

Securing DNS Environments

White Paper by FusionLayer Inc.



Copyright © 2020 FusionLayer Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of the copyright owners.

Securing DNS Environments by FusionLayer Inc., January 2021.

Any comments relating to the material contained in this document may be submitted to:

FusionLayer Inc. Annankatu 27, 00100 Helsinki, Finland.
or by email to: info@fusionlayer.com



Introduction

The venerable military strategist, General Karl von Clausewitz once suggested that possible engagements should be regarded as real ones because of their consequences [1]. With the emergence of modern warfare strategies that actively utilize cyberspace tactics, von Clausewitz's message is just as current today as it was 200 years ago.

The Internet was originally developed with availability and scalability in mind. At the time, security and safety of the network infrastructure was not as pressing an issue as it is today. As a result, many mission critical network services, such as the Domain Name System (DNS) providing name-to-IP-address resolution service in TCP/IP based networks, have inherent structural vulnerabilities in their protocol, making them easy targets for cyber attacks unless appropriate countermeasures have been taken prior to their deployment.

Whenever a human readable domain name is used to connect to an application over a TCP/IP based network, DNS comes into play. Yet despite its fundamental role, the insecure design of the DNS protocol makes this core network service prone to man-in-the-middle, DNS cache poisoning and other similar forms of malicious attacks. Due to their nature, these attacks can go undetected for significant periods of time while directing traffic to fraudulent sites. This allows the attacker to tamper with the name resolution process, facilitating pharming attempts and blocking of selected services at a critical moment.

To address some of the vulnerabilities in the DNS protocol and mitigate a number of the security issues associated with it, the networking industry has been working on developing, fine-tuning and adopting DNS Security Extension (DNSSEC) since at least 1993. While its deployment has been gaining momentum in the last number of years, in particular with the Root Zone being signed on 15th July 2010, the fact is that the majority of public and private sector DNS environments are still not DNSSEC-enabled.

The reasons for such a lukewarm response to the wide-scale deployment of DNSSEC could be that

1. deploying DNSSEC introduces network complexities and operational challenges that are yet to be clearly addressed and
2. DNSSEC alone does not solve all the DNS security problems and instead should be deployed alongside other measures so as to address a wider spectrum of cyber attacks.

The Internet was designed with scalability and availability in mind. At the time, security and safety of the network infrastructure was not as pressing an issue as it is today



On the other hand, while security measures such as the deployment of DNSSEC, are effective in their own right in improving the trust in the protocol, to truly increase network security and eliminate the possibility of human error, DNS management processes should be automated. In addition, this has to be taken one step further and the automated DNS services should be virtualized; so as to yet increase network security and efficiency while improving disaster recovery tactics.

In this white paper, we focus on our patent-pending DNS software appliance-based solution [2], FusionLayer DNS, which automates DNS management and administration processes while addressing its security issues out of the box. The main advantages of our approach are summarized below:

1. It provides a simple and secure DNSSEC solution that also includes other security measures, which together improve DNS and network security.
2. It automates and streamlines DNS management processes eliminating the possibility of human error, which is one of the main causes of network downtime.
3. The solution, delivered as a virtualization-ready software appliance [3], introduces a new level of platform independence, as it can be installed and run on either a virtual machine or other computing platform already in use.

The majority of public and private sector DNS environments are still not DNSSEC enabled



2. The Basics of the Domain Name System (DNS)

The Domain Name System (DNS) plays a mission critical role in TCP/IP-based communications. DNS is a hierarchically distributed database used to map domain names to IP addresses (name resolution) and vice versa [4]. All IP-based devices regardless of their usage and location rely on the accuracy and integrity of DNS in their interaction.

In 1983, Paul Mockapetris introduced the DNS standard based on an architecture designed with scalability and availability in mind, with the expectation that its infrastructure would be improved over time [5]. The existing vulnerabilities in the DNS infrastructure, which weaken the security and trust in TCP/IP communications, have been widely recognized since the late 1990s [6]-[7]. While the networking communities' efforts in addressing these issues culminated in the introduction of DNS Security Extensions (DNSSEC) in 1999, the original infrastructure enhancements envisioned by Mockapetris in 1983, are still waiting to be rolled out.

