AKAMAI WHITE PAPER

Is DNS Your Security Achilles Heel?

Akamai

Executive Summary

With the constant drumbeat of news reports about security breaches, cyber security is hard to ignore. Organizations understand that they need comprehensive solutions that prevent, detect, and respond to security threats. They often implement multiple layers of security controls to protect their IT systems.

Yet gaps remain. Many organizations have a blind spot when it comes to the Domain Name System (DNS). Although every action on the Internet relies on recursive DNS, many security organizations fail to install corresponding safeguards. Cybercriminals have been only too happy to exploit this security vulnerability.

This white paper explains how attackers take advantage of recursive DNS and provides best practices you can follow to mitigate the risks.

Cyber Attacks Proliferate

In 2016, 3,046,456 data records were lost or stolen every day; 126,936 every hour; 2,116 every minute. During that same period, 32% of companies admitted to being victims of cybercrime. The cost of these crimes is skyrocketing; the global price of cybercrime is expected to increase to more than $2 trillion by 2019.[1]

Defending against cyber attacks has always been a cat and mouse game between IT security and cybercriminals. Organizations understand that there are many attack vectors and typically use a layered security, or defense-in-depth, approach to eliminate security vulnerabilities.

With a layered approach, companies deploy multiple security controls to prevent, detect, and respond to a variety of cyber threats. In addition, multiple layers of software controls offset weaknesses in one security layer with strengths in another. These security measures include everything from user authentication and access policies, to anti-virus software, to proactive monitoring that uses machine learning to detect and defend against emerging types of attacks.

But even as information security personnel become more sophisticated about closing security gaps, cyber attackers continue to look for — and identify — vulnerabilities that they can exploit.

One area that increasingly draws the attention of cybercriminals is the Domain Name System (DNS). Recursive DNS is the perfect cyber attack vector because it is ubiquitous, open, often unprotected, and absolutely vital to the Internet's existence. Not surprisingly, cybercriminals have evolved their targeted threats to leverage this security vulnerability, and targeted threats that take advantage of this vector are growing in number.

What is DNS?

Any time that users or network-connected devices (including IoT devices) perform an Internet request — from web browsing to email to online retail to cloud computing — they use DNS. According to the Internet Corporation for Assigned Names and Numbers (ICANN), 30 to 50 million DNS servers exist.

These Internet requests literally would not work without DNS. DNS maps domain names that humans can understand (e.g.,www.companydomain.com) to machine-readable IP addresses (66.94.234.13). By using DNS, end users can connect to a website like www.mywebsite.com without having to know the IP address of the server where the website is hosted.

DNS relies on two types of servers:

- **Recursive DNS servers:** When an end user wants to find an IP address, this server looks for it.
- **Authoritative DNS servers:** The "yellow pages" of IP addresses, these give recursive DNS servers answers about the IP addresses mapping.

Cybercriminals promulgate many types of attacks that exploit both recursive and authoritative DNS servers. This white paper focuses on attacks that take advantage of recursive DNS servers to launch targeted threats, connect an infected endpoint with a command and control (CnC) server controlled by a cybercriminal, exfiltrate data, or spread malware to other endpoints on the network. For more information, see the appendix "How Does DNS Work?".

Why is Recursive DNS Prone to Exploitation?

The primary reason recursive DNS attacks are flourishing is that the DNS-based Internet infrastructure is largely unprotected — and attackers know and exploit that fact.

DNS is simply designed to resolve requests. It has no way to evaluate whether a website or other resource to which it connects the user or device is good or bad.

Moreover, firewalls often do not inspect TCP port 53, which DNS servers use to listen for queries from DNS clients. Indeed, in order to enable unfettered DNS communications, many firewalls between the Internet and a local network come with their defaults configured to accept all traffic sent to port 53.

Why are These Attacks Flourishing?

