DARKReading

MARCH 2022

Understanding DNS Threats and How to Use DNS to Expand Your Cybersecurity Arsenal

The goal of proactive DNS-layer security — such as using DNS data to mitigate threats that network-based controls may miss — is to block threats before they hit the enterprise network or endpoints.

Brought to you by informa tech

INSIDE:

Next

FEATURE

Understanding DNS Threats and How to Use DNS to Expand Your Cybersecurity Arsenal

The goal of proactive DNS-layer security — such as using DNS data to mitigate threats that network-based controls may miss — is to block threats before they hit the enterprise network or endpoints.

By Jai Vijayan, Contributing Writer, Dark Reading

With threat actors increasingly using DNS (that is, the Domain Name System) for malware command-and-control (C2), data exfiltration, and Web traffic redirection, DNS-layer security controls could soon become a core component of enterprise security strategies.

Phishing, malware, and ransomware attacks routinely exploit techniques like DNS beaconing to communicate with C2 servers and DNS tunneling to covertly deliver payloads and exfiltrate data from enterprise networks. To hide phishing and malware hosting sites, threat actors often use domain-generation algorithms and DNS fast flux (rapidly swapping out IP addresses associated with a domain) to churn out new domain names faster than static enterprise blocking tools can keep up.

A survey that analyst firm IDC conducted last year found 87% of organizations had experienced a DNS attack — mainly phishing, malware, and distributed denial-of-service (DDoS) — over the previous 12 months. Seventy-six percent of respondents viewed DNS security as critical to their threat mitigation efforts.

"As enterprises become more Internet-facing and cloud becomes the dominant form of computing, DNS is the front line for accessing data and information," says John Yeoh, global vice president of research at the Cloud Security Alliance (CSA).

"The openness of DNS and the Internet have made them commonly exploitable."

The basic goal of proactive DNS-layer security is to identify and block threats before they even hit the enterprise network or endpoints. The effort goes beyond just protecting the DNS infrastructure from attacks like DDoS, DNS hijacking, and cache poisoning. It includes using DNS data to mitigate threats that traditional network-based controls alone can't catch any longer.

The Promise of Protective DNS (PDNS)

Most approaches to leverage DNS for proactively protecting organizations against phishing, malware, and other threats are focused on recursive DNS, or the service that resolves IP addresses when Internet requests are made. Recursive DNS lookups are what happens when a user wants to navigate to a site and the DNS server either resolves the domain name on its own — if it has the information in its cache — or contacts another (authoritative) DNS server for the IP address. For example, when a user wants to navigate to *www dot darkreading dot com*, it's the organization's recursive DNS resolver that finds and returns an IP address for the requested URL. An organization might have its ISP handling DNS services or, as is often the case with large enterprises, it might have an internally operated DNS or designated DNS service managed by a third party.

Because recursive resolvers handle DNS queries and fetch requested IP addresses, many view them as the perfect spot for detecting and blocking connections to sites and infrastructure that are known to be malicious or suspicious in some way. For instance, when a user clicks on a malicious link in a phishing email, or mistypes the URL to a particular site, the recursive DNS resolver has information on both the client that made the query and the destination IP address. The idea is that if the IP address is known to be malicious — by checking it against threat intelligence on IPs associated with phishing, ransomware, and C2 activity, for example — policies can be enforced for preventing that connection being made. This ensures that threats are blocked before they even hit the enterprise network or endpoint devices.

Because recursive resolvers handle DNS queries and fetch requested IP addresses, many view them as the perfect spot for detecting and blocking connections to sites and infrastructure that are known to be malicious or suspicious in some way.

"Everything that happens in a network is usually visible in DNS, and some of the most simple and successful threat-hunting techniques revolve around DNS," says Johannes Ullrich, dean of research at the SANS Technology Institute. "In a simple, well-managed corporate network, internal recursive DNS servers provide all name resolution, and query logs collected from these servers can be used for analysis" and threat blocking.