- a. responds with an IP address for the requested domain name if it is authoritative for that zone
- b. answers with a delegation to another authoritative DNS server located closer to the requested domain name in the DNS hierarchy, or 3) provides a negative response if an answer to (1) or (2) is not known. Once the local recursive DNS server has obtained an IP address through this process, it updates its cache with the entry and forwards the response to the resolver.



Figure 1 depicts how a local recursive DNS server responds to a query for “foo.example.com” domain. As the first step, the DNS server checks its local cache to see if an entry for this domain already exists, stored based on a similar earlier query. If it does not, the server initiates an iterative query process, starting with the root’s authoritative DNS server, until it either resolves the IP address for the domain name or establishes that the requested domain name entry does not exist. At each point, an external authoritative DNS server receiving the query from the local recursive DNS server, either:

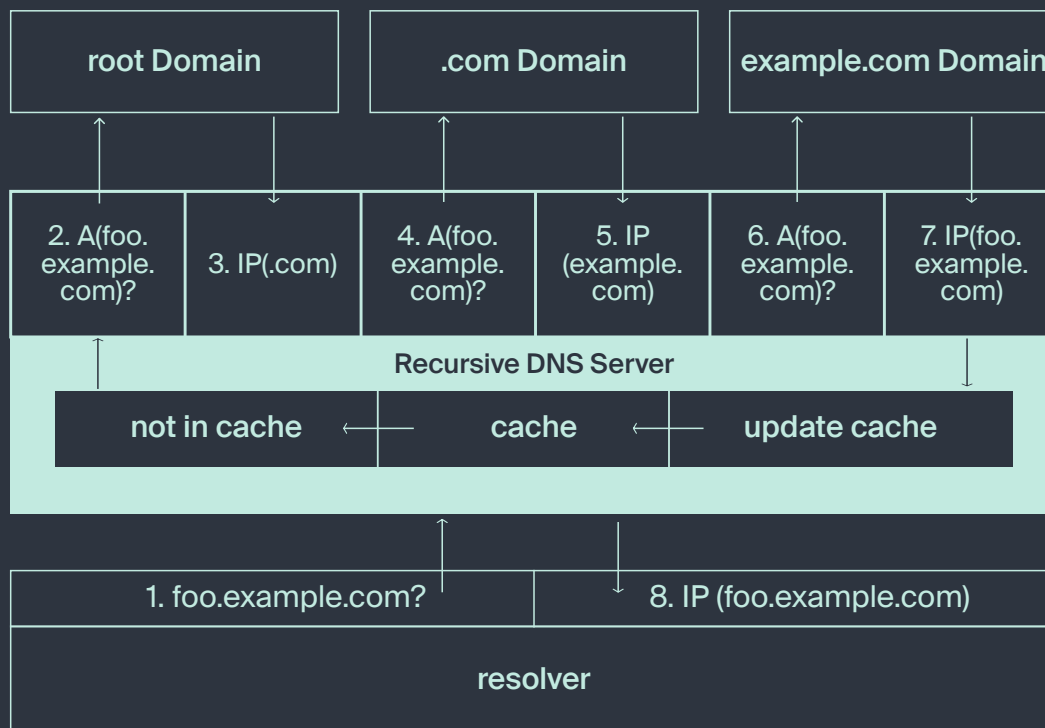


Figure 1:
The DNS name resolution proc



3. BIND and DNS Automation

BIND is by far the most widely used DNS implementation on the Internet. On the other hand, since it is based on unsupported open source software, it has to be maintained and patched manually. Using BIND for DNS services has some major security, administrative and management shortcomings [8]

1. The lack of administrative tools for managing day-to-day DNS routines, leads to input errors and DNS data corruption resulting in network downtime.
2. Software updates and security patches have to be implemented manually, which is extremely time-consuming and tedious.
3. The absence of a central management tool means that it is impossible to keep track of the IP address usage across the organization's network or to monitor use activities and changes made within the DNS environment.
4. All server backups, synchronization and disaster recovery activities have to be done manually, which increases the time to recovery.