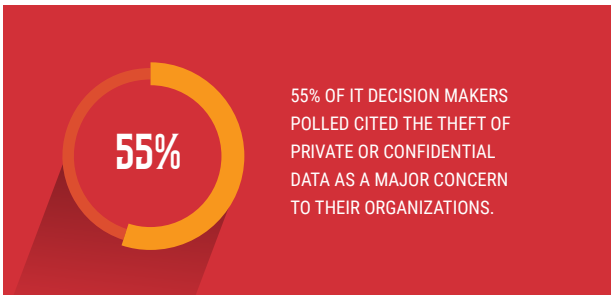

Despite these gaping security holes, IT teams persist in thinking that DNS is benign. Few CIOs make protecting DNS communications a priority. 44% of the respondents to a Vanson Bourne study found it difficult to justify investment in DNS security because senior management did not consider it to be an issue, despite the fact that 55% of IT decision makers polled cited the theft of private or confidential data as a major concern for their organizations.[2]

Because IT organizations don't implement security solutions that can automate the process of identifying malicious locations, the only way to monitor DNS ports is to manually review DNS server logs. Manual monitoring is not only incredibly time consuming due to the volume of DNS traffic, but also inefficient since the overwhelming majority of content is good. Regardless, it is difficult to spot problems or identify trends that indicate irregularities because companies only have the limited scope of their own DNS traffic to review.

Without protection, most organizations have no idea what types of queries are going over their DNS infrastructure. As a result, the typical DNS resolution process doesn't prevent users or connected devices from making requests to malicious domains.

How do Targeted Threats Exploit Recursive DNS?

The typical recursive DNS attack follows a well-worn pattern. Malware finds its way onto an end user machine or connected device using a variety of tactics, often using DNS for the download phase. Once on the device, the malware uses DNS to connect with its CnC infrastructure to download additional software or send back intelligence about the device on which it is installed. Because DNS has no way to determine whether the destination of the request is a malicious or safe domain, it simply resolves requests. Hackers can also exploit DNS to obfuscate these malicious CnC servers through techniques such as a Domain Generation Algorithm (DGA) and fast fluxing. Next, the CnC server downloads updates to the malware or additional components onto the compromised machine, exfiltrates confidential data out of the company, or even installs ransomware.

Targeted Threat Delivery

By far the most common means of infection is through generic phishing or targeted spear phishing emails. Phishing tricks the victim into opening an attachment or linking to a malicious website. Spear phishing is a more targeted form of phishing where emails are designed to appear as though they were sent by someone the recipient knows and trusts; they may include a subject line or content tailored to the victim's interests or industry.

End users' machines can also acquire malware when a user plugs in an infected USB key, browses to a malicious website directly, or clicks on a malicious ad or link on a page where they are already browsing. When the malware is dropped on the target computer and executed, it will use DNS to talk to a communications channel between the now infected computer and one or more CnC servers operated by the bad actors.

Command and Control (CnC)

Once the communication channel is established, the CnC server can retrieve information about what the malware finds on the device, such as unpatched OS vulnerabilities. It can then instruct the malware to download remote access tools or additional software components or updates that exploit unpatched vulnerabilities on the compromised device.

Now the CnC server is in a position to issue commands to the compromised systems on the target network. These communications can be as simple as maintaining a timed beacon or "heartbeat" so that the operators running the attack can keep an inventory of the systems they have compromised within the target network. Or, attackers can use the CnC communication channels for more malicious actions, such as remotely controlling the machine or exfiltrating data.

Once a machine on the network is compromised, the attacker will often attempt to move laterally within the target network to infect additional hosts. This lateral movement ensures that if one infected system is detected, the attacker continues to maintain access. These newly infected systems also begin to call home by sending beacons to the CnC servers. This pattern continues as the attackers perform reconnaissance on the targeted systems, establishing a shadow network within the enterprise infrastructure. Infected computers can also become part of a bot network of zombie devices that participate in attacks unbeknownst to the machine's owner.

