those resources. Although the DNS has a variety of uses, its most important function is to bind user-friendly names of Internet resources to corresponding IP addresses of the systems that host those resources. This allows end users to conveniently depict and access Internet resources using recognizable names. The DNS also creates a logical linkage between the name of an Internet resource and its IP address, allowing a resource to retain the same name, even though its IP address and point of attachment to the network changes over time.

4.1.1 Structure of Domain Names

A domain name denotes an Internet resource, such as a website, an email address, a database server, or any machine or service that is accessible through the Internet. Domain names are hierarchically organized in a tree structure as shown in Figure 2. Each node in the hierarchy represents a domain and has a label associated with it. A domain may be the parent of subordinate domains (subdomains). The root of the DNS tree has no formal name, but is generally referred to as the DNS root domain. Below the root domain are the top-level domains (TLDs) which comprise the first-level group of domains. The TLDs include generic top-level domains (gTLDs) such as .com, .net, .org, .edu, etc. and country code top-level domains (ccTLDs) such as .us, .uk, .br, .de, .se and so on.

The next subordinate levels in the tree structure include the second-level domains, third-level domains, fourth-level domains, etc. There can be up to 127 levels of subordinate domains in the hierarchy.

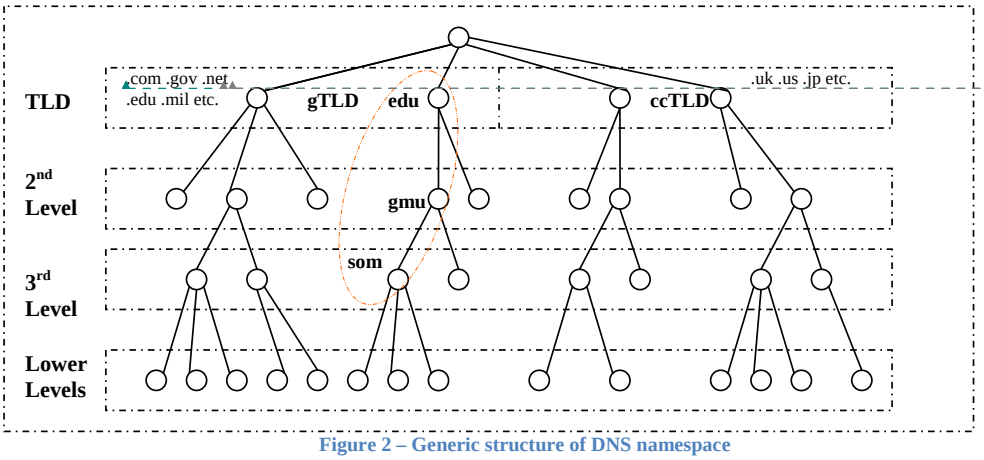

Figure 2 – Generic structure of DNS namespace

The administration of the DNS is decentralized. Each domain or subdomain can be managed by a separate organization. A domain administrator can delegate management of some of its subdomains to other entities—and this domain decomposition and delegation process can be

enacted recursively. Parent domains maintain only pointers to servers that contain information about their subdomains so that DNS queries can be referred to the appropriate data sources. Each autonomously managed domain is called a zone. The syntax of a domain name consists of a sequence of labels (designating nodes in the namespace) separated by dots. Essentially, a domain name is an index entry in the DNS database. For example “som.gmu.edu” refers to the “som” subdomain under “gmu” in the “edu” gTLD.

The DNS database is distributed across a very large number of geographically dispersed nameservers that are managed by independent organizations. Each nameserver contains information pertaining to a subset of the DNS namespace and pointers to other nameservers that can lead to information in other parts of the database. Nameservers store data associated with domain names in resource records (RRs). Broadly speaking, there are two types of nameservers: (1) authoritative and (2) caching. An authoritative server has complete knowledge about a subset of the domain namespace, while caching servers improve query response time by locally caching a subset of global DNS data for a specified time interval.

4.1.2 Operation of the DNS

Operation of the DNS is based on a client-server model. Each user device contains a resolver, which is a local agent that sends and receives DNS queries on the user's behalf. The device will also have one or more designated DNS nameservers whose IP addresses are configured either automatically (e.g. using the Dynamic Host Configuration Protocol [DHCP]) or manually by the user or a local administrator.

From a user's perspective, the operation of the DNS proceeds as follows. First, a user or front-end software inputs a URL (e.g. a website address) into a network application (e.g. a web browser). The resource name is sent to a local resolver on the user’s device. If the resolver has a locally cached copy of the domain's IP address and other pertinent RR details for the requested resource, it passes that data back to the application. Otherwise, the resolver will query a designated nameserver. If the designated nameserver has a cached copy of the required RR, it sends the information back to the user’s resolver. Otherwise, how the server behaves will depend on whether it is configured with DNS recursion:

- If the server is NOT configured with DNS recursion, it will send the user resolver a referral to another nameserver in the DNS hierarchy. The resolver will then query the new server and this process occurs iteratively until the requested IP address and associated resource record information are obtained from a nameserver in the system.
- On the other hand, if the designated nameserver is configured with recursion, it serves as an agent for the user and recursively submits queries to other nameservers in the DNS hierarchy (each server will either furnish the RR information or issue a referral to another server). Eventually, the recursive server will fetch the information from a nameserver in the system and pass it back to the end user's resolver.

4.2 What Is DNSSEC?

After an end user inputs an easily remembered URL such as www.examplebank.com into a web browser, a nameserver translates it into a string of numbers such as 192.168.0.3 or 2001:db8:ac10:fe01::. Usually, the browser then forwards the user's request to the server that

192.168.0.3 represents, which sends back examplebank.com's web page.

However, the web browser must first ask some other authority what IP address(es) www.examplebank.com translates into, since domain names change or are transferred, are created or destroyed, or hardware and software are updated in such a way (even dynamically to adjust to load or outages, or geographically to direct users to the closest host) as to make changing the address(es) necessary.

A lack of security in the Domain Name System (DNS) means that criminals and others can intercept the request for examplebank.com's address and send the request to their own servers, say 192.144.1.2, and send back an unwanted page (typically an advertisement) or worse, a carefully crafted fake examplebank.com web page that captures the user's innocently input credentials. This is called a "man in the middle" attack and the user may not realize they have given their information to criminals until it is too late.

