DNS Advanced Attacks and Analysis

Adam Ali.Zare Hudaib *adamhudaib@gmail.com*
Licensed Penetration Tester
Certified Ethical Hacker
Network Security Defence
Research &Troubleshooting
CEH , ECSA , LPT , WCNA
Poland

Esra'a Ali Zare Hudaib *israa_hudieb@eng.hu.edu.jo*
Computer & Engineering Department
The Hashemite University
Amman . Jordan

Abstract

Nowadays DNS is used to load balance, failover, and geographically redirect connections. DNS has become so pervasive it is hard to identify a modern TCP/IP connection that does not use DNS in some way. Unfortunately, due to the reliability built into the fundamental RFC-based design of DNS, most IT professionals don't spend much time worrying about it. If DNS is maliciously attacked — altering the addresses it gives out or taken offline the damage will be enormous. Whether conducted for political motives, financial gain, or just the notoriety of the attacker, the damage from a DNS attack can be devastating for the target.

In this research we will review different DNS advanced attacks and analyze them. We will survey some of the most DNS vulnerabilities and ways of DNS attacks protection.

Keywords: DNS, DoS, Cache Poisoning, DNSSEC, DNS Hijacking.

1. INTRODUCTION

Denial of Service (DoS) attacks can be classified into two major categories. In the first one, the adversary featly crafts packets trying to exploit vulnerabilities in the implemented software(service or protocol) at the target side. This class of attacks includes outbreaks like the ping of death[1]. In the second one, the aggressor attempts to overwhelm critical system's resources, i.e. memory, CPU, network bandwidth by creating numerous of well-formed but bogus requests. This type of attack is also well known as flooding. DoS attacks are a threat to almost every service in the Internet and DNS is no exception. These attacks against or related to DNS servers are also classified into two types. One is to directly flood DNS servers by sending a large number of DNS requests or other useless traffic.

Since the DNS servers cannot easily distinguish the legitimate requests from the attack traffic, they would simply accept both of them and send the responses [2]. The effective and deployable defense against this attack is to over-provision the network capacity and numbers of servers [3].The other attack strategy is to exploit DNS servers to amplify attack traffic. The attacker craftsa DNS request that gets a response significantly larger than the request itself, e.g., a 50-byterequest for a 500-byte response. The amplified response is replied to a spoofed third-party victim machine. Under this attack, both the amplifying DNS server's upstream bandwidth and the third-party machine's downstream bandwidth could be exhausted. Due to traffic amplification, an attacker can exhaust the bandwidth of its victims even if his bandwidth is 10 times smaller [4].An effective defense against spoofing-based DoS attacks on DNS servers requires source address spoof detection. Assuming a DNS server can distinguish between spoofed requests from real ones, it can selectively drop those spoofed ones with little collateral damage.

In this paper, we analyze different types of the DNS amplification attacks and ways of protection.

2. DNS ADVANCED ATTACKS AND ANALYSIS

2.1 DNS and The Most Common Security Issues

The Domain Name System (DNS) is a hierarchical, distributed database that contains mappings between names and other information, such as IP addresses.

DNS allows users to locate resources on the network by converting friendly, human-readable names like www.microsoft.com to IP addresses that computers can connect to. An often-used analogy to explain the Domain Name System is that it serves as the "phone book" for the Internet by translating human-friendly computer hostnames into IP addresses. For example, www.xyz.com translates to the addresses 20.52.88.12 (IPv4).

By default DNS works on port number 53 on TCP (Transmission Control Protocol) and UDP (user Datagram Protocol).

DNS is a crucial part of a network and hence securing DNS essentially become quite important. If a DNS is compromised, an attacker can easily prevent normal operations going in the network, can route computers to whatever spoofed IP address or resources he wants, steal information which and a lot of such malicious activity [6].

There are the essential functions of a DNS:

1. DNS is responsible for locating services like DC, etc. for authenticating the services on the network There are the essential functions of a DNS:
2. It is responsible for locating resources like Web Servers, Mail servers, etc. on the network.
3. And obviously, translating Computer names to IP and vice versa.

There are several most common security issues for DNS.

1. Unauthorized Authoritative DNS Record Changes – Changes to authoritative DNS records which point end users to computer systems outside of your control can have the most damage to your business's brand. This type of attack is typically done to either send users to a site which provides a negative marketing message, or to a location mirroring your site where account credentials can be harvested. This attack is particularly devastating because users are typically unaware anything untoward has happened [7].

