Stop Attackers from Using DNS Against You

Any modern organization requires the Domain Name System (DNS) to run its business, regardless of industry, location, size, or products. DNS is a protocol that translates user-friendly domain names, such as www.paloaltonetworks.com, into machine-usable IP addresses—in this case, 199.167.52.137. Without DNS, we'd have to memorize random strings of numbers, which our brains aren't well equipped to do. DNS is fundamental to every single modern organization, all over the world. Network operators cannot block DNS traffic, and firewalls have to let it through. Networks need DNS to function properly.

Many security professionals don't realize the ease and prevalence of DNS abuse by attackers. In fact, many security teams don't inspect DNS traffic for threats because they assume queries sent over DNS protocol and port 53 are benign. Other organizations don't inspect DNS traffic because the sheer volume of that traffic is overwhelming, and looking for a sign of something malicious in that traffic is like looking for a needle in a haystack. This takes a great deal of time and resources—often too great an investment for organizations, especially those that assume DNS does not pose a significant threat.

DNS is a massive and often overlooked attack surface that can be used for malware delivery, command and control (C2), or data exfiltration. Adversaries take advantage of the ubiquitous nature of DNS to abuse it at multiple points of an attack. According to Palo Alto Networks Unit 42 threat research team, almost 80% of malware uses DNS to initiate C2 procedures. Attackers establish reliable command channels that are difficult to take down or identify since DNS is such a reliable way to maintain a connection to DNS servers. As adversaries increasingly automate their attacks, it becomes almost impossible to identify and stop those threats.

80%

of malware uses DNS to initiate C2 procedures that can be used to steal data and spread malware

Figure 1: Unit 42 research on DNS traffic

At the same time, many security teams lack visibility into DNS traffic and how threats use DNS to maintain control of infected devices. Security teams are under pressure to enforce consistent protections for millions of new malicious domains while keeping up with advanced tactics like DNS tunneling. In addition to how prevalent and easy DNS abuse is, the sheer rate and volume of new malicious domains is enormous, and static signatures cannot be created quickly enough to keep up. If a system gets infected, networking and security teams are challenged to quickly identify that system and address the infection. By then, malware may have already spread, or data may have already been stolen.

Top Three Attacks Using DNS

Understanding how adversaries abuse DNS is the first step to stopping attacks on your network and minimizing your cybersecurity risks. Here are the top three ways cybercriminals abuse DNS to mask their C2 activity so they can deliver additional malware or steal data.

Malware Using DNS for C2

This is one of the most typical ways attackers take advantage of DNS. Attackers use common network protocols, including DNS, to spread malicious code. Malware can be sent to users through online ads, malicious URLs in emails, or other means. Once a user's computer is infected, the system sends a DNS request back to the attacker's control server. In this way, the infected computer becomes a bot the attacker can control. The malware can then steal personal or financial data and spread very quickly by issuing instructions to scan the network for other computers.

Recently, the hacker group WINDSHIFT launched a cyberattack that used DNS for C2 against government departments and critical infrastructure in the Middle East. To learn the technical details and timeline, read Unit 42's research on the WINDSHIFT attacks.

Malware Using Domain Generation Algorithms

Effective and growing quickly, domain generation algorithms (DGAs) randomly generate large numbers of slightly different domain names. A DGA can, for instance, create thousands of domains in a day that are each a slight variation of www.bigbadguys.com. Attackers developed DGAs so that malware can generate these domains and use them for C2. Unit 42 has observed that 18% of malware uses DGAs to automatically create thousands of C2 domains every day—of which attackers may use one—so that defenders can't block them. Malicious domains controlled by attackers enable rapid movement of C2 channels from point to point, bypassing traditional security controls like blacklists or web reputation filtering. Infected computers contact some of these new domain names to receive commands and updates. A key aspect of DGAs is that, even though thousands of domains can be generated in short order, not all of them need to be registered.