Domain Name Security Extensions (DNSSEC) addresses this problem; it is an enhanced level of security that allows websites (and other applications and protocols) and Internet service providers (ISPs) to validate domain names to ensure they are correct and have not been tampered with. For example, with DNSSEC, a domain name such as examplebank.com can be cryptographically signed in the Domain Name System (DNS). Then, when an end user tries to connect to that website, an ISP's DNS servers will check that the domain name and its security signature are verified and have not been tampered with by hackers. End users will then only be connected if this security verification has been passed. (This transaction occurs so quickly that end users do not even notice that it is being performed.)

So when DNSSEC is used for the examplebank.com domain, the user's ISP makes two requests: to .com to determine whether examplebank.com should have a DNSSEC signature (i.e. whether examplebank.com is "signed"), and then if that answer is yes, to examplebank.com for its signature. (If .com's answer is negative, that answer is also returned to the ISP.) The ISP then asks examplebank.com for its signature and then verifies that signature using examplebank.com's public key. It then forwards examplebank.com's content back to the end user as a DNSSEC-validated response.

In its most highly developed form, individual users will harness DNSSEC-enabled applications to perform this function as well, pushing responsibility for DNS security all the way to the edge of the worldwide network.

4.3 Current State of DNSSEC Deployment

4.3.1 Signing

4.3.1.1 Root and Registries

In July 2010, the Internet Corporation for Assigned Names and Numbers (ICANN), the organization that administers the global DNS and IP addressing for the Internet, signed the global root of the DNS. Subsequently, top-level domains (TLDs) such as .com, .net, .org, .edu, and .gov were signed in 2010 and 2011. (Crucially, these combine to account for a very high percentage of the Internet's zones.) New generic TLDs (gTLDs) issued by ICANN will be required to support DNSSEC from launch.[2] Almost all the major country-code top-level

[2] http://www.afilias.info/blogs/roland-laplante/icann-makes-progress-new-tlds-brussels

domains (ccTLDs) are signed.

In specific policy statements[3] and subsequent technical guidance,[4] the U.S. federal government has mandated that all civilian federal agencies adopt DNSSEC and to that end, the National Institute for Standards and Technology (NIST) now monitors DNSSEC adoption within the .gov zone on a weekly basis.[5] The Department of Defense's .mil domain is in the process of being signed.

4.3.1.2 Individual Domain Names

Once a TLD is signed, such as .com, then a domain name such as examplebank.com can be signed and will then enjoy a full chain of trust up to the global DNS root. As of early February 2012, Verisign Labs reported that over 5,000 .com domains and over 2,000 .net domains have been signed,[6] and the following two figures document increasing adoption over the past two years.

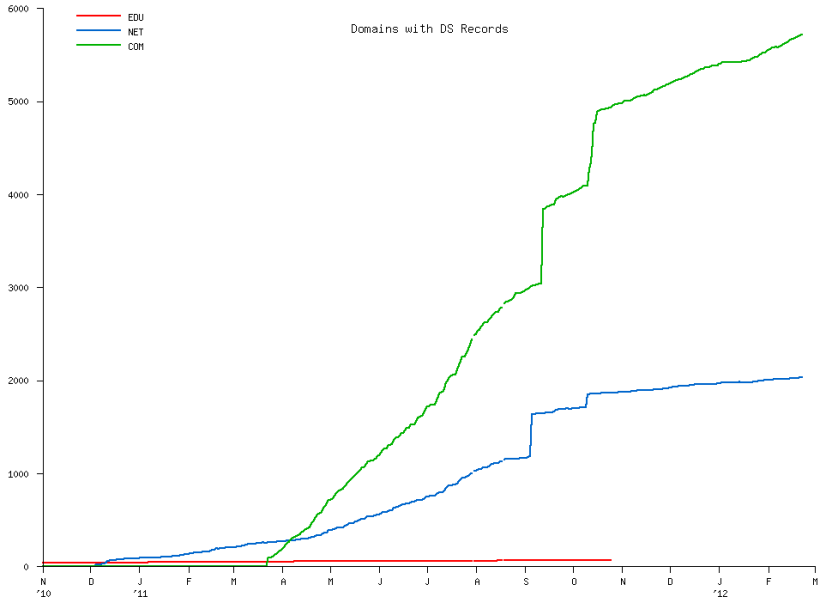

Figure 3 – Number of domains in the .edu, .net and .com TLDs that have DS records (Source: Verisign)

[3] For example, OMB M-08-23, downloadable at http://www.whitehouse.gov/omb/memoranda/fy2008/m08-23.pdf
[4] Including NIST SP-800-53, NIST SP800-57P3 and NIST SP-800-81.
[5] http://usgv6-deploymon.antd.nist.gov/cgi-bin/generate-gov
[6] http://scoreboard.verisignlabs.com/

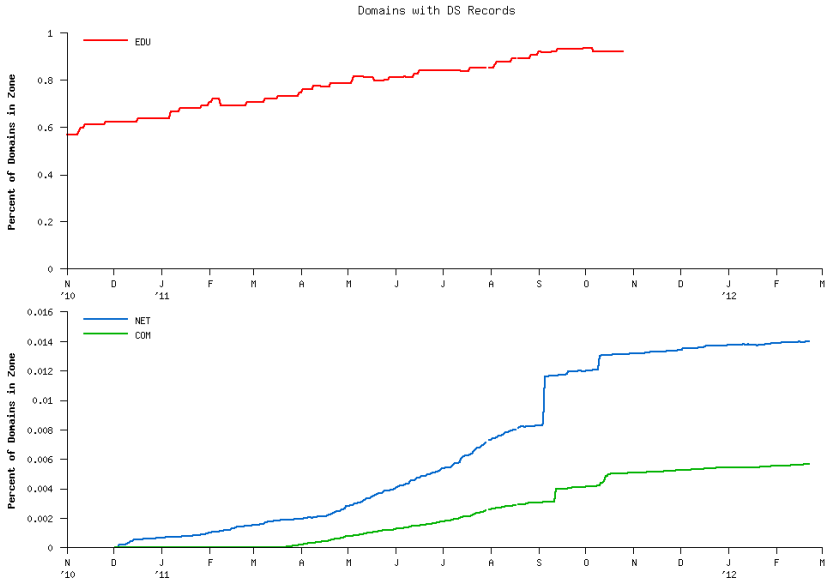

Figure 4 – Percentage of domains in the .edu, .net and .com domains that have DS records (Source: Verisign)

4.3.2 Validation

ISPs, the primary but not exclusive operators of recursive DNS resolvers, are in the early phases of DNSSEC deployment. To date, while some ISPs have completely migrated to DNSSEC,[7] the vast majority are either in the midst of testing and/or deployment planning or have not yet started work on DNSSEC validation. Best practices identified in this report reflect the work of major ISPs that have varying perspectives on security, DNS, and DNSSEC.