2. Denial of Service Attacks – Denial of Service (DoS) or Distributed Denial of Service Attacks (DDoS) are done to make your DNS service unavailable and thus create the impression your business is offline or closed down (website, portals, VPNs, FTP, VoIP, email, etc.). This type of attack is one of the easiest to perform and can be one of the hardest to defend against. One of the least recognized impacts to a business that suffers a DNS outage from a DDoS attack is the negative effect it has on your search engine rankings.

3. Recursive DNS Spoofing/Cache Poisoning – Outside of a business's control, the Recursive DNS server an end user utilizes is typically set by the user's network administrator. Recursive DNS servers communicate the Authoritative DNS records a business sets to an end user's device. Unfortunately, many Recursive DNS servers are not well maintained or protected and can be easily compromised to give out false responses. This has the same down stream effect of an Unauthorized Authoritative DNS record change [8].

The main vulnerabilities:

1. Denial of Service.
   a. Harm and block DNS traffic.
   DNS is an effective DOS attack vector for a few reasons: DNS usually uses the UDP as its transport; most of autonomous systems allow source-spoofed packets to enter their network; there is a lot of Open DNS Resolvers on the Internet.

   The attack overloads the system by using: DNS reflectors, amplification, botnet; DDOS, recursive malformed requests, impersonation

2. Data Modification.
   a. Query/Request Redirection.
   b. DNS cache poisoning.
   c. DNS ID hacking.

   Query/Request redirection uses man-in-the-middle position, breaks of the chain of trust. DNS spoofing forges a fake answer. DNS ID hacking succeeds in impersonating a DNS server. DNS cache poisoning sends user to malicious site.

3. Zone Enumeration.
   Not really considered as an attack. Most considered as a threat as it allows attackers to gather information Precedes an attempt at an attack.

4. Tunnels.
   Uses DNS TCP transport mechanism. DNS TCP is used for: failover transport: switch from UDP to TCP; secondary zone transfer; DNSSEC and IPv6 traffic; EDNS is often badly supported by customer network.

   Attacks use TCP channel to tunnel other protocol and run malicious software [9].

2.2 Types of DNS Attacks and How To Deal with Them

DNS servers work by translating IP addresses into domain names. When DNS is compromised, several things can happen. However, compromised DNS servers are often used by attackers one of two ways. The first thing an attacker can do is redirect all incoming traffic to a server of their choosing. This enables them to launch additional attacks, or collect traffic logs that contain sensitive information.

The second thing an attacker can do is capture all in-bound email. More importantly, this second option also allows the attacker to send email on their behalf, using the victim organization's domain and cashing-in on their positive reputation. Making things worse, attackers could also opt for a third option, which is doing both of those things.

There are three common types of DNS attacks.

The first type of DNS attack is called a cache poisoning attack. This can happen after an attacker is successful in injecting malicious DNS data into the recursive DNS servers that are operated by many ISPs. These types of DNS servers are the closest to users from a network topology perspective, von Wallenstein wrote, so the damage is localized to specific users connecting to those servers.

If DNSSEC is impractical or impossible, another workaround is to restrict recursion on the name servers that need to be protected. Recursion identifies whether a server will only hand out information it has stored in cache, or if it is willing to go out on the Internet and talk to other servers to find the best answer.

"Many cache poisoning attacks leverage the recursive feature in order to poison the system. So by limiting recursion to only your internal systems, you limit your exposure. While this setting will not resolve all possible cache poisoning attack vectors, it will help you mitigate a good portion of them," Chris Brenton, Dyn Inc.'s Director of Security [10].

The second type of DNS attack happens when attackers take over one or more authoritative DNS servers for a domain. In 2009, Twitter suffered a separate attack by the Iranian Cyber Army. The group altered DNS records and redirected traffic to propaganda hosted on servers they controlled. The ability to alter DNS settings came after the Iranian Cyber Army compromised a Twitter staffer's email account, and then used that account to authorize DNS changes. During that incident Dyn Inc. was the registrar contacted in order to process the change request. Defense against these types of attacks often include strong passwords, and IP-based ACLs (acceptable client lists). Further, a solid training program that deals with social engineering will also be effective. Unfortunately, all the time and resources in the world can be placed into securing a webserver, but if an attacker can attack the authoritative server and point the DNS records at a different IP address, to the rest of the world its still going to look like you've been owned. In fact it's worse because that one attack will also permit them to redirect your email or any other service you are offering. So hosting your authoritative server with a trusted authority is the simplest way to resolve this problem.

The third type of DNS attack is also the most problematic to undo. It happens when an attacker compromised the registration of the domain itself, and then uses that access to alter the DNS servers assigned to it.