Using a DGA (Domain Generation Algorithm) to Avoid Detection

Malware that depends on a fixed domain or IP address for CnC communications will likely be quickly detected and blocked either by the company or by a service provider. Network detection techniques can easily discover a predefined list of domain names. The company can then simply block the traffic in their firewalls or contact their Internet or cloud-service provider to instruct them to log or shut down the malicious behavior. When these domains are blocked or the server is taken down, it cuts the links between the infected device and the CnC servers.

If the links are cut, the hacker might have to go through the hassle of bringing out a new version of the malware or setting everything up again on a new server. But some families of malware take advantage of an easier alternative. They use a DGA to create new domains on-demand or on-the-fly to evade detection by blacklists, signature filters, reputation systems, intrusion prevention systems, security gateways, and other security methods.

DGAs generate a large number of domain names that can be used as rendezvous points with their CnC servers. This tactic is difficult to mitigate because someone would have to reverse engineer the DGA algorithm used to generate domain names to block them with a firewall blacklist. In addition, DGAs can regenerate domain names frequently using different keys or algorithms. Attackers can simply set up the CnC server briefly to allow infected machines to call home — then shut it down and set it up again as the need arises. This ensures that DGA domains do not live long enough to become known by blacklists.

Fast Fluxing

DNS fast flux is another way that malware delivery sites can hide domains used to download malware, host phishing websites, or communicate with CnC servers.

Fast flux allows an attacker to use constantly changing servers to provide a layer of redundancy in front of their actual CnC server, as well as to defeat IP-based firewall blocks. The DNS infrastructure allows an administrator to register multiple IP addresses to a single host name, and attackers corrupt this utility. Fast flux can be used to hide phishing and malware delivery sites behind an ever-changing network of compromised hosts acting as proxies. Essentially, the domain names and URLs for content no longer resolve to the IP address of a specific server. Instead, they fluctuate among many front-end redirectors or proxies, which in turn forward content to backend servers that serve malicious content.

Use of the proxy server stymies historically effective defense mechanisms — e.g., IP-based access control lists. This method also masks the attackers' systems, allowing cybercriminals to exploit the network through a series of proxies and making it much more difficult for IT security professionals to identify malicious networks.

Data Exfiltration

Once a machine is compromised, the CnC server can use DNS port 53 to exfiltrate data in a process called DNS tunneling.

In tunneling, malicious insiders or outside hackers use the DNS protocol as a pathway to get commands into the network and extract data. When stealing information, the criminals parse data into chunks sized to fit into a DNS query. They then send the queries to a specially modified "rogue" authoritative DNS server that the hacker has set up in advance and controls remotely. This authoritative name server is accessible from the Internet, as long as port 53 is passed to it. This server receives, unencodes, and reassembles the stolen information.

Data exfiltration may steal:

- Personally identifiable information (PII) such as social security numbers
- Regulated data related to Payment Card Industry Data Security Standard (PCI DSS) and HIPAA compliance
- Intellectual property that gives the organization a competitive advantage
- Other sensitive information such as credit card numbers, company financials, payroll information, email addresses, phone numbers, and so on

Best Practices to Prevent DNS Exploitation

How can organizations protect themselves from attacks that exploit DNS vulnerabilities?

To safeguard against targeted threats that take advantage of DNS to exfiltrate data or escalate horizontally, organizations can direct the company's external traffic to a service that checks requested domains against real-time risk scoring threat intelligence. Such a service blocks employees from accessing malicious domains and services, such as CnC servers. It validates traffic before the IP connection is made, to stop threats early in the kill chain.

But installing these protections should be done in conjunction with adopting best practices around DNS.

The following best practices prevent DNS exploitation to ensure that enterprise threat protection services properly identify, block, and mitigate threats, and enforce acceptable use policies across the organization.

1. Lock Down Employee Systems to Prevent Employees from Changing Local DNS Settings

Organizations adopt enterprise threat protection systems to block requests to known malware sites. When users type in a domain name, the request goes to the enterprise threat protection system and blocks any requests to known malware sites.

