THALES

Protecting the Internet with Thales nShield HSMs

Solution Brief – Securing signing keys for DNSSEC deployments

The domain name system (DNS) is a critical network infrastructure component responsible for the routing of both intranet and Internet connections. Because it was never designed with security in mind, inherent vulnerabilities in the system pose potential risks to all forms of Internet communications. In practice the DNS is made up of thousands of distributed servers that communicate and share IP addresses and domain name information with each other in the form of DNS queries. The DNS is effectively the Internet's master address book, enabling web addresses such as the Thales e-Security domain name (www.thales-esecurity.com) to be translated and matched to its corresponding registered IP address (98.129.76.138). Any illicit alteration to queries can potentially route users or services to rogue IP addresses with unfriendly servers impersonating legitimate sites.

This vulnerability in DNS has been known since the late 1990s. With the growing reliance on the Internet for all manner of services such as e-mail, banking, web services, voice over IP (VoIP), cloud services and more, the security of DNS is of increasing concern in order to manage the risk of severe outages and possible compromises to enterprise and government networks. In order to manage these risks, domains are now starting to deploy DNS security extension (DNSSEC) – an addition to the DNS standard that is designed to address these vulnerabilities and mitigate the risk of compromise.

Thales hardware security modules (HSMs) enable top level domains (TLDs), registrars, registries and enterprises to secure critically important signing keys used to validate the integrity of DNSSEC responses across the Internet, and protect the DNS from what are commonly referred to as "cache poisoning" and "man-in-the-middle" attacks. This solution brief highlights the growing concern over the security of the DNS for both internal organizational intranets where the integrity of local DNS records is critical, as well as for external Internet-based transactions where trusted communications are vital for continued growth in electronic commerce.

What is DNSSEC and how does it protect against these vulnerabilities?

As an addition to the DNS standard, DNSSEC mitigates the threat of cache-poisoning and man-in-the-middle attacks by establishing a mechanism to authenticate and verify the integrity of DNS responses to DNS queries. Cache poisoning is the accidental or deliberate introduction of incorrect records into the cache of a DNS server, causing incorrect routing information to be provided to users. Man-in-the-middle attacks, while not explicitly altering the records in the DNS database, intercept user requests and pose as a legitimate DNS server.

As shown in Figure 1, when a user requests a web page or other resource over the Internet using their browser in a non-DNSSEC enabled environment (step 1), the corresponding IP address sourced from the DNS records database or cache (steps 2 and 3) and provided by the DNS server (step 4) can be corrupted. As a result of receiving an incorrect DNS response, the user is directed to the impersonating server (step 5). As a result of the attack on the address cache, the user or application perceives that they are communicating with a legitimate server when in fact they are not.

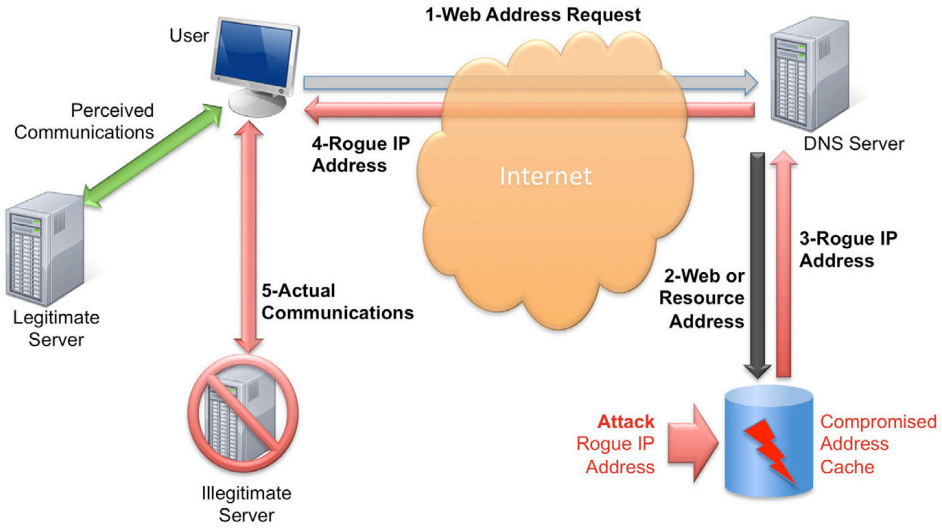

Figure 1: Non-DNSSEC scenario: User is incorrectly routed to an illegitimate server.