The solution to these problems is to fully automate DNS management and administrative routines. In practice, this would mean centralizing and streamlining DNS management process, providing automatic data validation and patch updates and improving network availability and disaster recovery measures. All this will lead to a more secure and stable DNS service.

BIND is the most widely used DNS implementation on the Internet but has some major security, administrative and management shortcomings.



5. Secure Architecture of FusionLayer DNS Software Appliance

In this section, we focus on the issue of DNS security and introduce some of the security principles, which have been incorporated in our DNS software appliance, FusionLayer DNS

Figure 2 represents the basic architecture of FusionLayer DNS. It has been built on a Linux operating system, which has been specifically hardened and optimized for DNS use. In order to truly enhance the level of network security, we have implemented various security measures not only in the design principles but also at the application and the OS level. The three principles of “Defence in Depth” [9], “Least Privilege” [10] and “Default Deny” form the foundation of our DNS software appliance.

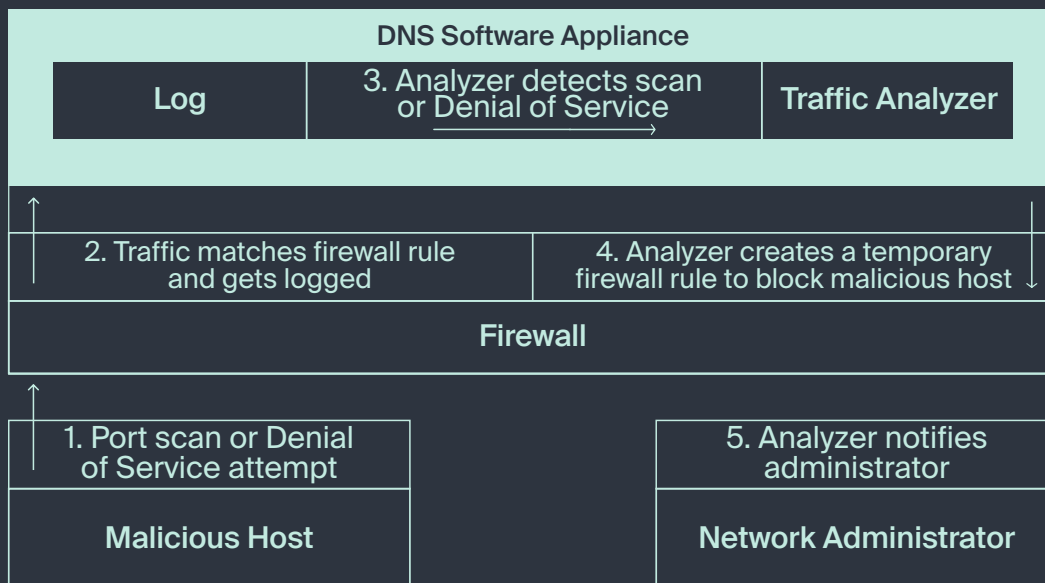
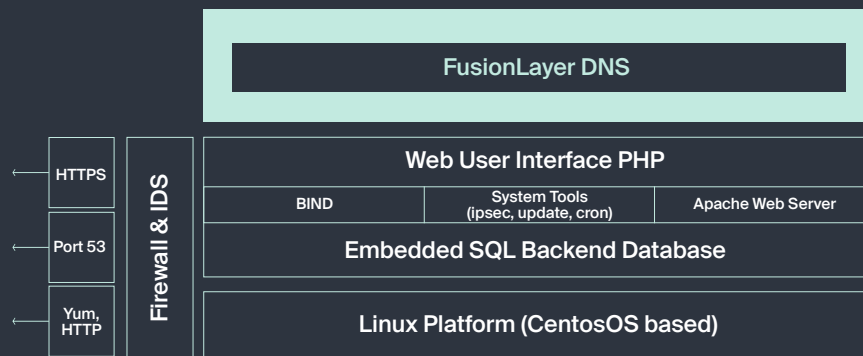


Figure 3: Intrusion Detection/Prevention System (IDS/IPS) deployed in FusionLayer DNS.



