SECURE64™

WHITE PAPER

Have I Reached The Party To Whom I Am Speaking?

Protect your business and your brand with DNSSEC

Introduction

"One ringy-dingy… two ringy-dingy… A gracious good morning to you. Have I reached the party to whom I am speaking?" Lily Tomlin often began her Ernestine the telephone operator comedy routine with these lines. Ernestine did not know (or care) whether she was talking to the right party, remarking that "the phone system consists of a multibillion-dollar matrix of space-age technology that is so sophisticated, even we can't handle it. But that's your problem, isn't it?"

When it comes to the internet, making sure you reach the right party can be a problem. Much like a telephone operator, the DNS works behind the scenes to translate names into IP addresses for services on the Internet. And like Ernestine's telephone company, today's DNS infrastructure cannot guarantee the answers the DNS system provides. Attackers can insert spoofed information into DNS responses, reroute requests to bogus name servers, and redirect DNS resolvers and email clients to servers under their control, leaving your organization vulnerable to a wide variety of fraudulent activities.

> Attacks on the internet infrastructure are a reality—it's estimated that 10 percent of servers in the network today are vulnerable to domain name system (DNS) attacks. And many technology experts believe that we will see a serious attack on the underlying infrastructure within the next decade. [1]

When attackers alter DNS information, individuals and organizations are exposed to a large set of risks:

- **Identity theft**. Attackers can manipulate the DNS to obtain sensitive data such as login information, passwords, account numbers, and identification numbers.
- **Damage to brand**. Organizations affected by DNS attacks can suffer great damage to their brand reputation if the attack is publicized.
- **Loss of sensitive information**. Attackers can re-route or block internet-based applications such as email and VoIP to disrupt operations or obtain sensitive information.
- **Malware distribution**. Attackers can redirect visitors to bogus web sites that secretly download Trojan horses, viruses, or other malware. The attacker can use the malware to send spam or launch other attacks.
- **Financial market manipulation**. Attackers can use the DNS to redirect internet users to bogus sites that publish information such as fake news stories and stock tickers to influence financial markets.

To address these problems and secure the critical name-to-address mapping function of DNS, a specification called DNSSEC (DNS Security Extensions) was developed. It is an extension to the DNS that can:

- Validate that a DNS query or response was sent by the source claiming to send it (authentication)
- Determine whether a DNS response has been tampered with (data integrity)
- Verify that a DNS record does not actually exist when a response is returned as unresolvable

[1] www.dnssec-deployment.org

In this whitepaper, we discuss the security vulnerabilities in the current DNS infrastructure. We then examine DNSSEC and how it can benefit organizations and the services that rely on the DNS.

DNS Vulnerabilities

The DNS is a distributed database of information managed by hundreds of thousands of name servers across the internet. Authoritative name servers hold the actual data that provides the answers to DNS queries. Resolving or caching name servers perform the task of locating authoritative name servers on behalf of DNS clients, such as web browsers and email applications.

In addition to finding and providing DNS answers, resolving or caching name servers store the answers for a period of time. This reduces the load on the DNS infrastructure and provides a quicker response to subsequent queries for the same information.

Figure 1 illustrates the normal query resolution process for a web site address. The steps include:

1. A web browser queries its caching server for the IP address of www.example.se.
2. The caching server does not have the answer in its cache. It does have the IP address for the .se authoritative server in its cache, so it queries the .se authoritative server for www.example.se and receives a referral to the authoritative server for the domain example.se in response.
3. The caching server queries the authoritative server for example.se and receives the IP address of www.example.se (1.2.3.4) in response.
4. The caching server stores the answer in its cache and returns the answer to the web browser on the client computer.
5. The web browser uses the IP address to connect to the web server.