"At this time, those authoritative nameservers answered all queries for the affected domains. What makes this attack so dangerous is what’s called the TTL (time to live). Changes of this nature are globally cached on recursive DNS servers for typically 86,400 seconds, or a full day. Unless operators are able to purge caches, it can take an entire day (sometimes longer) for the effects to be reversed," von Wallenstein wrote. The main advice for authoritative DNS is to host authoritative servers within the organization, allowing for complete control [11].

2.3 Amplification Attacks

The amplification attacks are some of the largest, as measured by the number of Gigabits per second (Gbps). That size of an attack is enough to cripple even a large web host. Even from a cost perspective, the attack doesn't end up adding to our bandwidth bill because of the way in which we're charged for wholesale bandwidth.

DNS Amplification Attacks are a way for an attacker to magnify the amount of bandwidth they can target at a potential victim. Imagine you are an attacker and you control a botnet capable of sending out 100Mbps of traffic. While that may be sufficient to knock some sites offline, it is a relatively trivial amount of traffic in the world of DDoS. In order to increase your attack's volume, you could try and add more compromised machines to your botnet. That is becoming increasingly difficult. Alternatively, you could find a way to amplify your 100Mbps into something much bigger [12].

The original amplification attack was known as a SMURF attack. A SMURF attack involves an attacker sending ICMP requests (i.e., ping requests) to the network's broadcast address (i.e., X.X.X.255) of a router configured to relay ICMP to all devices behind the router. The attacker spoofs the source of the ICMP request to be the IP address of the intended victim. Since ICMP does not include a handshake, the destination has no way of verifying if the source IP is legitimate. The router receives the request and passes it on to all the devices that sit behind it. All those devices then respond back to the ping. The attacker is able to amplify the attack by a multiple of how ever many devices are behind the router (i.e., if you have 5 devices behind the router then the attacker is able to amplify the attack 5x, see the figure 1below).

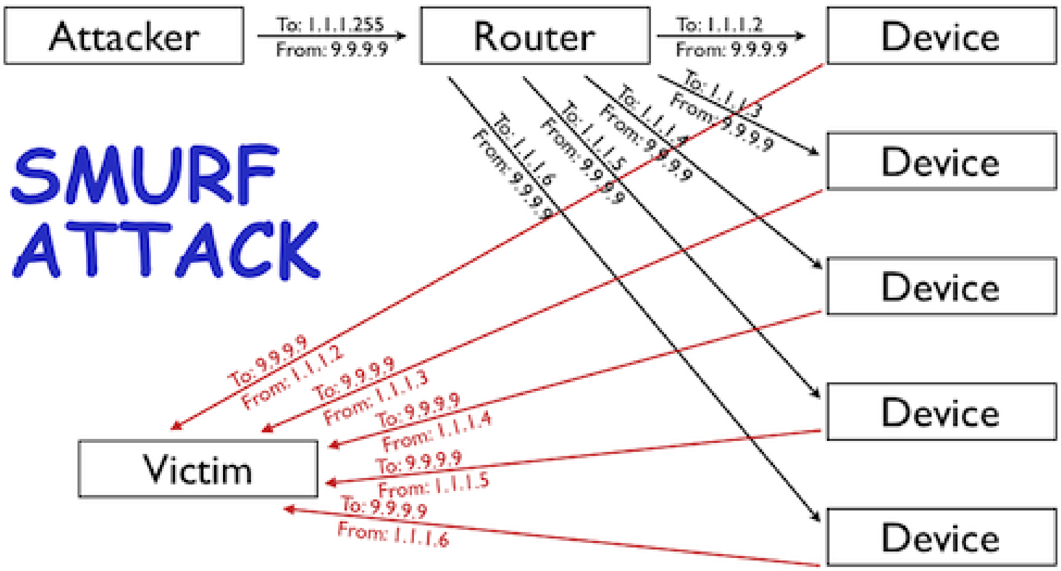

FIGURE 1: SMURF Attack.

SMURF attacks are largely a thing of the past. For the most part, network operators have configured their routers to not relay ICMP requests sent to a network's broadcast address. However, even as that amplification attack vector has closed, others remain wide open [13].

There are two criteria for a good amplification attack vector: 1) query can be set with a spoofed source address (e.g., via a protocol like ICMP or UDP that does not require a handshake); and 2) the response to the query is significantly larger than the query itself. DNS is a core, ubiquitous Internet platform that meets these criteria and therefore has become the largest source of amplification attacks.

DNS queries are typically transmitted over UDP, meaning that, like ICMP queries used in a SMURF attack, they are fire and forget. As a result, their source attribute can be spoofed and the receiver has no way of determining its veracity before responding. DNS also is capable of generating a much larger response than query.