Figure 3, is in fact the novelty of our approach that has proven to be quite effective in increasing DNS security.

Defence in Depth is the concept of protecting a computer network with a series of defence mechanisms such that if one fails, another will already be in place to block an attack. In practice, this has meant including such measures as encrypted connections, in-host firewalls, Intrusion Detection/Prevention Systems (IDS/IPS) [11], running services with secure configurations, hardening the OS and periodic backups. The IDS/IPS included in FusionLayer DNS; as depicted in

The principle of Least Privilege is concerned with granular user privileges and rights. The idea behind this principle is to grant as minimal privileges as possible to permit a legitimate action, thereby enhancing the protection of data and functionality from faults and malicious behavior. In practice, this means that in our DNS software appliance architecture, DNS services are not run as root, the user-interface and shell access have multiple levels of user accounts and each user has specific rights and restricted privileges.

The principle of Default Deny demands that anything that is not explicitly allowed is denied, regardless of whether it is related to access, privileges, security-related attributes or other functions. In practice, this involves having a firewall which blocks all but explicitly allowed connections, explicitly permitting shell access for users, disabling all but necessary services and removing all but necessary packages.

The three security principles form the backbone of FusionLayer DNS. We further strengthened its architecture by incorporating various security measures such as DNSSEC, which is discussed in the next section.



6. Examples of Malicious Attacks Against DNS Servers

In this Section, we present two examples of malicious attacks against DNS Servers and their consequences. Figure 4 depicts the flow of communications between the client and the server while a man-in-the-middle or cache poisoning attack is underway. One of the main objectives of a man-in-the-middle attack is to eavesdrop on the communication between the servers and forge the responses to DNS queries, directing Internet traffic to fraudulent sites.

Meanwhile, in order to increase efficiency, decrease network traffic and improve response time, recursive DNS servers temporarily cache the name resolution responses they have obtained for a period of time referred to as the Time-To-Live (TTL).

Cache poisoning occurs when fraudulent DNS data is provided and cached in a recursive name server. As a result of cache poisoning, all future queries to the “poisoned” recursive DNS server would be directed to the fraudulent site until the record’s TTL has expired

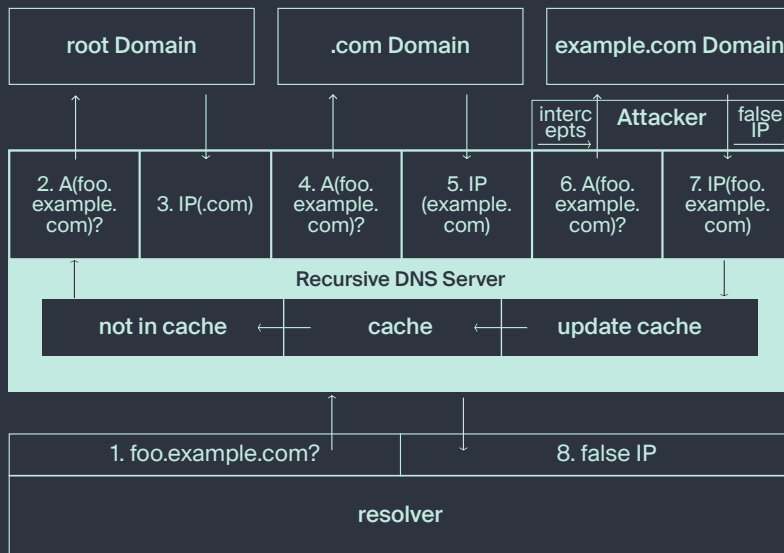


Figure 4: Vulnerabilities of DNS protocol.