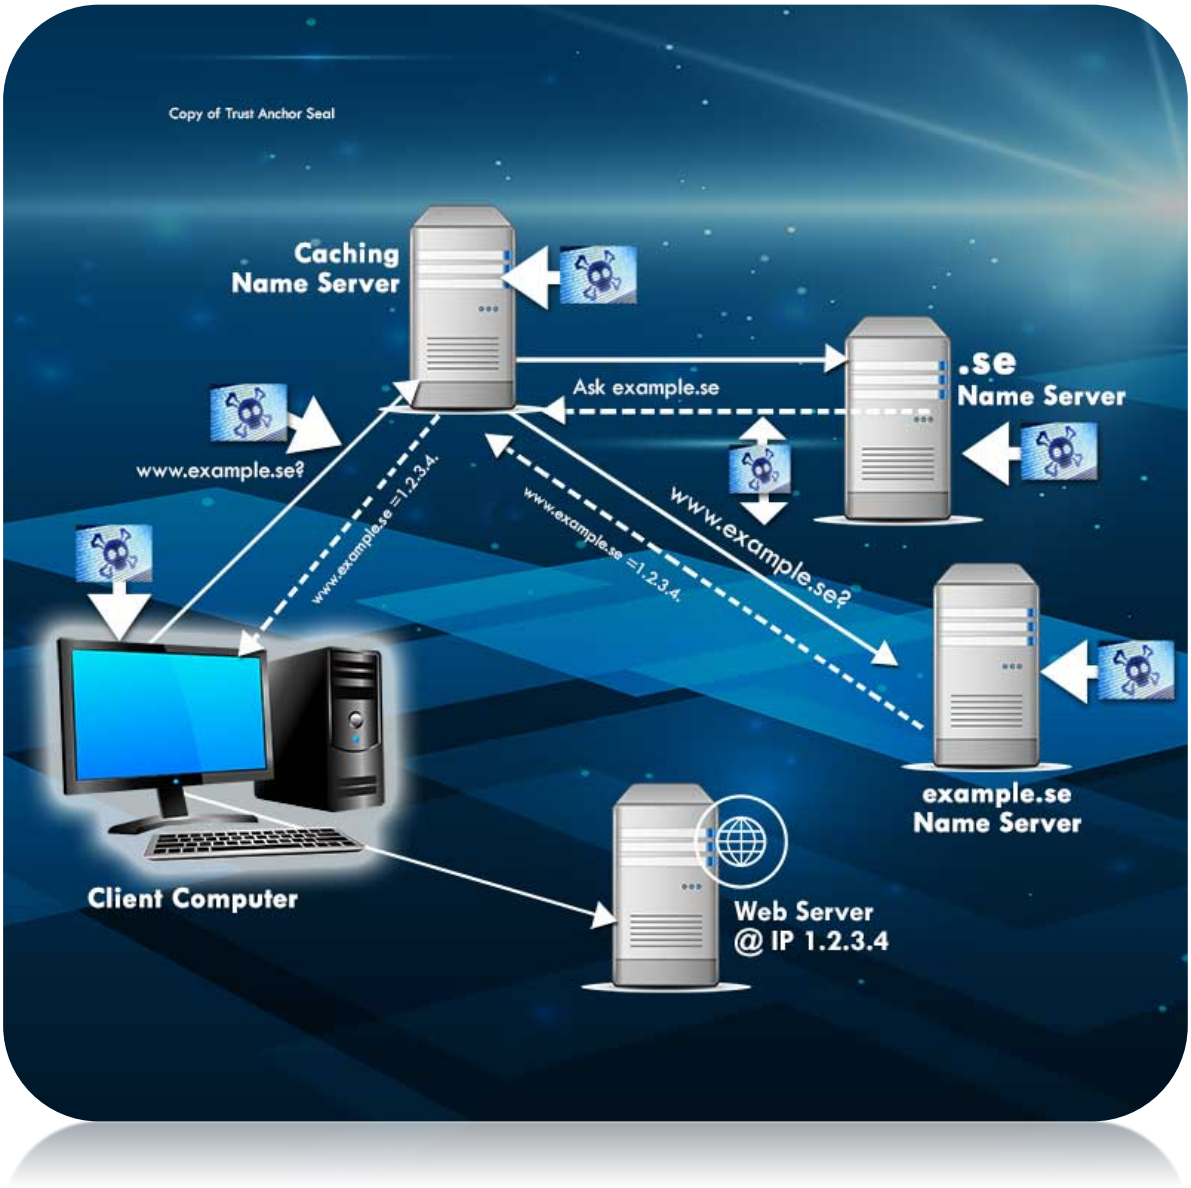

Figure 1. DNS query resolution process and potential attack targets

As Figure 1 also points out, this process is vulnerable to a number of different types of attacks that can cause incorrect answers to be returned to the client computer, leading users to visit fake web sites or email servers without their knowledge. Attacks can focus on any one of four different targets: the authoritative server, the caching server, the client computer, or the network itself. Examples of some of the most common of these attacks are provided below.