Services that analyze DNS queries and block access to known bad sites and infrastructure are generally referred to as protective DNS (PDNS). Organizations can implement the capability in-house — by integrating policy-based DNS filtering and blocking with a network firewall, for instance. Or they can tap a growing number of vendors that offer PDNS services.

With PDNS services, when a user wants to navigate to a site, the DNS provider's recursive resolver analyzes the request, acquires the requested IP address, and uses threat intelligence to determine if the requested connection is safe or not. Many PDNS services tap a combination of open, commercial, and government threat intelligence feeds to aggregate — and constantly update — information on malicious activity and the domains and infrastructure associated with that activity. If a query is deemed unsafe, the PDNS service then would apply customized policies contained in what is known

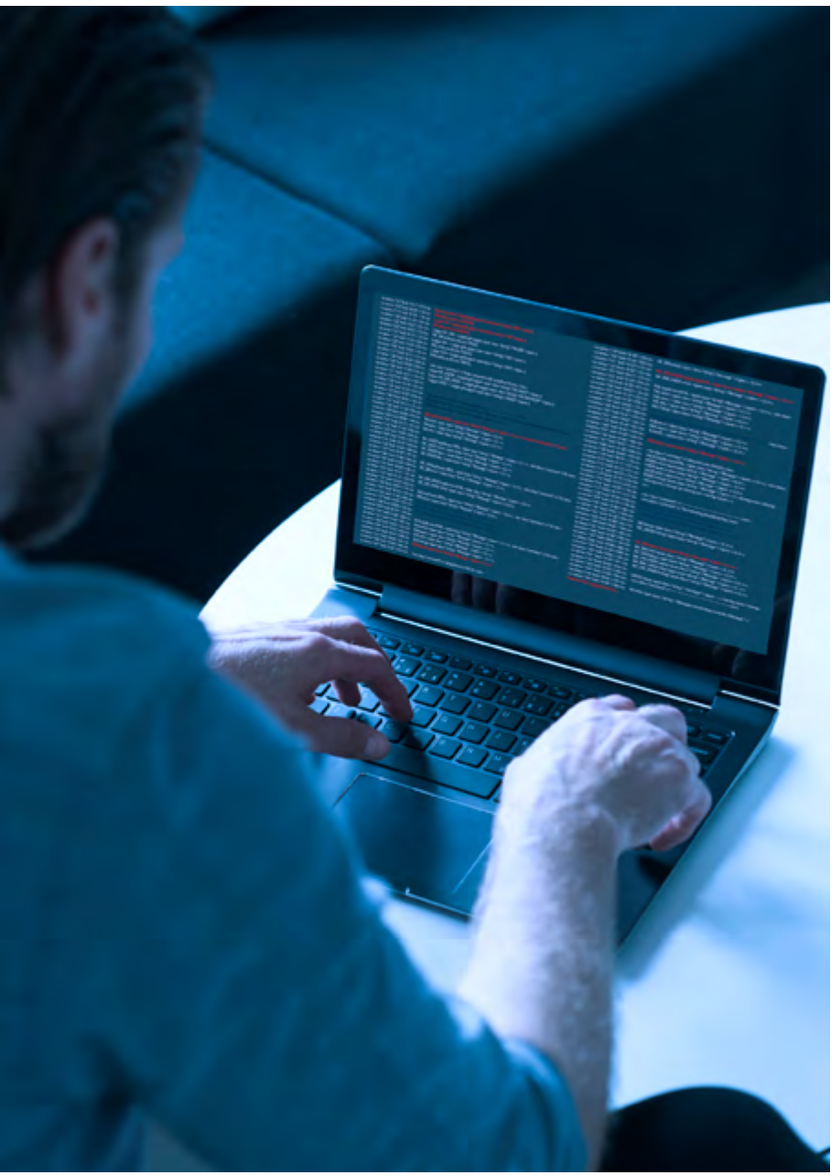

as a DNS response policy zone (RPZ) to either stop the IP from resolving or to redirect the user to a safe domain or to a block page. The RPZ is basically a collection of rules for how the DNS resolver should respond when users query a malicious domain.