Sender: signing and sending message

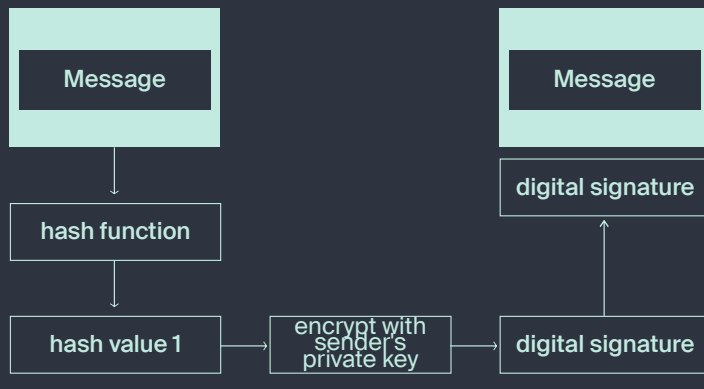
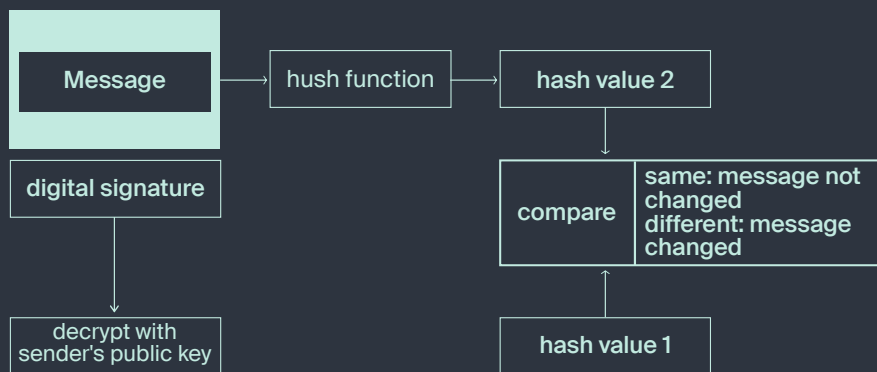


Figure 5: Asymmetric cryptography and digital signature

Receiver: receiving and verifying message





7. Mitigating Threats with DNSSEC

DNSSEC has been designed to solve some of the security problems associated with the DNS protocol. With the help of digital signatures [12], DNSSEC helps assure that the data received has in fact originated from the authorized source and that the message has not been tampered with in any way while travelling through the network. It also provides authenticated denial of existence, proving that the requested domain name entry does not exist. Having said that, DNSSEC does not protect the confidentiality of the data, nor does it prevent Denial of Service (DoS) attacks. As such, in order to further increase the trust in Internet and mitigate the latter two threats, in FusionLayer DNS we deploy DNSSEC alongside Secure Socket Layer (SSL) technology [13] and Intrusion Detection/Prevention Systems (IDS/IPS).

DNSSEC functions by requiring DNS zone data to be signed in order to assure the authenticity of the name resolution process. Its implementation is similar to Public Key Infrastructure (PKI) and relies on asymmetric cryptographic algorithm, which is represented in Figure 5. In asymmetric cryptography, a pair of private and public keys is used to encrypt and decrypt a message. The owner of the keys uses his/her private key to encrypt the message to be sent. The receiver, on the other hand, uses the public key of the sender, which he/she had obtained previously, to decrypt the message received. A message encrypted with a private key can only be decrypted with its corresponding public key and as such this approach can prove the authenticity of the source from which the data originated. Figure 5 also outlines how PKI can be used to prove the integrity of the message received, assuring that the data has not been modified while in transit.