DGAs offer an effective means for attackers to hide the locations of their C2 centers, which they use for financial fraud, identity theft, and other malicious activities. To learn more about DGAs, read Unit 42's DGA threat brief.

Attacks using DNS are effective and growing. Malware using domain generation algorithms (DGA) has grown

124% year over year

Figure 2: Unit 42 research on domain generation algorithms

DNS Tunneling

This technique, increasingly used by advanced persistent threat (APT) actors, lets attackers encode their payloads in small chunks within DNS requests to bypass security controls. Advanced attackers use DNS tunneling to hide data theft or C2 in standard DNS traffic. Once a victim's device is compromised, the infected device sends a request within the DNS traffic. The DNS server is instructed to connect to the cybercriminals' server, establishing a channel through which to steal and transmit data. With DNS tunneling, DNS requests pass through the normal DNS server, inside and outside a company's firewall. However, tunneled data hidden in the DNS requests goes unnoticed. Attackers including the threat group OilRig have used DNS tunneling extensively in recent years.

Why Current Security Approaches Fail

Current approaches to blocking malware attacks that use DNS are inadequate for several reasons. To begin with, it is difficult to address the many ways attackers can use DNS to compromise an organization. Many organizations focus solely on protecting their DNS infrastructure—and rightfully so. If DNS goes down, they can no longer access the internet. What they don't focus on is the hidden threat: attackers using DNS itself to spread malware or steal data. Some organizations do nothing to protect DNS, leaving it wide open for attackers. Many organizations don't have DNS monitoring and instead only block malicious domains, essentially doing nothing to address malware that abuses DNS.

Other security teams take a blacklisting approach to blocking attacks that use DNS, relying on relatively static threat feeds that work off known bad domains. However, as malware's use of DGA grows, the effectiveness of blocking known malicious domains alone becomes more limited. Using a list of randomly generated domains for C2 can overwhelm the signature capability of legacy tools and traditional security approaches. A limited set of signatures simply cannot scale to meet the growing threat of DNS-based attacks.

Additionally, relying on static lists limits the amount of context defenders can access to fully understand the attacks against their network. Although threat intelligence feeds are regularly updated with indicators or artifacts derived from a source outside the organization, daily or even hourly updates are too slow to keep up with the massive amount of DNS data. The sheer volume of DNS traffic often means defenders simply lack the visibility or resources to universally inspect it for threats. With a traditional approach, security teams don't have the resources to be proactive or scale their DNS security.

Some organizations use standalone point products to address threats to their DNS. These tools may adequately address specific facets of DNS security, but even "best-in-breed" technologies come with limitations. For instance, these tools often require changes to DNS infrastructure if they are to work effectively. Disparate products also create siloes of threat intelligence and data that may not work with other areas of an organization's security structure. As a result, overwhelmed teams drown in uncoordinated data from independent tools. Multiple tools become more things to own and manage, adding complexity and drain on already limited human resources.

Unit 42 Threat Research on OilRig

OilRig is an active, organized threat group first discovered by Unit 42. Operating primarily in the Middle East, OilRig carefully targets organizations to further its regional strategic goals across multiple industries, including supply chain-based attacks. As part of its adversary playbook, the group employs sophisticated, custom DNS tunneling for C2 and data exfiltration. The use of tunneling includes:

- **ALMA Communicator Trojan**, which uses DNS tunneling to receive commands from the adversary and exfiltrate data. The malware employs specially crafted subdomains to send data to the C2 server and specific IPv4 addresses to transmit data from the C2 to the Trojan over DNS requests.
- **Helminth PowerShell-based Trojan**, which can obtain files from a C2 server using a series of DNS TXT queries repeated every 50 milliseconds, essentially building malware on victim systems through hard-to-detect increments sent over DNS.

OilRig's use of DNS tunneling allows the group to establish reliable C2 that can potentially evade existing defenses to carry out further stages of the attack. Get the full details on OilRig from Unit 42's blog post series or the interactive Playbook Viewer.