The US Cybersecurity and Infrastructure Security Agency (CISA) and the National Security Agency (NSA) have described PDNS services as helping organizations detect and block access to known phishing sites, malware distribution and C2 sites, and sites that are programmatically generated using domain-generation algorithms. Organizations can also use PDNS services to prevent users from accessing sites that go against their policies, such as sites associated with online gambling. Unlike other DNS security enhancements, such as DNSSEC and DNS over Transport Layer Security, PDNS is not a protocol but a security service, the two organizations have emphasized.

CISA and NSA recommend that organizations tap PDNS services as part of a multilayered defense-in-depth strategy. One big benefit according to the agencies is that a PDNS can be set up easily just by changing the organization's recursive resolver to use the PDNS provider's recursive server instead. When setting up a PDNS capability, organizations need to ensure that the vendor provides enterprise dashboard views, malicious activity alerts, and historical logging and analytics capabilities, the NSA and CISA have noted. Administrators can integrate PDNS with enterprise SIEM systems or interact with them via APIs or through a Web interface.

"Protecting users' DNS queries is a key defense because cyber threat actors use domain names across the network exploitation lifecycle," the NSA and CISA noted last year in guidance around how to choose a PDNS provider. "The domain names associated with malicious content are often known or knowable, and preventing their resolution protects individual users and the enterprise."

The UK government also has emphasized the use of PDNS as an effective mechanism for addressing a broad range of modern threats. Federal government agencies in the country are required to use a PDNS created by the UK National Cyber Security Center (NCSC) to control access to known bad sites.

When considering a PDNS vendor, organizations should evaluate the vendor's adherence to privacy and security best practices, according to the NSA and CISA. Because the vendor can view the organization's DNS queries, security leaders must make sure they understand how the PDNS provider will use the DNS data and especially whether they will use it for any non-security-related purposes. One other caveat, according to the agencies, is that PDNS doesn't work in situations where IP addresses are used directly without any DNS lookups — or, in other words, without a DNS resolver. For this reason, organizations should not rely on PDNS alone to detect and protect against malicious traffic.

"Active defense against DNS threats requires protective DNS as recommended by both the US and UK governments," says Paul Vixie, an Internet pioneer and the designer of several DNS protocol extensions. Organizations can enable this capability by deploying a DNS firewall of some kind — either by operating a local DNS server for their network—or selecting a DNS operator that offers PDNS, he says.

Running a local DNS server ensures that private information, like what a user might have queried, remains private and internal to the organization, says Vixie, who most recently was CEO of Farsight Security before its acquisition by DomainTools. It also avoids reliance on distant third parties, such as cloud DNS operators, he says. In this vein, the most important question to ask when acquiring PDNS services is whether the vendor has an on-premises offering, such as DNS firewalls with RPZ, Vixie says.

The second most important question is whether the PDNS technology respects end-user privacy in the manner required by statutes such as the European Union's General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), he says. "It is vital that our defensive systems do not detract from individual privacy," Vixie says. "Data mining must be strictly avoided, or else the cure will be seen as worse than the disease."

With PDNS services, when a user wants to navigate to a site, the DNS provider's recursive resolver analyzes the request, acquires the requested IP address, and uses threat intelligence to determine if the requested connection is safe or not.

The Need for Passive DNS

Protective DNS services and techniques enable organizations to mount an active defense against DNS threats. But that's not always enough. Even with an effective PDNS layer, some attacks will get through, says Vixie. So, there's always going to be a need for at least some degree of after-attack forensics and threat-hunting capabilities. Passive DNS techniques offer a way to do this, he says.

Passive DNS is a logging technique that was developed in 2004 for storing anonymized, historical DNS records in a secure database for later analysis. It enables answers to questions such as where a domain name might have pointed to in the past, what domain names a particular nameserver might be hosting, what domain names point to a specific IP network, or how many subdomains a particular domain might have. Among other things, passive DNS can be a useful mechanism for uncovering potential security incidents and compromised systems and potential brand misuse. When combined with WHOIS records, they enable organizations to detect malicious and suspicious domains with considerable accuracy.