Figure 6 represents the flow of communications for a name resolution process with DNSSEC enabled. In order to deploy DNSSEC, each child domain creates a pair of public/private keys and participates in the chain of trust by uploading its public key to its parent's domain and having it verified. This information is stored in the parent zone's Delegation Signer (DS) record and is forwarded alongside any response to the recursive name server. The child domain would then use its private key to sign all zone data for which it is authoritative. After receiving any response, the recursive DNS server would compare the public key of the child zone, contained in the DNSKEY record and received with the response, against the DS record of the parent's zone received in the previous response, to confirm that the public key sent by the child zone is in fact correct. The recursive name server would then use this verified DNSKEY to decrypt the DNSSEC signature contained in the RRSIG record in order to verify the authenticity and origin of the response. It would then follow on with the query and repeat the process until either the IP address of the requested domain name is resolved or it receives an authenticated denial of existence reply which indicates that the requested domain name entry does not exist.

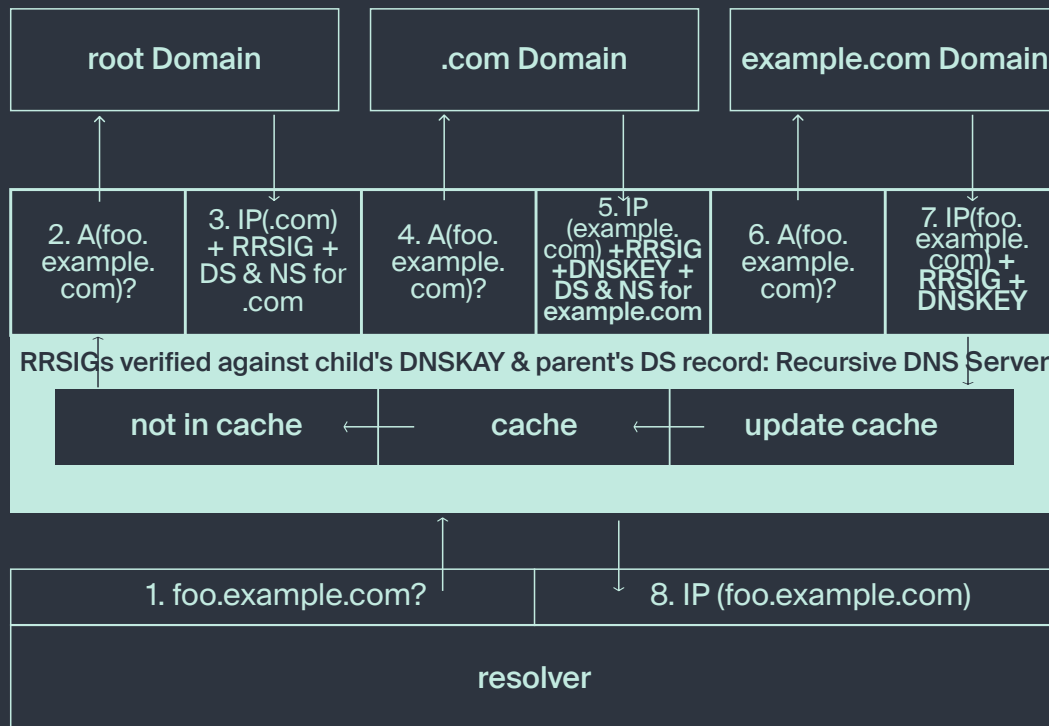


Figure 6: DNS name resolution with DNSSEC enabled.

Although the theory behind DNSSEC is relatively straight forward, implementing it introduces some operational challenges as listed below.

- When zones are signed, the number of resource records grows by a factor of 4-5, increasing the amount of data that has to be managed. In return, this will increase the amount of network traffic, although not significantly enough to warrant special measure.
- When signed, formerly static DNS entries will become dynamic because the private keys used to sign the zones have to be rolled over at regular intervals. According to DNSSEC Operational Practices [14], it is recommended that zone signing keys are rolled over once every month. This practice is aimed at increasing the level of security by preventing theft and brute force attacks on private keys.
- As signed zones contain resource records that are effectively very long keys, managing the keys manually becomes more error-prone. This is due to the fact that as the key length grows, so does the possibility of human error while manually inputting, updating and managing the keys.