4.3.3 Client Support

Client support for DNSSEC, such as in browsers, is available via plug-ins or other software but is not yet mainstream or widely adopted. However, these solutions may have some of the same vulnerabilities as Secure Socket Layer (SSL) implementations in browsers, as when end users ignore security warnings associated with failed SSL certificate validation and either "click through" to a site anyway, or reflexively direct questions to their ISP rather than to the site in question.

4.4 How ISPs Use DNSSEC

The majority of Internet users have their computers configured to use their ISP's recursive DNS resolvers. This means that ISPs operate the DNS servers that most people utilize, and also means that good DNSSEC adoption by end users hinges upon ISP adoption of DNSSEC.

ISPs may hold two separate DNSSEC-related roles: signing and validating. The role of signing

[7] http://blog.comcast.com/2012/01/comcast-completes-dnssec-deployment.html

applies to cryptographically signing their own domain names, as well as their hosted customers' domain names. The role of validating refers to authenticating each answer down the line via cryptographic signatures, so that the end user's request is resolved correctly (and not redirected to a phishing, criminal, or other website).

4.4.1 Validating Names

In the process of performing DNSSEC validation, an ISP relies on and validates a complete chain of trust, from signed root to signed TLD to signed domain name (and perhaps beyond). When an ISP is ready to perform validation, it may choose to enable this function on existing DNS recursive resolvers or to install new servers and migrate end users to those new servers. DNSSEC involves a "heavier," i.e. somewhat larger, packet size than plain DNS, so some upgrades or changes to DNS infrastructure likely will be necessary. In the latter case, ISPs can automatically update DNS server IP addresses, such as via DHCP lease updates or other configuration changes.

(When performing validation, an ISP's DNS recursive resolver typically sets the DNSSEC OK bit ["DO bit"] to 1, which indicates that DNSSEC is understood and requested by the resolver. If a name validation is successful, the query is answered for an end user. If a name validation fails, the query will result in a failure [response of "SERVFAIL"].)

4.4.2 Categories of DNSSEC service

Only a small percentage of ISPs currently validate DNS queries, and it is, indeed, an objective of this report to recommend steps toward increased implementation of DNSSEC validation for ISPs. Toward that end it is useful to categorize the levels of DNSSEC-related service provided by an ISP. Table 2 is a list of four categories; a preliminary survey suggests that very few ISPs are in category A, with Comcast being the most visible U.S. example. Some ISPs are in category B, and most are in category C. These categories are relevant to the analysis, findings and recommendations below.

Category	DNSSEC service
A	Fully validating, where the ISP performs all validation on the end user's behalf (although the validating server can be configured in two distinct modes[8])
B	DNSSEC-aware but non-validating, so that end systems may validate but the intermediate resolver doesn't
C	DNSSEC-unaware but able to handle large (e.g., EDNS0, IPv6) packets
D	DNSSEC-unaware but unable to handle large packets (this category is largely obsolete but is included for completeness)

Table 2: Categories of DNSSEC validation service levels in ISPs

[8] A fully validating resolver can be configured in two modes: strict or permissive. Strict mode will prevent an answer from being returned to a client when validation fails. Permissive mode will return a non-validated answer to the client, but will not set the authenticated data flag. Permissive mode only offers security protections for DNSSEC-aware client-side software, but does not prevent access to non-DNSSEC-aware applications when DNSSEC validation fails, either due to misconfiguration or security abuse. Permissive mode is considered a transitory setting with the end goal considered to be a strict mode.

4.4.3 Signing Names

ISPs may potentially set up and sign large numbers of their customers' domains (i.e., act as registrars) and may also handle validation on their behalf. In addition, large ISPs typically own many domains of their own, and must maintain and sign them. These two classes of behaviors constitute separate businesses.

In order to perform cryptographic signing of domains owned or managed by an ISP, the ISP's authoritative nameservers must sign a domain, such as examplebank.com. In addition, the ISP must work with its registrar to insert a delegation signer ("DS") record into a registry such as .com in order to create a trust linkage between the domain (examplebank.com) and the TLD (.com).

4.4.4 How End Systems Use DNSSEC

Validation at the ISP is not the total picture of validation. End systems may rely on the ISP to perform validation, or they may rely on the ISP's DNS service to fetch the DNSSEC keys and signatures so the end system can do the validation itself. The behavior of the end system falls into one of five different levels, detailed in Table 3.

Level	Query	End System Behavior
1	RD off, EDNS0	End system acts as its own recursive resolver and goes around ISP to query authoritative nameservers directly
2	RD, EDNS0, DO, CD	End system does its own validation but depends on the ISP's resolver to fetch and return the chain of keys and signatures
3	RD, EDNS0, DO	End system depends on the ISP to carry out validation
4	RD, EDNS0	End system does not request or depend on validation
5	No options	End system conforms to RFC 1034 and does not recognize DNSSEC related options. Deprecated and not recommended for current day operations, but included here for completeness

Table 3: End-system levels of DNSSEC behavior

These levels, and the implications for end users of their interaction with how ISPs validate names above, will be discussed later in this report since they affect the rationale for ISPs to deploy DNSSEC.

5 Analysis, Findings and Recommendations

5.1 Analysis: What Are ISPs' Key Drivers and Challenges Regarding DNSSEC?

5.1.1 Key Drivers

5.1.1.1 Protection Against Cache Poisoning

The Kaminsky Vulnerability,[9] disclosed in early 2008, exploits a fundamental flaw in the DNS protocol. This flaw means an attacker could "poison" the cache of any DNS recursive resolver, replacing the correct IP addresses of a particular domain name with the addresses of malicious servers under their control. If the user intended to access examplebank.com, the attacker could direct them to a malicious server that they controlled and the web browser's address bar would still say www.examplebank.com. The end user would have no idea they were connecting to a different server than intended unless they paid attention to SSL protections (or the lack thereof) provided by the site.

Thus, unlike with phishing, an end user would still see the real and intended domain name in their client software or on their device. This also means that most end-user software intended to prevent access to known malware and phishing sites would not work, since a seemingly valid site address was being accessed.

Every DNS recursive resolver in the world had to be patched against the Kaminsky Vulnerability, but the fixes implemented in 2008 provided only temporary protection and may become less effective over time without detection and blocking of attack attempts. Furthermore, attackers have not given up trying, as many ISPs observe regular attempts to poison the caches of their DNS recursive resolvers.