DNSSEC uses proven public-key cryptography, which has been widely employed and trusted to protect many other network security applications, to digitally sign DNS resource records. In this way it enables DNS servers to attest to the origin and integrity of records tying domain names to their corresponding IP addresses. By using DNSSEC, organizations are able to minimize the risk that a service or user is redirected to an incorrect IP address masquerading as a legitimate site where they can fall victim to other compromises.

Just as secure socket layer (SSL) has become the de facto standard for encrypting sensitive data to protect privacy over the Internet, DNSSEC is expected to become the default mechanism for protecting the integrity of routing instructions. As DNSSEC is deployed, a "chain of trust" is created that spans the multiple layers in the DNS hierarchy, from the highest level root down through the top level domain (TDL) and then further down to the enterprise DNS, and can be extended even further to localized DNS servers within the enterprise. As shown in Figure 2, roots are aware of the addresses of their TLDs. Accordingly TLDs distinguish the addresses of their Internet Service Providers (ISPs), registrars and registries and these in turn identify Enterprise Level Domain (ELD) customers – each of which has the capability to further segment the registries. As with any hierarchical infrastructure, the security of the system is only as good as the weakest link in the chain of trust.

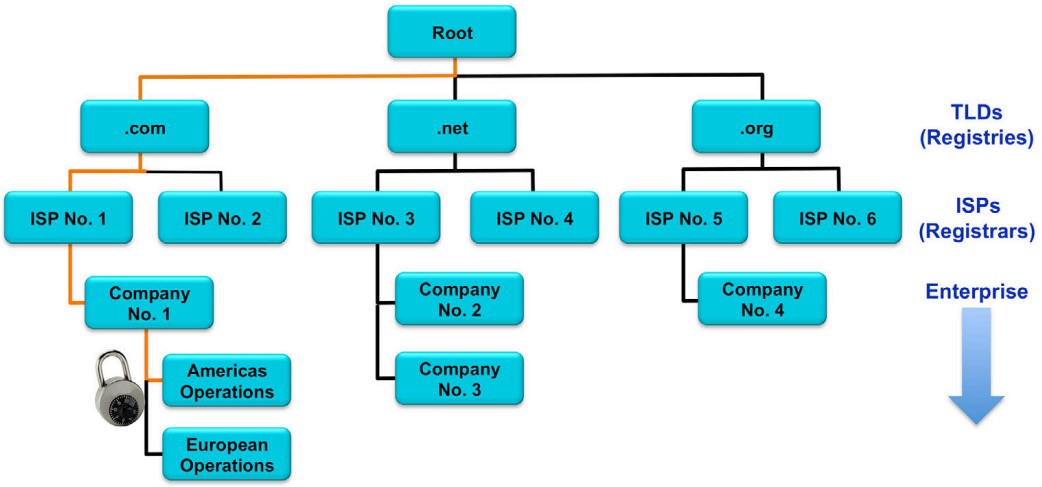

Figure 2: Chain of trust extending from the DNS Root to the TLDs and down to the enterprise DNS.

Who is implementing DNSSEC and what are the challenges for wide-scale deployments?

DNSSEC is being deployed today in multiple domains and at various levels in the DNS hierarchy. The .com, .gov and .org domains are now DNSSEC-enabled and according to the Internet Corporation for Assigned Names and Numbers (ICANN) that manages the IP address space on a global basis, more than 20% of the world's top-level domains pertaining to country codes have now implemented DNSSEC in their zones. Independent research published by Gartner indicates that DNSSEC will see widespread adoption in 2011[1] and the adoption of cloud-based services is expected to further drive DNSSEC-enablement. However, the deployment of DNSSEC is not without its challenges, and in particular the management and protection of the highly sensitive cryptographic keys used to sign DNS records and responses is rapidly emerging as a topic that requires careful planning and robust operational and security practices.

Why are HSMs an important component of DNSSEC integrations?