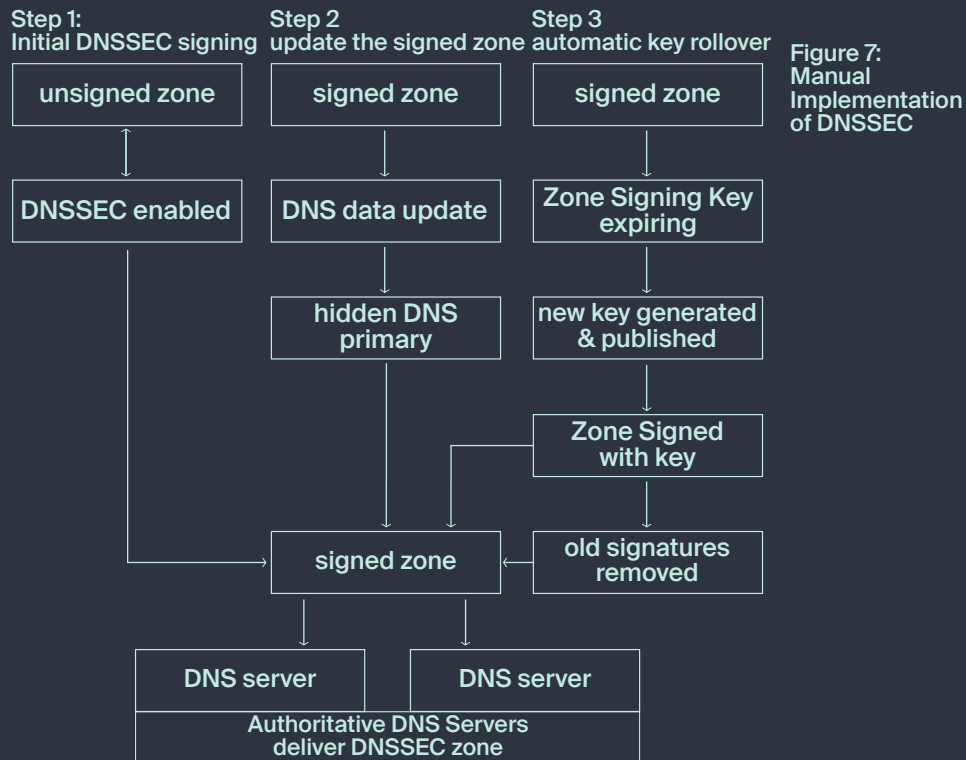


Figure 7: Manual Implementation of DNSSEC

Figure 7 depicts the steps that have to be taken in order to manually deploy DNSSEC. As elaborated above, it is these operational challenges of manually implementing DNSSEC that have prevented many organizations from deploying DNSSEC in their network. As such, deploying DNSSEC, on either external or internal networks, is a process that greatly benefits from automation. In the next section, we elaborate on the criteria necessary to successfully deploy a secure DNS environment with DNSSEC enabled using software appliances.



8. Countering Predatory Practices in DNS Environments

To successfully deploy a DNS environment that addresses the forms of cyber attacks discussed in this paper, we have identified three key criteria that should be met.

1. The deployed DNS service should support DNSSEC and include automated key management functionality used in the authentication of the name resolution process.
2. It should include Secure Socket Layer (SSL) technology and built in Intrusion Detection/Protection System (IDS/IPS) providing protection against Denial of Service attacks.
3. To improve network performance, security and availability, DNS services should be based on a dedicated and optimized server stack with just enough Operating System (JeOS) [15] used to minimize the number of attack vectors.

Combining DNSSEC with the methodology described above has numerous advantages. Firstly, it improves the reliability and availability of core network services while reducing operating expenses. Secondly, when deployed based on the security principles discussed previously, this approach is proven to be extremely effective in mitigating predatory practices discussed in this paper. Thirdly, the fact that FusionLayer DNS is distributed as a software appliance means that it could be deployed either on a virtual machine or any x86-based hardware platform already in use. This platform independence supports recent trend towards cloud computing initiatives that reduce computing costs. Finally, instead of requiring dedicated physical servers, our approach allows numerous instances of software appliances to be consolidated as virtual machines on any computing platform already in use and thus helps reduce the number of physical servers required.