The key term that I used a couple times so far is "open DNS resolver." The best practice, if you're running a recursive DNS resolver is to ensure that it only responds to queries from authorized clients. In other words, if you're running a recursive DNS server for your company and your company's IP space is 5.5.5.0/24 (i.e., 5.5.5.0 - 5.5.5.255) then it should only respond to queries from that range. If a query arrives from 9.9.9.9 then it should not respond.

The problem is, many people running DNS resolvers leave them open and willing to respond to any IP address that queries them. This is a known problem that is at least 10 years old. What has happened recently is a number of distinct botnets appear to have enumerated the Internet's IP space in order to discover open resolvers. Once discovered, they can be used to launch significant DNS Amplification Attacks.

Nowadays there's been an increase in big DDoS attacks. It's in large part because the network operators listed above have continued to allow open resolvers to run on their networks and the attackers have begun abusing them.

2.4 DNS Hijacking

DNS hijacking or DNS redirection is the practice of subverting the resolution of Domain Name System (DNS) queries. This can be achieved by malware that overrides a computer's TCP/IP

configuration to point at a rogue DNS server under the control of an attacker, or through modifying the behavior of a trusted DNS server so that it does not comply with internet standards. These modifications may be made for malicious purposes such as phishing, or for self-serving purposes by Internet service providers (ISPs) to direct users' web traffic to the ISP's own web servers where advertisements can be served, statistics collected, or other purposes of the ISP; and by DNS service providers to block access to selected domains as a form of censorship [14].
A number of consumer ISPs such as Cablevision's Optimum Online, Comcast, Time Warner, Cox Communications, RCN, Rogers, Charter Communications, Verizon, Virgin Media, Frontier Communications, Bell Sympatico, UPC, T-Online, Optus, Mediacom, ONO, TalkTalk and Bigpond (Telstra) use DNS hijacking for their own purposes, such as displaying advertisements or collecting statistics. This practice violates the RFC standard for DNS (NXDOMAIN) responses and can potentially open users to cross-site scripting attacks.

The concern with DNS hijacking involves this hijacking of the NXDOMAIN response. Internet and intranet applications rely on the NXDOMAIN response to describe the condition where the DNS has no entry for the specified host. If one were to query the invalid domain name (fakeexample.com), one should get an NXDOMAIN response - informing the application that the name is invalid and taking the appropriate action (for example, displaying an error or not attempting to connect to the server). However, if the domain name is queried on one of these non-compliant ISPs, one would always receive a fake IP address belonging to the ISP. In a web browser, this behavior can be annoying or offensive as connections to this IP address display the ISP redirect page of the provider, sometimes with advertising, instead of a proper error message. However, other applications that rely on the NXDOMAIN error will instead attempt to initiate connections to this spoofed IP address, potentially exposing sensitive information.

Examples of functionality that breaks when an ISP hijacks DNS:

- Roaming laptops that are members of a Windows Server domain will falsely be led to believe that they are back on a corporate network because resources such as domain controllers, email servers and other infrastructure will appear to be available. Applications will therefore attempt to initiate connections to these corporate servers, but fail, resulting in degraded performance, unnecessary traffic on the internet connection and timeouts.

- Many small office and home networks do not have their own DNS server, relying instead on broadcast name resolution. Many versions of Microsoft Windows default to prioritizing DNS name resolution above NetBIOS name resolution broadcasts; therefore, when an ISP DNS server returns a (technically valid) IP address for the name of the desired computer on the LAN, the connecting computer uses this incorrect IP address and inevitably fails to connect to the desired computer on the LAN. Workarounds include using the correct IP address instead of the computer name, or the DhcpNodeType registry value to change name resolution service ordering.

- Browsers such as Firefox no longer have their 'Browse By Name' functionality (Where keywords typed in the address bar take you to the closest matching site.).

- The local DNS client built into modern operating systems will cache results of DNS searches for performance reasons. If a client switches between a home network and a VPN, false entries may remain cached, thereby creating a service outage on the VPN connection.DNSBL anti-spam solutions rely on DNS; false DNS results therefore interfere with their operation [15].

- Confidential user data might be leaked by applications that are tricked by the ISP into believing that the servers they wish to connect to are available.

- User choice over which search engine to consult in the event of a URL being mistyped in a browser is removed as the ISP determines what search results are displayed to the user; functionality of applications like the Google Toolbar do not work correctly.