As DNSSEC begins to be fully deployed across internal and external networks, it important to employ proper generation and storage techniques for signing keys to assure the integrity of the DNSSEC validation process. Compromise of private signing keys can allow rogue sites to successfully impersonate legitimate ones – damaging the confidence in the DNS with possible serious economic impact. Although it is possible to deploy DNSSEC in purely software-based systems, this introduces tangible risk of compromise for DNS signing keys and therefore the signing process. This issue is not unique to DNSSEC and applies equally to a host of other security applications, such as public key infrastructure (PKI), the global payments network and numerous data encryption scenarios. In common with these other instances, HSMs are deployed to mitigate these risks, providing the only proven and auditable way to protect these valuable private keys. HSMs secure these valuable assets within carefully designed cryptographic boundaries that employ robust access control mechanisms with enforced separation of duties to ensure keys are only used by authorized entities. Furthermore, from an operational resilience perspective, they utilize sophisticated key management, storage and redundancy features to guarantee keys are available when needed. Finally, from a performance perspective, as DNSSEC adoption increases, DNS servers will experience increasing signature and verification transaction volumes; HSMs offload the computational load for these transactions from the server central processing unit (CPU) and will play an increasingly important role in DNSSEC performance optimization.

[1] Gartner RAS Core Research, Market Scope for DNS, DHCP and IP Address Management, March 2011.

Americas: 2200 North Commerce Parkway, Suite 200, Weston, Florida 33326 • Tel: +1 888 744 4976 or +1 954 888 6200 • Fax: +1 954 888 6211 • E-mail: sales@thalesesec.com
Asia Pacific: Units 4101, 41/F, 248 Queen's Road East, Wanchai, Hong Kong • Tel: +852 2815 8633 • Fax: +852 2815 8141 • E-mail: asia.sales@thales-esecurity.com
Australia: Level 2, 103-105 Northbourne Avenue, Turner, ACT 2601, Australia • Tel: +61 2 6120 5186 • Fax: +61 2 6120 5101 • E-mail: sales.australasia@thales-esecurity.com
Europe, Middle East, Africa: Meadow View House, Long Crendon, Aylesbury, Buckinghamshire HP18 9EQ • Tel: +44 (0)1844 201800 • Fax: +44 (0)1844 208550 • E-mail: emea.sales@thales-esecurity.com

Why do Thales HSMs offer unparalleled advantages?

Thales is a leading provider of HSMs; thousands of customers around the world have successfully deployed them in a wide variety of PKI-related applications, including DNSSEC. The Thales nShield family of general purpose HSMs has been proven to deliver robust security and high performance in the most demanding operational environments. All nShield HSMs perform key management and cryptographic operations such as encryption and digital signing within a trusted and certified tamper-resistant security boundary on behalf of a variety of common DNSSEC-enabled software systems. Offering unparalleled security over alternative software solutions where an equivalent security boundary cannot be established, deployment of Thales nShield HSMs within a DNSSEC environment provides:

- Certification to FIPS 140-2 Level 3 and Common Criteria EAL4+ providing independent assurance of the security properties of the system[2]
- Robust tamper-resistant hardware protecting key material even when archived
- Strong authentication of administrators and dual controls through the use of advanced quorum techniques to mitigate the threat of single 'super users'
- Advanced separation of duties of key management activities between DNS, IT and security administrators to facilitate regulatory compliance

In addition to enhancing security and compliance, these operational benefits can be realized by ISPs, registrars, registries, and vendors of IP address management (IPAM) appliances and applications. Operational benefits include:

- Centralized key management to support multiple DNS servers
- Scalability to add HSMs dynamically and balance load as capacity requirements increase
- High availability and disaster recovery with unlimited secure key backup and retrieval
- Cryptographic CPU offloading to improve DNS server performance

Thales' unique approach to key management protects keys from loss, provides unlimited storage, supports replication between data centers and ensures continuity of operations. Thales ensures that DNSSEC signing keys always remain protected with FIPS-certified security no matter where they reside in the system. The Thales Security World architecture is unique in the market and enables the automation of burdensome and risk-prone administrative tasks and labor-intensive key backup processes to guarantee secure key recovery, and eliminates the vulnerability of a single point of failure. Thales nShield HSMs provide certified best-of-breed security that is easy to use, easy to implement and ensures cost-effective deployment with minimal operational overhead.

How do nShield HSMs integrate into a DNSSEC deployment?

With DNSSEC enabled, specific DNS servers cryptographically sign their zone records. This permits other name servers in the system to verify the identity and integrity of query responses. The widely deployed Internet Systems Consortium-Berkeley Internet Name Domain (ISC-BIND) is an example of a DNS server application that runs in DNSSEC protected environments and protects against the vulnerabilities outlined earlier in this brief. Thales nShield HSMs have been integrated with the BIND DNS server to support secure private signing key storage. For more details on this this integration and supported functionality please consult the Thales ISC BIND Integration Guide.