In addition, there are examples of rogue systems administrators in ISP networks poisoning cache records as well,[10] meaning both internal and external actors can execute cache-poisoning attacks. The only permanent protection is to implement DNSSEC, which protects against both of those scenarios.

5.1.1.2 Security Increasingly Resonates with ISP Customers

Security-related features and protections increasingly resonate with ISP customers. At least one ISP has noted that, in surveys of new customers, security is rated roughly on par with speed when choosing an ISP. The existence and increasing sales of supplemental security services by ISPs and third-party providers is further evidence of this, and represents a shift from several years ago when concerns about speed and pricing largely dominated such survey responses. Clearly, DNSSEC may be a part of a wider security strategy for ISPs, and can be explained as such to current and prospective customers.

[9] http://en.wikipedia.org/wiki/Dan_Kaminsky - Flaw_in_DNS
[10] See relevant stories at
http://www.securelist.com/en/blog/208193214/Massive_DNS_poisoning_attacks_in_Brazil,
http://threatpost.com/en_us/blogs/major-dns-cache-poisoning-attack-hits-brazilian-isps-110711 and
http://www.thetechherald.com/articles/Insider-arrested-after-DNS-poisoning-attack-targets-Brazilian-ISPs.

5.1.1.3 DNSSEC Can Be a Market Differentiator for Early Adopters

ISPs that adopt any particular security initiative early on can use that adoption as a market differentiator. For example, Comcast describes DNSSEC, malware detection, and other security tools and protections as part of its Constant Guard™ system. This system is then used in marketing and communications with prospective customers, as well as in communications with existing customers. Over time, ISPs that do not offer comparable security protections and features, or that are the target of exploits such as DNS cache poisoning, may suffer in the marketplace.

5.1.1.4 DNSSEC Can Help to Avoid Some Costs

ISPs are also concerned about maintaining their reputation and with minimizing costs. An ISP that experienced a successful DNS cache-poisoning attack could experience damage to its reputation and goodwill as well as to its trusted relationship with customers. Such damage could be especially acute if an attack affected non-DNSSEC-adopting ISPs but not their DNSSEC-adopting competitors. While reputation and goodwill are generally difficult to quantify, it is likely that a non-adopting ISP might lose current or potential customers to competitors, affecting its future growth and health.

However, ISPs that were attacked might also experience direct costs in the form of increased customer-support calls. For example, if the DNS records for a major website such as Facebook.com or Google.com were "poisoned" so that users were directed to a page saying “You’ve been hacked,” those users would likely direct a significant volume of telephone calls to their ISP's support line. Since both of those sites are among the top five most-popular sites at peak usage times, the resulting costs might be quite high since some number of customers will call in to complain and each of those calls will cost an ISP money.

Should any DNS cache-poisoning attack be sustained for several days, the target ISP might also have to offer service credits to some customers on top of customer-service costs, which might also be a substantial cost.

Finally, there is the risk of business-to-business complaints. This could involve simple private complaints from the affected domain to the ISP in question, but could also escalate to complaints on the affected domain’s blog, press releases by affected customers, and action by applicable governmental bodies such as the FCC. A complaint might be expressed more formally as a lawsuit, particularly if the ISP could have or should have done something that it did not do (for example, implement DNSSEC).

The risk of escalations in any such complaints may rise or fall depending upon the popularity of the affected domain, and whether the ISP is perceived as acting aggressively enough should the affected domain be thought to compete with one of the ISP's own service offerings. For example, if ISP A offered video services but was perceived as not acting quickly to solve an extended cache poisoning of a video-service domain such as netflix.com, complaints might rapidly escalate toward litigation.

5.1.1.5 ISP Adoption of DNSSEC Is Vital for End-Point Validation

As noted in Tables 2 and 3, there are four categories of DNSSEC “maturity” for an ISP’s DNS service, and there are five levels of DNSSEC maturity for the end user’s system. These

combinations are displayed in Table 4, which shows those combinations that will lead to successful validation.

	ISP Categories			
End-System DNSSEC Behavior	A	B	C	D
1	NA	NA	NA	NA
2	V or D	D	P	C
3	V	D	P	C
4	V	P	P	C
5	P	P	P	P

Table 4 - Possible outcomes of ISP and end-user validation

Key
V = Validated
D = DNSSEC chain returned
P = Plain DNS, no validation or DNSSEC related records
C = Compatibility mode (essentially the same as P)
NA = Not applicable

This table shows that an ISP must be running a DNSSEC-aware resolver, i.e. category B, even if it doesn't do validation itself, for DNSSEC-validating applications to work. The only other possibility is for the end system to be its own recursive resolver and go around the ISP's DNS service completely. Thus, it's clear that ISP service must be at level A or B as described in the table above. Thus even if one believes the ultimate goal is to have end systems do their own validation, substantial support is required at the ISP. Otherwise, the end system would have to act as its own recursive resolver.

5.1.2 Key Challenges

5.1.2.1 Unclear U.S. Federal Government Policy Regarding Use of DNS Redirection for Security

Some federal programs are recommending the use of DNS redirection as a means to protect organizations against other substantial security threats, including advanced persistent threats (APTs) and botnet command-and-control (C&C) channels.[11] One benefit of DNS redirection as a security measure is the ability to help identify and notify or remediate victims of attack quickly while preventing the attack from being successful. (Such measures are necessary when the associated DNS registrars are uncooperative.) Due to the hierarchy of DNS, identifying the infected devices in sophisticated attacks is much more effective when redirecting domains. However, domain redirection is generally inconsistent with implementation of DNSSEC validation—whose primary purpose is to prevent such redirection—so DNS service providers would benefit from either a clarified redirection policy or alternative approaches to avoiding DNS fraud.

[11] See the PDF at http://www.dhs.gov/xlibrary/assets/privacy/privacy_nppd_jcsp_pia.pdf.