Passive DNS makes it possible to "pierce the veil around threat-actor infrastructure by finding related and historical indicators," Vixie says. "Passive DNS makes this possible despite the growth of DNS privacy technologies, such as DNS over HTTPS, simply by having enough sensor operators around the world to reliably reconstruct the content of DNS as witnessed by threat victims," he says.

The idea of using passive DNS data for forensic, investigative, and analytical purposes is not new. But it has received greater attention recently. Last year, for example, concerns over breaches like the one at SolarWinds prompted a White House memo requiring all federal civilian agencies to implement a passive DNS capability for, among other things, monitoring the most frequently accessed or looked-up hostnames in their environments. "Effective forensics and hunting absolutely requires passive DNS as well as historical WHOIS" records, Vixie notes.

John Bambenek, principal threat hunter at Netenrich, says organizations should also consider using passive DNS along with domain registration information to look

for attempts to phish their brand. The information also can be useful to detect and block attempts to phish key partners and cloud providers, he notes.

One concern that has been associated with passive DNS is whether it enables reconstruction of end-user Internet activity by maintaining a record of their DNS behavior, particularly in situations where the DNS behavior might be considered personally identifiable information (PII). Some regulations, such as the CCPA, include IP addresses under the definition of PII, if an IP address can be reasonably associated with a particular individual. "As always, privacy must be considered," Vixie says. Any passive DNS tool that organizations use should be completely privacy friendly and not allow observers to identify which end user made a particular DNS query, while at the same time allowing for the tracking of threat actors and infrastructure, he says.

Taking a Zero-Trust Approach

CSA's Yeoh advocates that organizations take a zero-trust approach to their DNS. All sources and connections to the DNS need to be validated and monitored for access and unusual behavior, he says. Security managers should consider using common protocols such as Dynamic Host Configuration Protocol (DHCS) and Internet Protocol Address Management (IPAM) along with secure DNS to log sources and activity on the DNS network. DHCP is a protocol for automatically assigning an IP address and some related information to any device on the network so it can communicate over the Internet. IPAM is a protocol for managing IP addresses.

Several so-called DDI platforms — for DNS, DHCP, and IP address management — are available that allow organizations to implement the kind of zero-trust approach to DNS that Yeoh recommends.

Shamun Mahmud, a senior research analyst at CSA, describes the DDI combination as giving organizations a way to log users on the network, where they might be navigating to, and where they might have been previously. When this information is coupled with threat intelligence, organizations have a way to set and enforce polices at the DNS layer and the control plane, he says. "An integrated approach to DDI is key here," Mahmud notes. "Oftentimes, the three components are implemented separately. This can — and often does — lead to inherently insecure systems due to interoperability and codependency issues."

Yeoh says that CSA also recommends that organizations use a software-defined perimeter (SDP) to enforce security policies and to enable early detection and blocking of malware within the DNS network. "The DNS and DDI protocols share device, location, and network behavior into a zero-trust SDP system that can enforce identity-centric and context-aware access control to the protected sources."

Protecting the DNS Infrastructure

In leveraging the DNS for defensive purposes, organizations should not overlook the need to protect DNS services itself from attacks. DNS servers and infrastructure continue to be prime targets for attackers for a variety of reasons.

The fact that every network connection begins with a DNS query makes DNS a target for a lot of bad actors, Yeoh says. Within the enterprise, DNS is used to direct traffic from one compute source to another via an IP address. In the cloud, IP addresses point to all networks, servers,

Last year, concerns over breaches like the one at SolarWinds prompted a White House memo requiring all federal civilian agencies to implement a passive DNS capability for, among other things, monitoring the most frequently accessed or looked-up hostnames in their environments.

and computer instances hosting enterprise data. "Compromising a DNS server can give an attacker unauthorized access to a source and its data, cause redirection of traffic to an alternate source, or simply disrupt the operational connection to a source," he says. Mahmud points to new threats, such as attacks on DNS records publishing and DNS cache poisoning or spoofing, which introduces false data into the DNS system.