Attacks on Authoritative Servers

AUTHORITATIVE NAME SERVER COMPROMISE

Because the authoritative name server stores all of the DNS information for domains for which it is authoritative, any attack that compromises the authoritative server itself can potentially alter the DNS

data at its source. For this reason, best practices for managing DNS servers call for hardening and patching authoritative servers and using a hidden master in order to minimize the possibility of a compromise.

ZONE TRANSFER ATTACKS

Even if an attacker cannot compromise the authoritative server itself, it may be able to alter DNS data by fooling the authoritative server into accepting a zone transfer from an attacker controlled server. If successful, this attack can alter the DNS data and redirect users to unauthorized servers.

DOMAIN HIJACKING

This type of attack often involves the use of social engineering to alter the registration information for a given domain. Once having administrative control over a domain, an attacker can change the DNS information.

Attacks on Caching Servers

CACHE POISONING

Cache poisoning attacks use various means to insert incorrect information into the cache of DNS caching or resolving servers. When it receives a query from a client, the compromised server provides a response that contains the incorrect information, without the end-user's knowledge.

DNS FORGERY

A DNS forgery attack is a specific type of cache poisoning attack in which the objective of the attacker is to redirect a user to a rogue web site for the purpose of stealing personal or financial information. Unlike phishing attacks, where the URL of the web site is visibly changed, a DNS forgery attack is nearly impossible for a user to detect.

Attacks on Client Computers

CORRUPTED HOSTS FILE

If an attacker is able to get malware running on the client computer, the malware can alter the hosts file, which contains a list of hostnames and corresponding IP addresses that the client computer checks before querying the caching server. If this information is altered, the attacker can direct the client to bogus servers.

CORRUPTED RESOLUTION PATHS

This type of attack alters the IP address of the caching server as stored on the client computer (a registry entry on Microsoft Windows machines) to point to a caching server under the attacker's control. The attacker's caching server is then able to direct the client computer to bogus web sites or divert email to an attacker-controlled email server.

Attacks through the Network

DNS PACKET INTERCEPTION

Attackers can eavesdrop on DNS queries and intercept DNS packets. Then, by responding to the query faster than the intended caching server, an attacker can send a forged response to the client computer.

In addition, attackers can intercept DNS client requests and respond with bogus information by spoofing the authoritative name server that holds the answer to the query.

DNS MAN-IN-THE-MIDDLE

In one type of DNS packet interception attack, the attacker sniffs DNS query packets and sends a fake IP address (one the attacker controls) in response. As a result, all of the traffic from the client to the destination and vice versa travels through the attacker's system, which is known as a man-in-the-middle attack. By secretly harvesting information such as usernames, passwords, or the contents of email messages, the attacker can obtain unauthorized access to financial accounts, discover sensitive government or business data and use it for financial or political gain, or sell the information to identity thieves. Attackers can also use DNS cache poisoning to launch a man-in-the middle attack.

What is DNSSEC?

DNSSEC is a set of IETF (Internet Engineering Task Force) standards outlined in RFCs 4033, 4034, and 4035. Its purpose is to authenticate responses to DNS queries through the use of a trusted chain of name servers. It does this in part by utilizing public-key cryptography to digitally sign DNS data.

DNSSEC protects DNS clients (such as web browsers and mail clients) from forged DNS data. If an attacker attempts to alter any part of the DNS resolution process, then a DNSSEC-aware client can detect the altered response. Note that this protection does not prevent an attacker from injecting false data into DNS communications; instead, it allows the client to detect with certainty when this has happened. [2]

A Transparent, Sealed Envelope

As described by Kolkman et al.[3] and illustrated in Figure 2, DNSSEC resembles a sealed, transparent envelope:

- The seal is applied by whoever closes the envelope. Because the seal is publicly known, anyone who receives the message knows where it came from (source authentication).
- Anyone can read the message. Because the envelope is transparent, all recipients can read the message. But no one can change the message without breaking the seal (data integrity).
- The seal is applied to the envelope, not the message itself. The message is not encrypted. (The DNS information is not private or confidential.)