5.1.2.2 Inaccurate Impression of DNS Infrastructure Security

DNSSEC deployment will not prevent the predominant cases of DNS fraud that occur today. Since 2008, there have been no documented cases in which a strong argument could be made that DNSSEC would have prevented or detected an attack. Examples of more prevalent types of domain-name fraud include:

- Unauthorized access to authoritative DNS service registrars, which would likely provide those users' authorized DNSEC signatures. These events occur due to weaknesses in the authorization or security of user interfaces with authoritative DNS service providers
- DNS Changer malware, which will bypass trusted DNS infrastructure and use its own rogue DNS infrastructure, bypassing DNSSEC validation or blocking DNSSEC forwarding
- Phishing websites where domain names are made to appear legitimate via subtle misspellings, which deceive users into becoming victims of identity theft

5.1.2.3 Loss of Non-Existent Domain Revenues

Some ISPs perform redirection of non-existent domain names to a search portal that carries advertising and enables the ISP to monetize clicks on some links (commonly known as NXDOMAIN redirection). It is generally thought that NXDOMAIN redirection is incompatible with DNSSEC, and, indeed, the specific purpose of DNSSEC is to assure the end system that the responses to its queries are the ones provided by the authoritative servers and have not been modified in transit or in intermediate caches. A more detailed examination the interaction of DNSSEC and redirection of NXDOMAIN suggests there is some room for coexistence.

The DNSSEC protocol includes proof of non-existence of a name using NSEC or NSEC3 records and, in general, any name that falls within the span of one of these records is affirmatively known not to exist. However, the NSEC3 record includes an option called “opt-out,” which treats names that exist but are unsigned the same as non-existent names. This feature of the NSEC3 record greatly reduces the cost of introducing DNSSEC in large zones, and has the side effect that an ISP could rewrite an NXDOMAIN response that occurs within an NSEC3 span that has the opt-in bit set, and return an unsigned address record. We do not know of any implementation that works this way, so it would have to be demonstrated and tested, but this may be a possible path for those ISPs that wish to deploy DNSSEC validation but also wish to continue using NXDOMAIN redirection.

We note that NXDOMAIN redirection is not without controversy. Further, revenue from NXDOMAIN redirection has been steadily declining for several years, a process that may also be accelerated by web browsers that increasingly display their own search pages when a user enters a domain name that does not exist.

5.1.2.4 Validation Failures Due to Misconfiguration

One challenge during the time when only some ISPs perform DNSSEC validation is that some domains may not properly sign their domain, may mismanage key rollovers, or may make other DNSSEC-related configuration errors. This will very likely render their domain unreachable via those ISPs that perform DNSSEC validation (though ISPs may have tools that provide some recourse; see Section 6.1.4, "Negative Trust Anchors," for further discussion). End users may perceive this as the ISP nefariously blocking access to the misconfigured domain name, as

Comcast observed during the DNSSEC validation failure of the nasa.gov domain.[12]

However, as more ISPs perform DNSSEC validation and domain owners consequently acquire more experience with signing, DNSSEC simply becomes another aspect of DNS configuration to manage in the normal course of business. Over time, this will gradually reduce the frequency of validation failures due to misconfiguration and reduce the impact of any such failures on ISPs, since domain owners are ultimately responsible for ensuring their DNS records (A, CNAME, MX, DS, RRSIG, etc.) are configured correctly.

5.1.2.5 Poor WHOIS Contact Information Complicates Troubleshooting

In the course of troubleshooting misconfigured domains, ISPs or their customers may attempt to use WHOIS data to contact the domain in question. This could prove challenging as many domains do not update WHOIS contact data; that data can be difficult to find since it is not centralized; and some domains hide their contact data in WHOIS.

The ICANN WHOIS Review Team studied the WHOIS process and how to improve it during 2011, and issued a draft report in December.[13] That report found that policies, practices and responsibility for the accuracy of WHOIS were diffuse, outdated and occasionally contradictory. It recommended (among other things) that the ICANN Board act to create a single authoritative WHOIS policy document and take other actions to improve WHOIS data accuracy and access.

5.1.2.6 Effectiveness and Impact of DNS Amplification Attacks Are Exacerbated by DNSSEC

DNS amplification attacks are distributed denial-of-service (DDoS) attacks in which small queries are engineered to provoke much larger UDP responses, which are then misdirected at a designated target.[14] Technically, DNS amplification attacks are a problem with the DNS protocol and not with DNSSEC, but they have some characteristics that will make DNS amplification attacks a greater threat in a world of widespread DNSSEC deployment:

- Widespread deployment of DNSSEC significantly increases the range of standard queries that can be used in attacks. With many query types to chose from, attackers may significantly diversify their attacks and make mitigation by filtering nearly impossible. (Attackers can also create their own DNS TXT record to provide amplification, but these are easy to recognize and filter for.)
- Before DNSSEC, “open” DNS resolvers were the primary means of amplifying attacks since the variety of queries that provide significant amplification is limited. With substantial DNSSEC services deployed, however, any domain that is signed will significantly expand the scope of amplifying servers on the Internet since any DNSSEC-enabled authoritative server may act as an amplifier.
- Typically, filters are used to identify the source addresses of DDoS attack activity. This practice comes into play during DNS amplification attacks where authoritative resolvers are used, meaning authoritative DNS services could be blocked from the victims’ network. Victims may be inclined to complain to upstream ISPs about the attack activity

[12] See the PDF at http://www.dnssec.comcast.net/DNSSEC_Validation_Failure_NASAGOV_20120118_FINAL.pdf.
[13] https://community.icann.org/display/whoisreview/Draft+Report
[14] Randal Vaughan and Gadi Evron's 2006 description of this type of attack can be downloaded from http://isotf.org/news/DNS-Amplification-Attacks.pdf, while Matsuzaki Shinobu's shorter PDF may be downloaded at http://meetings.ripe.net/ripe-52/presentations/ripe52-plenary-dnsamp.pdf.

and consequently, valid authoritative DNS servers could be blocked from the Internet. This could result in other, indirect denial-of-service (DoS) implications.

DNS amplification attacks are in common use today and predominantly use a small set of known DNSSEC-enabled domains in their attack activity. It is not clear whether these attackers know that DNSSEC contributes to the amplification effect; however, when attackers do discover this, the volume and diversity of their exploits may worsen even as other attackers adopt the same tactics. Some analysis of existing DNS amplification attack activity conducted by one ISP is included in the Appendix of this report. (As of the date of this report, these findings have intentionally been left unpublished.)