As illustrated in Figure 3, zone administrators digitally sign their resource records in the address cache with assigned private keys and publish the digital signatures along with the matching public keys in the DNS. DNSSEC clients validate digital signatures using the zone administrator public key. Successful validation of the digital signature provides confidence that DNS response is authentic and that it accurately provides routing information to the legitimate site. By extending the chain of trust across the domain name hierarchy and deploying HSMs at all levels, DNSSEC enables clients to fully validate the authenticity of DNS responses.

[2] Information on these certification programs is available at http://csrc.nist.gov/groups/STM/cmvp/index.html and http://www.commoncriteriaportal.org/.

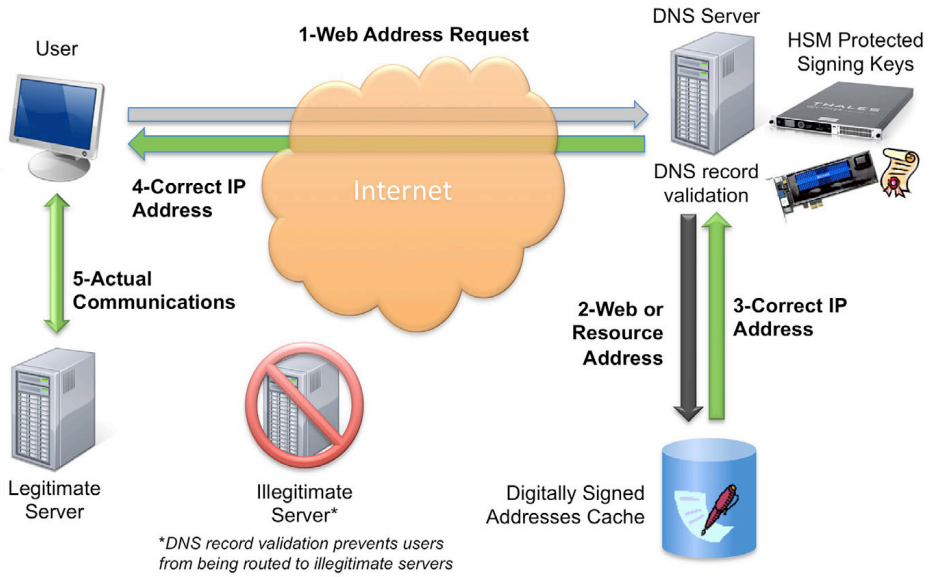

Figure 3: DNSSEC-enabled scenario: Thales HSM(s) securely generate and store private signing keys used to sign resource records.

HSMs can be deployed at all levels in the domain hierarchy to provide enhanced protection for the signing keys used to verify IP addresses. As the most security sensitive component of the DNSSEC, signing keys should always be HSM-protected, enabling users to trust the legitimacy of DNS responses they receive and the corresponding web sites and servers that they reach. DNSSEC thus provides users of critical applications such as online banking with the confidence that they are exchanging their sensitive user IDs and passwords with legitimate sites.

About Thales e-Security

Thales e-Security is a leading global provider of data protection solutions. With a 40-year track record of protecting the most sensitive corporate and government information, Thales encryption and key management solutions are an essential component of any critical IT infrastructure. Thales makes it easy to enhance the security of software-based business applications and reduce the cost and complexity associated with the use of cryptography across the enterprise and out to the cloud.

In addition to the nShield HSM products described in this brief, Thales e-Security offers products for network security, storage security and payments security. To learn more, please visit http://www.thales-esecurity.com

Americas: 2200 North Commerce Parkway, Suite 200, Weston, Florida 33326 • Tel: +1 888 744 4976 or +1 954 888 6200 • Fax: +1 954 888 6211 • E-mail: sales@thalesesec.com
Asia Pacific: Units 4101, 41/F, 248 Queen's Road East, Wanchai, Hong Kong • Tel: +852 2815 8633 • Fax: +852 2815 8141 • E-mail: asia.sales@thales-esecurity.com
Australia: Level 2, 103-105 Northbourne Avenue, Turner, ACT 2601, Australia • Tel: +61 2 6120 5186 • Fax: +61 2 6120 5101 • E-mail: sales.australasia@thales-esecurity.com
Europe, Middle East, Africa: Meadow View House, Long Crendon, Aylesbury, Buckinghamshire HP18 9EQ • Tel: +44 (0)1844 201800 • Fax: +44 (0)1844 208550 • E-mail: emea.sales@thales-esecurity.com