- Computers configured to use a split tunnel with a VPN connection will stop working because intranet names that should not be resolved outside the tunnel over the public Internet will start resolving to fictitious addresses, instead of resolving correctly over the VPN tunnel on a private DNS server when an NXDOMAIN response is received from the Internet. For example, a mail client attempting to resolve the DNS A record for an internal mail server may receive a false DNS response that directed it to a paid-results web server, with messages queued for delivery for days while retransmission was attempted in vain.

- It breaks Web Proxy Autodiscovery Protocol (WPAD) by leading web browsers to believe incorrectly that the ISP has a proxy server configured.

- It breaks monitoring software. For example, if we periodically contact a server to determine its health, a monitor will never see a failure unless the monitor tries to verify the server's cryptographic key.

In some cases, the ISPs provide subscriber-configurable settings to disable hijacking of NXDOMAIN responses. Correctly implemented, such a setting reverts DNS to standard behavior. Other ISPs, however, instead use a web browser cookie to store the preference. In this case, the underlying behavior is not resolved: DNS queries continue to be redirected, while the ISP redirect page is replaced with a counterfeit dns error page. Applications other than web-browsers cannot be opted out of the scheme using cookies as the opt-out targets only the HTTP protocol, when the scheme is actually implemented in the protocol-neutral DNS system.

2.5 DoS Attacks

Perpetrators of DoS attacks typically target sites or services hosted on high-profile web servers such as banks, credit card payment gateways, and even root nameservers. This technique has now seen extensive use in certain games, used by server owners, or disgruntled competitors on games, such as server owners' popular Minecraft servers. Increasingly, DoS attacks have also been used as a form of resistance. Richard Stallman has stated that DoS is a form of 'Internet Street Protests’. The term is generally used relating to computer networks, but is not limited to this field; for example, it is also used in reference to CPU resource management [16].

One common method of attack involves saturating the target machine with external communications requests, so much so that it cannot respond to legitimate traffic, or responds so slowly as to be rendered essentially unavailable. Such attacks usually lead to a server overload. In general terms, DoS attacks are implemented by either forcing the targeted computer(s) to reset, or consuming its resources so that it can no longer provide its intended service or obstructing the communication media between the intended users and the victim so that they can no longer communicate adequately.

Denial-of-service attacks are considered violations of the Internet Architecture Board's Internet proper use policy, and also violate the acceptable use policies of virtually all Internet service providers. They also commonly constitute violations of the laws of individual nations.

The United States Computer Emergency Readiness Team (US-CERT) defines symptoms of denial-of-service attacks to include:

- Unusually slow network performance (opening files or accessing web sites);
- Unavailability of a particular web site;
- Inability to access any web site;
- Dramatic increase in the number of spam emails received—(this type of DoS attack is considered an e-mail bomb);

- Disconnection of a wireless or wired internet connection;
- Long term denial of access to the web or any internet services [17].

Denial-of-service attacks can also lead to problems in the network 'branches' around the actual computer being attacked. For example, the bandwidth of a router between the Internet and a LAN may be consumed by an attack, compromising not only the intended computer, but also the entire network or other computers on the LAN.

If the attack is conducted on a sufficiently large scale, entire geographical regions of Internet connectivity can be compromised without the attacker's knowledge or intent by incorrectly configured or flimsy network infrastructure equipment.

A denial-of-service attack is characterized by an explicit attempt by attackers to prevent legitimate users of a service from using that service. There are two general forms of DoS attacks: those that crash services and those that flood services [18].

A DoS attack can be perpetrated in a number of ways. Attacks can fundamentally be classified into five families:

1. Consumption of computational resources, such as bandwidth, memory, disk space, or processor time.
2. Disruption of configuration information, such as routing information.
3. Disruption of state information, such as unsolicited resetting of TCP sessions.
4. Disruption of physical network components.
5. Obstructing the communication media between the intended users and the victim so that they can no longer communicate adequately.

A DoS attack may include execution of malware intended to:

- Max out the processor's usage, preventing any work from occurring.
- Trigger errors in the microcode of the machine.
- Trigger errors in the sequencing of instructions, so as to force the computer into an unstable state or lock-up.
- Exploit errors in the operating system, causing resource starvation and/or thrashing, i.e. to use up all available facilities so no real work can be accomplished or it can crash the system itself
- Crash the operating system itself.

In most cases DoS attacks involve forging of IP sender addresses (IP address spoofing) so that the location of the attacking machines cannot easily be identified and to prevent filtering of the packets based on the source address.