5.1.2.7 Unanticipated Abuses of DNSSEC-Enabled Services

It is not clear whether an objective evaluation has been performed of potential abuses of a DNSSEC-enabled infrastructure. Theoretical examples include:

a) The same techniques that allow a cache-poisoning attack to occur could be used to generate a DoS attack against domains. If an attacker can execute a cache poisoning, then validating resolvers and/or validating end users may reject the resulting response. The impact of this may be to block access to the domain. While subsequent responses are unclear, this problem suggests that some sort of detection and prevention may be necessary despite DNSSEC deployment. In this scenario, it may be better to implement detection and prevention as the first-order solution rather than DNSSEC.
b) Malware may use DNSSEC validation to mislead or prevent some security measures. For instance, some attackers use DNS to direct malware toward their command-and-control or exfiltration drop servers, sometimes dynamically. One countermeasure involves “sinkholing” the domain names used by the malware, which allows victims to identify victim devices (a potentially complex process beyond the scope of this report). However, it is sometimes necessary to use DNS redirect (or a "sinkhole") if the DNS providers for the malicious domains are uncooperative. If the domain names are redirected, they will fail any DNSSEC signing and validation; consequently, attackers would be able to detect that they are being sinkholed, disable their malware, and hide their tracks, meaning that infected machines might never be identified.

Potential abuses such as these should be considered further.

5.1.2.8 Lack of Direct Financial Benefit

Like many security-related protections, DNSSEC is largely prophylactic in nature. As such, there does not appear to be any way for an ISP to charge directly for performing DNSSEC validation and the ISP's costs may in fact rise somewhat as they begin to implement DNSSEC. However, ISPs may be able to avoid some costs outlined in this section, and adoption has other indirect benefits. Ultimately, though, DNSSEC is being implemented globally and will become a normal and expected part of any Internet service.

5.1.2.9 More Signed Domains Needed

The value to ISPs of validating lookups increases as the number of signed domains increases, so it is important that when signing is considered, both large numbers of "ordinary" domains are signed as well as the smaller number of "high impact" domains.

One prominent example of the latter category is PayPal, which in December 2011 said it had signed all the zones in its top-level domains. This is a high-impact event considering that the company has over 100 million active accounts; PayPal customers can now validate PayPal's DNSSEC signatures and be confident they are dealing with the company and not an impostor or hijacker.

Meanwhile there are European examples in which registrars have either worked closely with registries to sign domains *en masse*, or where registries charge less to create a signed zone than an unsigned one, creating a clear economic incentive for registrars to offer DNSSEC signing as an option or even a default. The Czech Republic registry, nic.cz, and associated registrars have taken both these paths and the result has been a startlingly high 35 percent of .cz domains signed as of January 2012.[15]

This is a potential path for U.S. ISPs that, to the extent that they are also registrars, may want to negotiate similar agreements with registries, thus creating the same type of incentive. It is also a way for ISPs that have made the investment in DNSSEC infrastructure to leverage that investment economically.

5.1.2.10 End-System Validation Obviates the Need for ISP Validation

Some ISPs have questioned the need to perform DNSSEC validation if the end goal of DNSSEC deployment is to have an end-to-end chain of trust that includes the last mile from the ISP to the end user. This line of thought concludes that ISPs may simply act as conduits for DNSSEC signature information that end users will then validate on their own, removing the necessity for ISPs to deploy DNSSEC themselves.

However, as the discussion above in Section 5.1.1.5, "ISP Adoption of DNSSEC Is Vital for End-Point Validation," indicates, even if only end users performed validation, this system would still require that ISPs retrieve and return the DNSSEC records the end system needs to carry out its own validation. If the ISP is not at least up to category B as shown in Table 4 above, the end system would have to do all of the queries itself, thereby bypassing the DNS service provided by the ISP. To put this more strongly, if the ISP is not providing at least category B service, the service it provides will not be usable by any end system that needs a validated response.

[15] Personal communication of Steve Crocker with Jaromir Talomir, technical manager for .cz; Talomir reported that 314,088 of 894,033 domains were signed as of January 31, 2012.

6 ISP Considerations and Specific Best Practices for Deploying DNSSEC

When determining whether and what kind of DNSSEC service to deploy, ISPs should consider their customer base and whether customers are likely to be attracted to a DNSSEC-aware but non-validating service, which would allow customers to do their own validation (typically large enterprises), or to a fully validating service (typically consumer end users happy to leave end-point validation to the ISP).

It is likely that ISPs will want to provide at least the capacity to perform validation on their customers' behalf as a service differentiator, and they may want to give consideration to securing the last mile of the transaction to provide complete end-to-end security.

Other considerations are included in the following sections.

6.1.1 Evaluate and Test Network Equipment, Servers and Software

As a first step, an ISP should evaluate its network equipment, servers, software, and capacity. This includes design evaluation, lab testing, and capacity forecasting. ISPs might want to ask the following types of questions:

- Will firewalls and load balancers permit both UDP and TCP traffic on port 53?
- Can firewalls and load balancers handle larger DNS responses? For example, UDP responses are usually up to 512 bytes long, but if EDNS0 is used those responses may be up to 4,000 bytes.
- Can routers, firewalls and load balancers properly handle any potential fragmentation caused by larger queries and/or responses, including traffic for recursion?
- Will the current version of the ISP's DNS server software support DNSSEC or is an upgrade necessary?
- For DNS recursive resolvers, when validation is turned on, is the load on the server or its capacity-handling level affected? For example, if turning on validation increases CPU utilization by some amount, this could reduce the peak queries/second capacity of each server, necessitating the deployment of additional servers to properly service DNS queries.
- As more domains are signed, how does this affect the average user's average query size and query frequency? Does that in turn affect the ISP's capacity planning model, potentially requiring greater investment in DNS server capacity?

6.1.2 Evaluate Domain-Signing Tools, Processes and Authoritative Server Capacity

As noted earlier, ISPs also sign their own domains, so they will need to evaluate the ability of their domain registrar(s) to accept DS records (i.e., whether there is an automated process or standard interface for doing so or whether this will involve manual processing by ISP and/or registrar personnel).

Also, the ISP may have a significant number of zones to sign and keys to roll over at specific

times, so automated software tools are critical to ensuring a successful DNSSEC deployment. While some authoritative DNS software will be able to perform signing activities (though a version upgrade may be necessary), other software may not be able to do so, or may not be able to do so efficiently or effectively. This would call for an additional functional component in the authoritative-server architecture. A variety of open-source and commercial signing services will be important to addressing these needs.

Should an ISP's authoritative DNS software be capable of performing automated signing, it's important to note that this activity can be resource-intensive, which may necessitate adding authoritative DNS server capacity. Finally, the processes used to sign domains, manage chaining to subdomains, and manage the insertion of and updates to DS records can be complicated. ISPs should take the time to plan this process out carefully, especially for large, complex, and/or mission-critical domains.