[2] "DNSSEC-The Theory," Geoff Huston, http://ispcolumn.isoc.org/2006-08/dnssec.html

[3] "DNS Risks, DNSSEC," Olaf M. Kolkman and Allison Mankin, http://www.nlnetlabs.nl/downloads/DNSandDNSSEC.pdf

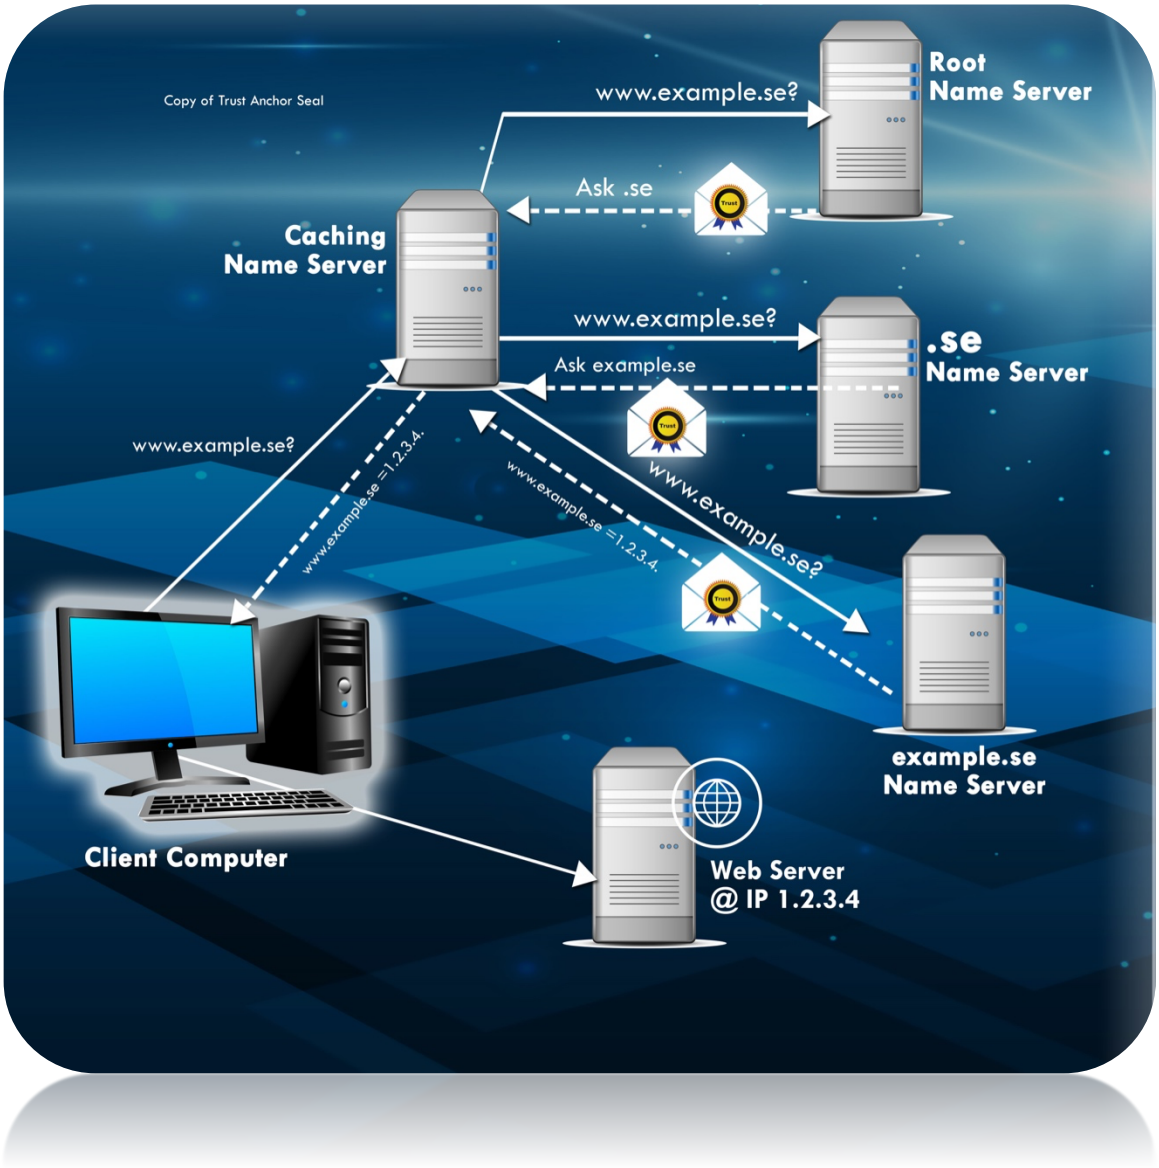

Figure 2. Securing DNS query resolution with DNSSEC

Using this seal and sealed envelope metaphor for public key cryptography, Figure 2 shows how DNSSEC adds the following important elements to the query resolution process:

- **Trust anchors**. In order to validate DNS responses, clients must be configured with one or more trust anchors. A trust anchor is a seal for a known, trusted domain that is obtained through a secure process. In Figure 2, a copy of the seal of the .se authoritative server stored on the client computer is the trust anchor.
- **Sealed responses**. Each authoritative server in the query resolution process signs its responses with its own seal. If the seal is missing, broken, or spoofed, the recipient knows that it cannot trust the source of the response. In Figure 2, responses from .se and example.se are both sealed.

- **Chain of trust**. Each DNSSEC-enabled parent domain has a copy of the seal of any DNSSEC-enabled children domains, which it obtains through a secure process. A DNSSEC-enabled child has copies of its DNSSEC-enabled children's seals, and so on. This creates a chain of trust within a set of chained DNSSEC domains. In Figure 2, the .se domain has a copy of the seal of the example.se domain, which creates a chain of trust from .se to example.se.