But end users may want to circumvent the threat protection and go to the site anyway. For example, they might want to view photos on the site. End users can easily reach any restricted site by simply using a free resolver (such as Google or other companies) to bypass the DNS settings on their local computer.

Organizations should therefore lock down employee computers to prevent employees from changing local DNS settings. One way to accomplish this is to create a group policy in the Active Directory that prevents users from changing the DNS settings.

2. Lock Down Employee Systems to Prevent Employees from Installing Third-party VPNs That Can Bypass the Enterprise Threat Protection System

When organizations secure their communications over the Internet, they have the option to use two types of VPN tunnels:

- **Split VPN tunnels:** With split VPN tunnels, the DNS query hits the local system and the local system makes an outbound query that adheres to company policies set by the administrator. The query might be directed to an ISP or to the local DNS server in the enterprise.
- **No-split VPN tunnels:** With this type of tunnel, the DNS queries go over the tunnels themselves. If an employee were to download a free or paid third-party, no-split VPN service, he or she could bypass the enterprise threat protection filters installed by the organization.

Best practice is to lock down employee systems to prevent employees from installing third-party VPNs that can bypass enterprise threat protection or network firewalls. Such a lock-down denies remote VPN access from any employee workstations unless they meet business requirements.

3. Lock Down the Perimeter Firewall to Allow DNS Queries to Come Only from Local DNS Systems

Best practice is to use a firewall to deny outbound traffic on DNS port 53 unless it's from a trusted source (e.g., the local DNS server) and bound to a trusted destination (e.g., the resolver virtual IP [VIP] that the enterprise threat protection system assigns to you). That local DNS server uses the enterprise threat protection system to look up all local queries for the enterprise to ensure that they only go to specified, legitimate addresses.

4. Look for Suspicious Patterns in Your DNS Logs

Suspicious patterns in DNS logs indicate the presence of malware. Malware patterns can include:

- **Queries outside of normal business hours:** If a query comes from an employee's laptop at 2 a.m., chances are it is not the employee working late. Instead, it is a suspicious query.
- **Queries that use non-standard naming conventions:** Bad actors use domain names with non-standard naming conventions because they know that these domain names will not be taken or registered. They can register the names on the fly. So, look at the naming conventions of the domains and try to rationalize them.
- **Long-tail log queries:** To locate queries at the tail of the log, first sort all the domain names alphabetically. Eliminate any duplicates to ensure that all of the domain names are unique. Count the requests to each domain name. Look at the tail end of the log and you will find a number of domain names that are accessed only once or twice. You will find that these are also often accessed outside of normal business hours and use non-standard naming conventions. These are the domain names you should investigate to determine who owns the company and how long it was registered to determine whether the domain is legitimate. For example, if it was only registered 10 hours ago and accessed 30 minutes later, that is not normal.

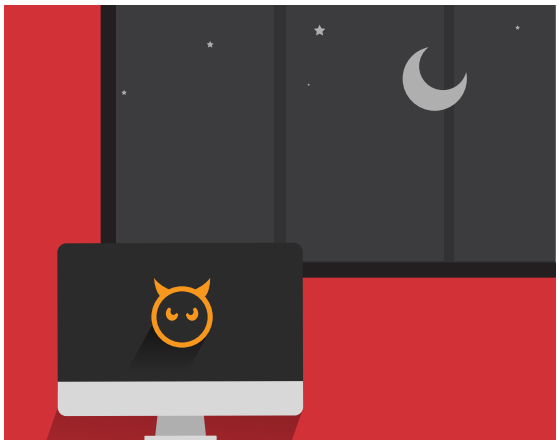

5. Separate Your User Traffic from Data Center Traffic

Separate your outbound traffic from your data center traffic, and have them egress from different points in your network. This allows you to set much stricter controls over traffic leaving your data center. In addition, look at the DNS traffic exiting your data center to make sure that it is not going to a suspicious domain.