6.1.3 Operational Monitoring

ISPs should consider subscribing to relevant industry email lists that discuss DNSSEC-related issues, such as the dns-operations@lists.dns-oarc.net list maintained by the DNS Operations Analysis and Research Center. This is a good way to be aware of DNS-related issues as well as report on or seek assistance with DNSSEC problems.

ISPs already track key performance indicators (KPIs) or operational metrics related to their DNS servers, and they may find it helpful to track the rate of DNSSEC validation failures. Above certain failure-rate thresholds, alarms could be triggered that alert ISP engineers to potential problems requiring investigation. Depending on the issue, an ISP may proactively contact the domain owner to help them identify and correct the problem, as some ISPs have already begun doing in the past few years. In some cases, an ISP may even implement a negative trust anchor (see Section 6.1.4 below) to temporarily bypass DNSSEC validation for a misconfigured domain.

In addition, as an ISP migrates to DNSSEC, it should carefully monitor KPIs related to server CPU utilization. Any changes should be checked against capacity-forecast models to ensure that the ISP will not exhaust peak-hour query response capacity or response times as its migration progresses. Should capacity-forecasting model changes hint at increasing future demand, an ISP may need to pause further migration work until capacity can be added, or rush new capacity into production to meet projected system demands.

6.1.3.1 Signature Expiration Alerts

One practice that would help smooth the process for all parties involved in DNSSEC adoption would be mechanisms for alerting administrators to the upcoming expiration of DNSSEC signatures. These could take the form of both software-based alerts and the publication of signature expiration dates, which would enable those concerned to highlight to zone operators the need to act promptly.

6.1.4 Negative Trust Anchors

An ISP cannot correct misconfigured records of domains for which it is not authoritative (i.e., when it is not the domain owner). In the case of DNSSEC, though, some ISPs have the ability to use a so-called negative trust anchor on a case-by-case basis to address DNSSEC

misconfigurations.

When a domain has been confirmed to be failing DNSSEC validation due to a DNSSEC-related misconfiguration, an ISP may in some cases consider using a negative trust anchor for a domain (if their DNS software supports this feature). This instructs an ISP's DNS recursive resolvers to temporarily not perform DNSSEC validation for a given misconfigured domain. This immediately restores access to the domain for the ISP's customers while the domain's administrator corrects any misconfiguration(s).

The long-term, frequent use of such a tool is not scalable, but it is useful in the near term as DNSSEC-adopting organizations mature their operational practices. Thus, as DNSSEC evolves into 'just another standard' for DNS configuration, domain owners will ultimately become fully competent at ensuring that all their DNS records are correctly configured.

6.1.5 Education of Customers and Customer-Service Personnel

ISPs should evaluate, update and augment frequently asked questions (FAQs) relating to DNS and DNSSEC as needed. A customer FAQ related to why a customer should not switch to an alternative (non-validating) resolver if DNSSEC validation fails may be especially worthwhile, and might read something like: "We recommend strongly against changing your DNS servers to ones that do not perform DNSSEC validation in order to attempt to access the domain name which has failed that validation. Such a failure may indicate a security problem that could result in your computer being infected with malware, your login credentials for a site being compromised, and other security problems."

6.1.6 Recommended Diagnostic and Other Tools

An ISP should expect to periodically investigate and diagnose DNSSEC validation failures. When they do so, two online tools are particularly useful and straightforward:

- DNSViz, created and maintained by Sandia National Laboratories, at http://dnsviz.net
- DNSSEC Debugger, created and maintained by Verisign Labs, at http://dnssec-debugger.verisignlabs.com

If an ISP desires to see and possibly employ existing DNSSEC tools can see a survey of available tools and resources at: https://www.dnssec-deployment.org/wiki/index.php/Tools_and_Resources. Additionally, ISPs or other organizations that want to make use of a wide range of freely available tools may consult the DNSSEC Tools project website at: https://www.dnssec-tools.org/.

6.1.7 Initial Beta Testing

ISPs beginning the first phase of production network testing with validating resolvers may wish to consider beta-testing this service on an opt-in basis. Thus, end users may manually reconfigure their DNS settings to point to the beta-testing servers. This can enable an ISP to validate its capacity-model assumptions and observe real-world traffic on a controlled basis.

6.1.8 Gradual Deployment

Like any significant new functionality, ISPs are well-advised to gradually enable DNSSEC validation in their networks. Simultaneously turning validation on for all users and on all servers would likely pose a significant operational risk; for example, DNS servers or other network elements such as load balancers might suddenly become overwhelmed.

A better approach is to manage operational risk more prudently by undertaking a gradual, incremental deployment over some extended period that allows operational and support personnel to respond to any potential problems.

6.2 Finding

The Working Group generated the following finding:

- ISP support for DNSSEC is necessary even in a future in which end points perform all validation. They must be able to, at a minimum, recognize DNSSEC-related traffic and allow it to pass for the smooth functioning of an end-to-end, DNSSEC-secured system.

6.3 Recommendations

The Working Group recommends that:

1. ISPs implement their DNS recursive nameservers so that they are at a minimum DNSSEC-aware, as soon as possible.
2. Key industry segments, such as banking, credit cards, e-commerce, healthcare and other businesses, sign their respective domain names. The FCC ask industry-leading companies in key sectors commit to doing so, in order to create competitive pressure for others to follow. These industries may be prioritized based on the prevalence of threats to each one, which would mean focusing on financially related sites first, followed by other sites that hold private user data.
3. Software developers such as web-browser developers study how and when to incorporate DNSSEC validation functions into their software. For example, a browser developer might create a visual indicator for whether or not DNSSEC is in use, or perhaps only a visual warning if DNSSEC validation fails.

7 Appendix: Samples of Data Types

7.1 Analysis of Recent DNS Amplification Attack Activity

As an example of ISP efforts to monitor and mitigate DNS amplification attacks, one U.S. ISP has detectors in place to identify such activity. It uses a fairly conservative threshold that limits the number of false positives, although there are a small number of false positives in the data (<3%). Today, these attacks are relatively easy to detect since there is limited diversity in the types of queries used and little public recognition that DNSSEC could have an amplifying effect.

DNS amplification attacks occur frequently; over a recent 90-day period (11/14/11–2/14/12), this ISP recorded approximately 6,400 alarms. While some of those attacks may have targeted this ISP or its customers, most simply transited the ISP's network en route to another destination.