A Distributed Denial of Service Attack (DDoS) occurs when multiple systems flood the bandwidth or resources of a targeted system, usually one or more web servers. This is the result of multiple compromised systems (for example a botnet) flooding the targeted system with traffic. When a server is overloaded with connections, new connections can no longer be accepted. The major advantages to an attacker of using a distributed denial-of-service attack are that multiple machines can generate more attack traffic than one machine, multiple attack machines are harder to turn off than one attack machine, and that the behavior of each attack machine can be stealthier, making it harder to track and shut down. These attacker advantages cause challenges for defense mechanisms. For example, merely purchasing more incoming bandwidth than the current volume of the attack might not help, because the attacker might be able to simply add more attack machines.

Malware can carry DDoS attack mechanisms; one of the better-known examples of this was MyDoom. Its DoS mechanism was triggered on a specific date and time. This type of DDoS involved hardcoding the target IP address prior to release of the malware and no further interaction was necessary to launch the attack.

A system may also be compromised with a trojan, allowing the attacker to download a zombie agent, or the trojan may contain one. Attackers can also break into systems using automated tools that exploit flaws in programs that listen for connections from remote hosts. This scenario primarily concerns systems acting as servers on the web. Stacheldraht is a classic example of a DDoS tool. It utilizes a layered structure where the attacker uses a client program to connect to handlers, which are compromised systems that issue commands to the zombie agents, which in turn facilitate the DDoS attack. Agents are compromised via the handlers by the attacker, using automated routines to exploit vulnerabilities in programs that accept remote connections running on the targeted remote hosts. Each handler can control up to a thousand agents. In some cases a machine may become part of a DDoS attack with the owner's consent, for example, in Operation Payback, organized by the group Anonymous [19].

These collections of systems compromisers are known as botnets. DDoS tools like Stacheldraht still use classic DoS attack methods centered on IP spoofing and amplification like smurf attacks and fraggle attacks (these are also known as bandwidth consumption attacks). SYN floods (also known as resource starvation attacks) may also be used. Newer tools can use DNS servers for DoS purposes. Unlike MyDoom's DDoS mechanism, botnets can be turned against any IP address. Script kiddies use them to deny the availability of well known websites to legitimate users. More sophisticated attackers use DDoS tools for the purposes of extortion – even against their business rivals.

Simple attacks such as SYN floods may appear with a wide range of source IP addresses, giving the appearance of a well distributed DoS. These flood attacks do not require completion of the TCP three way handshake and attempt to exhaust the destination SYN queue or the server bandwidth. Because the source IP addresses can be trivially spoofed, an attack could come from a limited set of sources, or may even originate from a single host. Stack enhancements such as syn cookies may be effective mitigation against SYN queue flooding, however complete bandwidth exhaustion may require involvement [20].

If an attacker mounts an attack from a single host it would be classified as a DoS attack. In fact, any attack against availability would be classed as a Denial of Service attack. On the other hand, if an attacker uses many systems to simultaneously launch attacks against a remote host, this would be classified as a DDoS attack.

2.6 Prevention of The Attacks

There are some ways of DNS attacks prevention.

1. Usage of the best practices configurations.
    a. Run software in secure environment.
    b. Identify data flow.
    c. ACLs.
    d. Stealth Architecture.

2. Enabling DNSSEC.

3. Monitoring DNS Traffic.
    a. Short term analysis (peak detection).
    b. Long term analysis (abnormal behavior).

By server secure environment is meant: running up-to-date software version; checking that the operating system is also having all security fixes; efficient IP comes into an appliance format with a single upgrade process that updates: operating system, services, software.

Also you must identify data flow; run caching, resolver, authoritative server. You should separate the functions as possible and disable unwanted features It will help into preventing attacks. A public authoritative server should never be recursive [21].

Access control list is very important too.

ACLs are used to control what information will be published. With data flow identification, you can choose who will be able to:

- Allow query (server and zone level);
- Allow query cache (server level);
- Allow transfer (server and zone level);
- Allow update (zone level);
- Blackhole (server level);
- Negative Cache (zone level).

There are library of SmartArchitecture DNS templates. One of them is DNS Stealth: State of the Art Internet DNS architecture (see the figure 2).

DNSSEC is used to protect against query/request redirection. DNSSEC creates a chain of trust between the client and the authoritative server. Based on key exchange inside specific signed resource records.