9. Conclusion

DNS is one of the core services necessary for network connectivity. Despite its mission critical role, its inherent design has made it an easy target for numerous forms of cyber attacks. DNSSEC was developed as a protocol that mitigates some of the security issues associated with DNS. Its successful deployment though requires automation; both in terms of installation and management processes. At the same time, to be truly effective in improving DNS security, DNSSEC has to be deployed in combination with other security measures to address a wider spectrum of cyber attacks.

In this whitepaper, we introduced FusionLayer DNS, which is developed based on a methodology that addresses DNS security issues out of the box. FusionLayer DNS provides a simple and secure DNS solution designed based on stringent security principles which meets the standard requirements of today's networking environments. It has already been deployed in various organizations world wide, both in public and private sector networks in all industry verticals, creating significant cost saving while introducing an enhanced level of network security.

About FusionLayer

Managing complex corporate and telecom networks is a challenge where the cost of failure is enormous. FusionLayer collates all network information into a single Network Source of Truth, accessed securely by both engineers and automation to eliminate the chance of network downtime – on-premise, at the edge, and in the public cloud. This provides our customers with reassuring realtime information, so their digitalized operations can connect 24x7x365.



References

1. Carl Von Clausewitz, Colonel F.N. Maude and J.J. Graham, "On War", p.152, Wilder Publication, 2008.
2. WIPO, IP Services, (WO/2007/107634) "Applianced Domain Name Server", <http://www.wipo.int/pctdb/en/wo.jsp?WO=2007107634>
3. Ganguly, J. Yin, H. Shaikh, D. Chess, T. Eilem, R. Figueiredo, J. Hansom, A. Mohindra and G. Pacifici, "Reducing complexity of software deployment with delta configuration", Proceedings of the 10th IFIP/IEEE International Symposium on Integrated Network Management, pp. 729-732, Munich, Germany, May 2007.
4. P. Mockapetris, "Domain names - concepts and facilities", Internet RFC 1034, November 1987.
5. S. Lawson, "Inventor looks to DNS's past, future", Computerworld June 2003, www.computerworld.com.au/article/70040/inventor_looks_dns_past_future/
6. S. M. Bellovin, "Using domain name system for system break-ins", Proceedings of the 5th Usenix Security Symposium, Salt Lake City, UT, 1995.
7. R. Chandramouli and S. Rose, "Secure Domain Name System (DNS) deployment guide - recommendations of the National Institute of Standards and Technology", Sponsored by the Department of Homeland Security, April 2010.
8. P. Vixie, "DNS and BIND security issues", Proceedings of the 5th Usenix Security Symposium, Salt Lake City, UT, 1995.
9. L. Smith, "Understanding concepts in the defence in depth strategy", Proceedings of IEEE 37th International Carnahan Conference on Security Technology, pp. 8-16, 2003.
10. N. Jing-xuan, H. Hong-jun, L. Li, L. Peng and D. Li-ming, "Discussion on minimizing file access privilege", Proceedings of International Conference on Multimedia and Information Technology, p. 801, 2008.
11. J. McHugh, A. Christie and J. Allen, "Defending yourself: the role of intrusion detection systems", IEEE Software, Vol. 17, Is. 5, PP. 42-51, September/October 2000.
12. H. L. Kesterson, "Digital signatures, whom do you trust?" IEEE Proceedings of Aerospace Conference, Vol. 4, p. 159, February 1997.
13. A.C. Weaver, "Secure Sockets Layer", Computer, Vol.39, Issue 4, pp. 88-90, April 2006.
14. O. Kolkman and W. Mekking, "DNSSEC operational practices, Version 2", RFC 4641 (Proposed Standard), September 2006.
15. <http://www.ietf.org/rfc/rfc4641.txt>
16. D. Geer, "The OS faces a brave new world", Computer, Vol. 42, Issue 10, pp. 15-17, October 2009.