Stop Attackers from Using DNS Against You

How can you regain control of your DNS traffic and prevent attackers from using DNS to attack your organization?

Security Data, and Lots of It

You need massive quantities of real-world security data, either that you collect yourself or gather through threat intelligence or cyberthreat alliances. With data from a large and expanding intelligence-sharing community, your protection will continue to grow.

Analytics and Machine Learning

Your security teams need to be able to run analytics on that data. To address the dynamic nature of domains or DNS tunneling, your teams must employ machine learning to dynamically identify unknown bad domains. Without analytics, it is impossible to predict highly dynamic malicious domains. Behavioral analytics can also help determine a baseline of activity, understand general patters, and find what is normal. When

defenders see signals that require action, analytics can help determine how manual or automated that action should be. Analytics can also understand which signals need to be acted upon, helping your teams prioritize time and resources.

Integration with a Next-Gen Firewall for Automated Action

Because many DNS-based attacks happen so quickly, it is imperative that security teams spend less time manually responding to attacks. To stand a chance, defenders need automation. Automation can help quickly determine infected machines, automate responses, and contain threats before they spread to other areas of a network. Security teams need integrated innovations that extend the value of existing security investments without complicating operations.

Cloud-Based Protection

Using the cloud, your DNS protections can scale infinitely and always stay up to date, giving you a critical new control point from which to stop attacks that use DNS. Cloud-based innovations enable your defenders to develop and deploy new detection techniques that your organization can take advantage of instantly. Cloud-based protections update instantly without requiring you to update or make changes to software, which means less work for your security operations center (SOC) teams.

Avoid Standalone Point Products

Finally, your security teams must avoid deploying disparate tools that are poorly integrated or require changes to DNS routing. Many of these tools weren't designed for automation, forcing your analysts to manually stitch together insights from multiple disparate sources before acting. These products also don't automatically share data or insights, and they won't let you coordinate alerts across your entire security stack. As a result, your teams can't approach protection holistically, resulting in slower responses to threats.

DNS Security Best Practices

In addition to deploying the right technology, there are other best practices your organization can follow to protect your network from DNS-based threats.

Train Your Staff to Be Security Aware

Implement a security education and awareness program to train your staff on what to look for in suspicious emails. Encourage them to take care when following links to avoid installing malware. Phishing training can help them learn to recognize, avoid, and report email-based attacks.

Implement a Threat Intel Program

Understand the threat landscape and set up a threat intelligence program to understand what threats and techniques exist. With this knowledge, you can ensure you have the right technology stack to keep your network safe.

Learn What the Logs Can Tell You

Don't just look at DNS traffic. Collecting DNS logs has little value unless you understand what you're looking at, what the data is telling you, and what you can do to secure your network from DNS-based attacks.

Don't Blindly Rely on a DNS Resolver

If a DNS server is compromised, it may feed you false responses meant to direct your traffic to other compromised systems or enable a man-in-the-middle attack.

Plan for Mobile Employee Risk

Develop a strategy for your mobile employees as they can put company data at risk. Warn them against using unsecured, free, or public Wi-Fi as adversaries can easily put themselves between employees and the connection point. Integrate multi-factor authentication. Assume a high risk of devices being lost or stolen, and have a plan in place.

Approach Network Security Holistically

Don't rely on a single product that promises to solve all your security problems. Instead, take a holistic approach to network security and ensure you have all the right tools to combat modern threats. Look at your security tools' capabilities and whether you can use them together effectively. You need tools with multiple capabilities that address various threat vectors, including intrusion prevention, URL filtering, and file blocking.

Automate Response, Not Just Alerts

Require automated response, not just signals. Threats move so fast that alerts or signals alone are ultimately not helpful. By the time an analyst has prioritized an alert, confirmed a threat, and identified the threat and its source, it may already be too late. Your security systems must be able to automatically determine threats and quarantine potentially infected systems before more damage is done.

Is your organization implementing best practices in your DNS security strategy? Take a Best Practice Assessment to be sure.

paloalto NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com