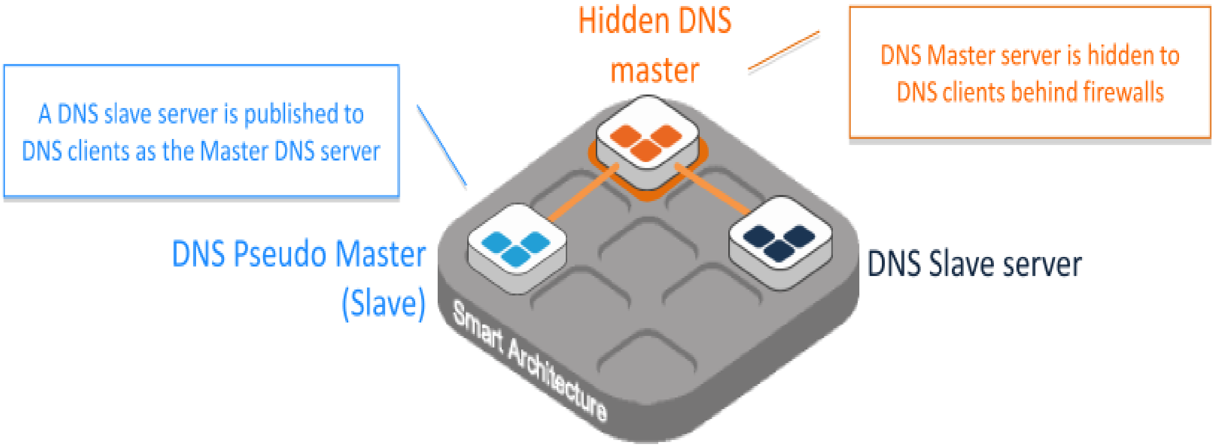
Hidden DNS master
DNS Master server is hidden to DNS clients behind firewalls
A DNS slave server is published to DNS clients as the Master DNS server
DNS Pseudo Master (Slave)
DNS Slave server
Smart Architecture

FIGURE 2: DNS Stealth Architecture.

2.7 Security Best Practices

Registrar Lock Your Domain Names – One of the simplest protections that can be used is lock all of your domain names at your registrar [22].

Outsource DNS Services – In today's world it is typically unrealistic to maintain your own DNS name servers in a way that both protects them from attacks and maintains global performance, and it is naive to use the free DNS services of a domain registrar. Cloud based managed service providers are your best bet for both Authoritative & Recursive DNS. Neustar (UltraDNS), DynDNS, Verisign, Amazon (Route 53), and Community DNS (European focused) are some of the top IP Anycasted Authoritative DNS providers to consider. OpenDNS, Neustar (UltraDNS), DynDNS & Google are the top Recursive DNS providers to consider. The investment in a cloud based DNS provider will protect your business from many of the common attacks, and free you from having to manage the devices yourself.

Utilize Strong Access Controls – As with any critical IT infrastructure, only allow users access to DNS administration for what they need to manage, lock down access to these critical accounts to known IP ranges, utilize strong password controls, and whenever possible use two factor authentication.

Activate DNSSEC On Your Domain Names – DNSSEC counters cache poisoning attacks by verifying the authenticity of responses received from name servers. It effectively prevents responses from being tampered with, because in practice, signatures are almost impossible to forge without access to the private keys.

Continuously Monitor Your Critical Services & DNS Records – Utilize an advanced SIEM like the one available from Savanture to monitor all of critical services and monitor your DNS records for changes from outside your network. UltraTools.com provides a free DNS monitoring service that many top organizations use. Additionally, monitoring the activity level on your services can show when traffic suddenly gets directed away.

Promote The Use of Protected Recursive DNS Servers – Usage of one of the top Recursive DNS providers for network. Many times there is no cost to this, only a configuration change [23].

Protect DNS Service Against DDoS Attacks – Usafe of one of the top Authoritative DNS providers that provides DDoS protection for DNS service. For public facing services that require DDoS protection, lower your DNS Time to Live (TTLs) settings to 300 (5min) so it can redirect traffic quickly if you come under attack and need protection.

3. CONCLUSIONS

We presented our analysis for DNS attacks. We found serious logic flaws in advanced attacks mechanisms. We discussed the weaknesses in the DNS systems and ways of its protection.

Much of the Internet’s DNS infrastructure remains open and unprotected—characterized by a lack of dedicated security personnel, poor traffic visibility and unrestricted access to DNS recursors. Yet security threats against DNS infrastructure are serious—and growing.

We believe that our study takes some steps in the security problem space that DNS infrastructure has brought. We believe that our study brings some new chain of trust between the client and the authoritative server in DNS security. In future work we are considering the security challenges that come with other advanced DNS attacks. Fundamentally, we believe that vulnerabilities of DNS demands new research efforts on ensuring the security quality of the systems.

4. REFERENCES

[1] ”Denial of Service Attack via ping”. Internet: http://www.cert.org/advisories/CA-1996-26.html [Dec, 1996].