6. Beware of Tor

The Tor network is a group of servers operated by volunteers that allow organizations and individuals to improve the privacy and security of information on the Internet. The Tor network has many legitimate uses. But there is no good reason to use Tor on a corporate network. If a connection is observed going to a Tor entry node, it's likely bad. Many targeted threats use Tor to communicate with their CnC servers. As a result, it is good practice to block Tor entry nodes at the edge firewall. Do this either by blocking all Tor entry nodes (based on publicly available lists of entry nodes) or doing deep packet inspection (DPI) on HTTPS traffic in your firewall.

Conclusion

Cyber attacks are likely to remain a significant threat in the years to come. In the security realm, your best defense remains a layered defense. But in today's environment, where cybercriminals are constantly on the lookout for vulnerabilities, a layered defense is only as secure as its weakest link.

Don't let unprotected recursive DNS be the Achilles heel in your layered defense strategy. Take the threat from cyber attacks on recursive DNS seriously and protect this very real vulnerability.

For more information on how to stop recursive DNS attacks, read the solution brief **"Targeted Threat Protection in the Cloud."**

APPENDIX

How Does DNS Work?

Machines on the Internet use IP addresses to communicate. These addresses consist of strings of numbers. But humans do not want to keep track of many complex numbers whenever they want to perform a simple task. Enter DNS, which is used like a phone book to simplify the process of finding the numerical IP address that corresponds with each human-readable domain name.

How are Domain Names Structured?

DNS uses a hierarchical system to resolve the domain name that a user enters and communicate the corresponding IP address back to that user.

The domain name consists of several parts:

- **Root -** The root is represented by a "." at the end of a URL, which is not always shown.
- **Top-level domains -** Top-level domains are classified into two subcategories: organizational (.com, .edu, .gov, .mil, .net, .org, .int) and geographic (.uk).
- **Second-level domains -** This is the main part of the domain name.
- **Sub-domains -** These are subsets of the main domain.

To allow computers to recognize each part properly, dots are placed between the parts of the domain name. The hierarchical tree is searched from right to left, starting with a root domain name server.

How Does DNS Find the Right IP Address?

The DNS resolution process begins when you ask your computer to access a hostname. Say you want to visit www.companydomain.com.

Your computer starts by looking for the IP address in its local DNS cache, which stores information that your computer has recently retrieved. If the address is in the cache, the host name is resolved and the information is passed to your browser, which opens the connection to the requested web server.

If your computer does not find the answer in the cache, it launches a DNS query to find it. Your computer asks a recursive DNS server to perform a DNS query on your behalf to find the proper IP address of the domain name you want to access.

Recursive servers have their own caches, so the process usually ends here, and the information is returned to your computer.

If the information is not in the cache, recursive DNS servers turn to name servers and authoritative servers. Name servers and authoritative name servers act like phone books for each domain in the hierarchy — root, top-level, second-level, and sub-domains.

The recursive server starts by querying the root name server. There are 13 root name servers on the Internet. If the root name server doesn't know the answer, it directs the query to a top-level domain (TLD) name server that knows where to find it. The TLD server in turn refers the query to the authoritative server for the specific second-level domain or sub-domain.

An authoritative name server is responsible for knowing all the information about a specific domain or sub-domain. It stores this information in DNS records and provides the recursive DNS servers with the IP mapping of the intended website. An authoritative server will have all the information about the domain it is responsible for, or referral information for zones within the domain that have been delegated to other name servers.

The recursive server retrieves the record from the authoritative name servers, stores the record in its local cache, and then returns the record back to your computer. Your computer stores the record in its own cache, reads the IP address from the record, and passes the information to your browser. The browser opens the connection to the web server and receives the website. The entire process takes only milliseconds to complete.

Sources

1) https://www.cyberark.com/blog/noteworthy-cyber-security-statistics/

2) http://www.securityweek.com/nearly-50-percent-organizations-hit-dns-attack-last-12-months-survey