VALIDATING THE CHAIN OF TRUST

In order to validate the response from an authoritative server, the client must validate the seals of each of the authoritative servers involved in the query resolution hierarchy, starting with the trust anchor and ending with the server that provided the authoritative response.

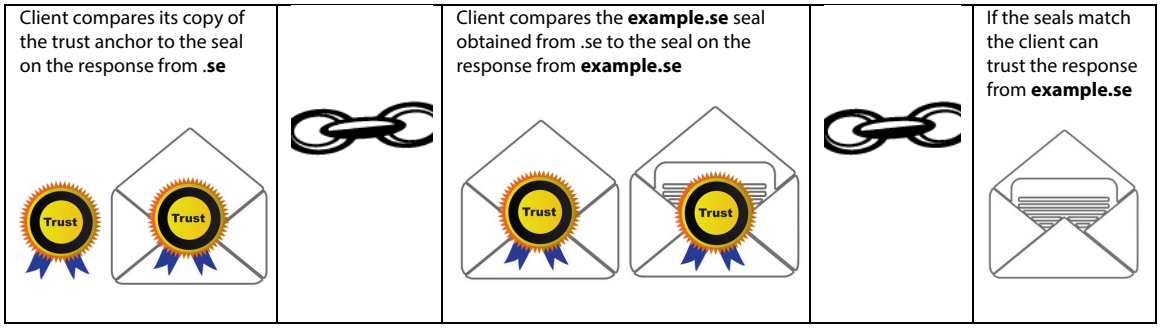

As noted previously, each authoritative server vouches for the name server it refers by including a copy of the seal of the referred authoritative server in the response.
This chain of trust is illustrated in Figure 3:

- When the client receives the sealed response from .se (the trust anchor), it compares its own copy of the .se seal with the seal on the response from .se. If the seals match, the client knows it can trust that the response came from .se and, because the seal is not broken, it knows that the response has not been altered in transit.
- The trusted response from .se includes a copy of the seal of the next authoritative server in line (example.se).
- The client compares the seal it retrieved from .se to the seal on the response from the example.se authoritative server. If they match, then the client knows it can trust the answer (www.example.se=1.2.3.4) from example.se.

Note that DNSSEC does not prevent an attacker from impersonating a DNS server or tampering with DNS information. However, if the seal is missing, broken, or from a source different from the source of the message, the client knows it cannot trust the authenticity or integrity of the message.