[2] Sun Changhua, Liu Bin, Shi Lei. “Efficient and low-cost hardware defense against DNS amplification attacks”. IEEE Global Telecommunications Conference, GLOBECOM 2008 [May, 2008].

[3] Li M, Li J, Zhao W. “Simulation Study of Flood Attacking of DDOS”, Icicse: International Conference on Internet Computing in Science and Engineering, Proceedings [June, 2008].

[4] Guo Fanglu, Chen Jiawu, Chiueh Tzi-Cker, Spoof detection for preventing DoS attacks against DNS servers, 26th IEEE International Conference on Distributed Computing Systems, ICDCS [Feb, 2006].

[5] Kambourakis G., Moschos T., Geneiatakis D., Gritzalis S, Detecting DNS Amplification Attacks, Critical Information Infrastructures Security, v(5141), pp. 185 – 196.

[6] Bau J., Mitchell J., A security evaluation of DNSSEC with NSEC3, Citeseer [May, 2010].

[7] Li Wei-min, Chen Lu-ying, Lei Zhen-ming, Alleviating the impact of DNS DDoS attacks , Proceedings of the 2010 2nd International Conference on Networks Security, Wireless Communications and Trusted Computing (NSWCTC 2010), pp. 240-243 [Dec, 2010].

[8] Scalzo F, Recent DNS Reflector Attacks Verisign. Internet: http://www.nanog.org/mtg-0606/pdf/frank-scalzo.pdf [Dec, 2006].

[9] Sen J, A Robust Mechanism for Defending Distributed Denial OF Service Attacks on Web Servers, Arxiv preprint arXiv: 1103.3333 [Jul, 2011].

[10] Dittrich D, Distributed Denial of Service (DDoS) Attacks/tools. Internet: http://staff.washington.edu/dittrich/misc/ddos [Oct, 2012].

[11] The Measurement Factory, Domain name servers: pervasive and critical, yet often overlooked, The Measurement Factory DNS Survey. Internet: http://dns.measurement-factory.com/surveys/sum1.html [Nov, 2005].

[12] Singh A, Singh B, Joseph H, Vulnerability Analysis for DNS and DHCP, Vulnerability Analysis and Defense for the Internet, pp. 111-124 [Dec, 2008].

[13] Beverly R and Bauer S, The spoofer project: inferring the extent of source address filtering on the Internet, USENIX workshop on Steps to Reducing Unwanted Traffic on the Internet, 2005.272 X. Ye et al. /Journal of Computational Information Systems 9, pp. 265–272 [May, 2013].

[14] V. Paxson, “An analysis of using reflectors for distributed denial-of-service attacks, ” SIGCOMM Comput. Commun.Rev., vol. 31, no. 3, pp. 38 – 47 [May, 2011].

[15] K. Rikitake, “A Study of DNS Transport Protocol for Improving the Reliability, ” Ph.D. dissertation, Graduate School of Information Science and Technology, Osaka University [Oct, 2005].

[16] M. de Vivo, G. O. de Vivo, R. Koeneke, and G. Isern, “Internet vulnerabilities related to TCP/IP and T/TCP, ” SIGCOMM Comput. Commun. Rev., vol. 29, no. 1, pp. 81 – 85 [Dec, 1999].

[17] V. Ramasubramanian and E. G. Sirer, “The design and implementation of a next generation name service for the internet, ” SIGCOMM Comput. Commun. Rev., vol. 34, no. 4, pp. 331 – 342 [Feb, 2004].

[18] H. Yang, H. Luo, Y. Yang, S. Lu, and L. Zhang, “HOURS: Achieving DoS Resilience in an Open Service Hierarchy, ” in Proc. IEEE DSN04 [March, 2004].

[19] ICANN SSAC, SSAC Advisory SAC008 DNS Distributed Denial of Service (DDoS) Attacks. Internet: http://www.icann.org/committees/security/dns-ddos-advisory-31mar06.pdf [Feb, 2006].

[20] Huiming Yu, Xiangfeng Dai, Baxliey T, Xiaohong Yuan, Bassett T, A Visualization Analysis Tool for DNS Amplification Attack, Proceedings of the 2010 3rd International Conference on Biomedical Engineering and Informatics (BMEI 2010) [May, 2010].

[21] IPTraf - An IP Network Monitor. Internet: http://iptraf.seul.org/ [Jan, 2014].

[22] S. Murdoch and R. Anderson. “Verified by Visa and MasterCard SecureCode: or, How Not to Design Authentication”. Financial Cryptography and Data Security, pp. 42-45 [Jan, 2010].

[23] “SSL: Intercepted today, decrypted tomorrow”. Netcraft, pp. 10-12 [May, 2013].