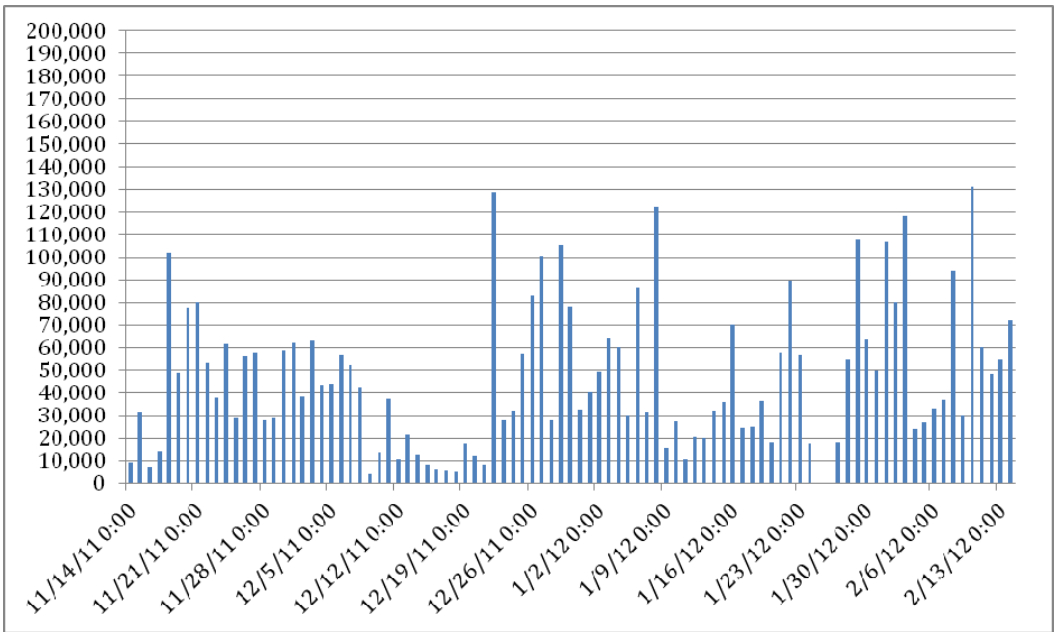

Figure 5 – Estimated DNS query response rates /second or packets/second (Source: a U.S. ISP)

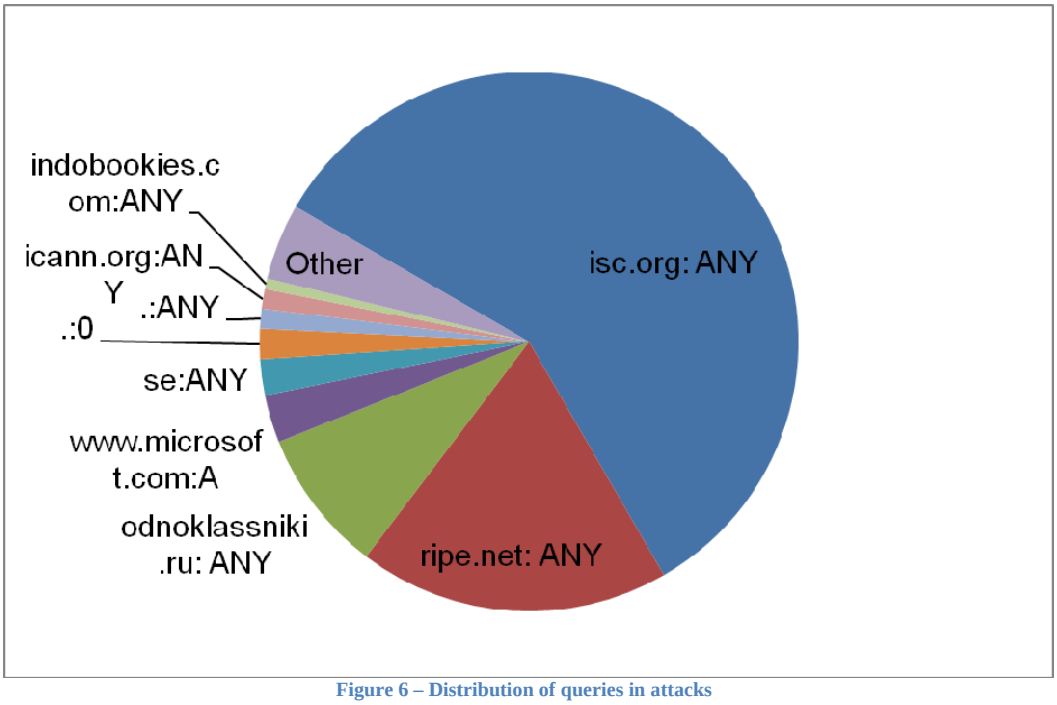

Figure 6 – Distribution of queries in attacks

Of the top nine queries detected in alarms, all but three are DNSSEC-enabled. The one query that is not DNSSEC-enabled is www.microsoft.com, which appears to a false-positive detection and not related to attacks.

The following visualizes a recent DNS amplification attack against a U.S. ISP's customer:

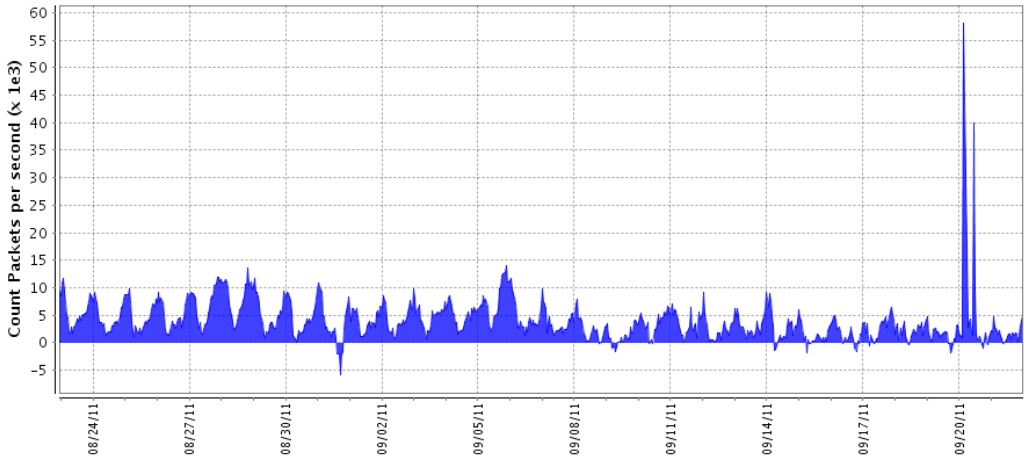

Figure 7 – A recent DNS amplification attack against a U.S. ISP's customer (Source: a U.S. ISP)