In summary, DNSSEC provides:

- Authentication of the source of the response to a DNS query
- Verification of the DNS data integrity
- Validation of non-existing DNS data

SOURCE AUTHENTICATION

DNSSEC utilizes public-key cryptography and digital signatures to validate the authenticity of a query response. Starting with a trusted DNS server (.se in our example), DNSSEC-aware clients or servers verify the digital signature of pubic keys (the seals) to establish an unbroken chain of trust to the source of the response. This chain of trust allows the recipient system to know for sure that the response came from an authorized DNS server and not from an attacker.

DNS DATA INTEGRITY

Digitally signing DNS data also provides proof that the data has not been altered or tampered with. It confirms that the data sent by the queried DNS server is the same data received by the recipient of the response (the seal is not missing, broken, or spoofed).

NON-EXISTING DNS DATA

In addition to validating DNS data in a response, DNSSEC provides a mechanism for validating non-existent DNS data. A special response is returned to authenticate that the requested data does not exist. This proof of non-existence can prevent attacks that remove valid DNS data, resulting in a query failure or denial of service.

Benefits of DNSSEC

DNSSEC adds certainty and trust to internet communications, protecting consumers and organizations from a wide variety of fraudulent activities. Benefits from deploying DNSSEC include:

- Increasing consumer confidence in online transactions
- Protecting brand reputation from damage
- Protecting valuable intellectual property from loss
- Reducing the likelihood of costly litigation
- Adhering to external security and audit requirements
- Increasing the security of Internet-based supply chain transactions
- Positioning the organization as a trusted entity

DNSSEC deployment is critical for financial institutions, government agencies, and security-conscious enterprises doing business on the internet or securing internal networks. Regardless of the type and amount of zone data or query load, these benefits are valuable at all levels of the DNS hierarchy from top-level domain (TLD) operators with a few large zones to individual ISP's that want to deploy DNSSEC for hundreds or thousands of customer zones to organizations wanting to secure the answers provided by their authoritative servers. Additionally, implementing DNSSEC is the first step in deploying DANE (Domain Authentication of Named Entities) which enables secure transit of email.

DNSSEC Deployments Today

DNSSEC has been deployed on the root servers, the .arpa servers and virtually all of the major generic top level domains like .com, .org, and .net. Most of the largest country code top level domains like .de, .uk, .nl and .eu have deployed DNSSEC.

DNSSEC is required by FISMA (Federal Information Security Management Act), which was established in January 2003 to produce several key security standards and guidelines required by Congressional legislation. FISMA requires incremental deployment of DNSSEC across U.S. governmental agencies and their IT

Conclusion

DNS is behind the billions of requests made on the internet every day. But today's DNS infrastructure cannot authenticate the answers the DNS system provides, leaving organizations and their customers, employees, and business partners vulnerable to a wide variety of fraudulent activities.

In practice, DNSSEC implementation is occurring in pieces, with some areas of the query resolution process such as the last mile (responses from caching servers to client computers) being less developed than other areas, such as the responses from authoritative servers to caching servers. And DNSSEC is not a panacea that solves every DNS-related security issue. Indeed, administrators must employ other security measures to prevent other types of attacks such as denial-of-service or malware injection.

Nevertheless, implementing DNSSEC eliminates risks and reduces DNS attack vectors compared to continuing on the same unsecured DNS response path. Only through the concerted efforts of the internet community, the demands of businesses and the government, and the solutions developed by commercial providers can the protections of DNSSEC be truly realized.

About Secure64

Secure64 brings trust to the internet through its suite of purpose-built, secure, DNS-based network security products. The company was built on a foundation of security and has forged solutions that are self-protecting and immune to malware. Secure64 secures the DNS infrastructures of leading service providers, government agencies and enterprises globally,

Secure64 is a privately held company founded by HP veterans and boasts deep technical and global experience in its leadership and staff. It is the only DNS solution provider that has authored a secure micro OS, automated the deployment of DNSSEC and built self-protecting DNS servers. For more information, visit us at www.secure64.com

For More Information:
(303) 242-5890
www.secure64.com

Secure64 Software Corporation
5600 South Quebec Street, Suite 320D
Greenwood Village, CO 80111