Ullrich says another serious current issue is loss of control over domains. Attackers often try to phish credentials for domain registrars and domain management systems and use that information to manipulate DNS information, he says. "The most obvious of the attacks will redirect the company's website to a site the attacker control," he says. "More subtle attacks will add mail servers or other records."

Companies also need to pay attention to things like not allowing their domains to expire because attackers can re-register expired domains and use them for follow-up attacks, he says. Look-alike domains continue to be a problem in phishing and other impersonation attacks that are often used, for example, in business email compromise or fake job offer scams, Ullrich says.

Most threats to DNS infrastructure and the controls for addressing them are well understood. Typical measures include using DNNSEC to validate DNS data, using a DNS firewall, and using strong authentication to control access to DNS servers.

The NSA also advocates that organizations encrypt DNS in enterprise environments to protect the integrity of DNS records and to keep DNS lookup information safe from prying threat actors. It has pointed to the DNSSEC protocol, for example, as being useful in protecting the privacy and integrity of DNS queries and responses, and protocols such as DNS over HTTPS (DoH) or DNS over TLS (DoT) as being useful for encrypting DNS requests to the recursive resolver.

"A lot of it comes down to good cybersecurity hygiene," says Jon France, the CISO at (ISC)^2. On the simpler side of things, protecting enterprise DNS is about implementing measures such as network segmentation, hardening in-house DNS servers, applying good access controls, and having well-configured firewalls. It's also important to limit zone transfers or copying the contents in a primary DNS server to a secondary server. "At the more advanced end, it's about log analysis, traffic profiling and inspection and more advanced tooling," as long as it aligns with risk management goals, France says.

Understanding a deployed base and utilization of DNS is a key element when implementing DNS security controls. So, too, is choosing the right tool set and ensuring that consumers of encrypted DNS services can access it. "Security and convenience are often in direct conflict, and achieving the right balance is key," France notes.

The Challenge That Privacy-Focused DNS Enhancements Present

"Monitoring DNS has always been a challenge and an opportunity," says Ullrich, from SANS.

Privacy regulations and advancements around DNS privacy-enhancing protocols ironically are emerging as the biggest challenges to efforts to tackle security issues at the DNS layer. Privacy regulations such as CCPA and the Children's Online Privacy Protection Act, for instance, consider IP addresses to be PII in some circumstances. And there's considerable ambiguity over how other regulations define IP addresses from a PII standpoint. There's concern that these statutes could limit the ability of organizations to collect and use DNS data to protect against security threats.

Privacy-focused DNS enhancements — such as DoT for encrypting data between a DNS server and a DNS client and DoH for communicating DNS information in encrypted fashion over HTTPS — are another problem. The goal behind these enhancements is to ensure that sensitive DNS data, which includes user behavior, is protected and cannot be misused.

But the technologies are also making it harder for security vendors and defenders in general to gain visibility into what might be going on at the DNS level. "The worst DNS-related threats at the moment are what we're doing to ourselves," Vixie says. "We should all contemplate the effect the IETF's efforts around DNS privacy, such as DNS over HTTPS, will inevitably have."

He predicts that the privacy-focused DNS enhancements could soon compromise the ability of organizations to enforce PDNS filtering or to monitor the DNS activities of local devices, apps, and users. "DNS over HTTPS has already been used by more than one botnet, to illustrate that point," he says.

Conclusion

The "phonebook of the Internet," as the DNS is often described, has become a focal point of industry efforts to protect organizations against modern cyber threats. Many, including the US and UK governments, perceive PDNS approaches as offering a way to detect and block threats before they have an opportunity to hit the enterprise perimeter. Multiple vendors currently offer PDNS as a managed service that is relatively easy to implement and can offer significant benefits, according to the NSA and CISA. However, privacy-focused DNS enhancements such as DoT and DoH could complicate efforts to implement and enforce PDNS.

About the Author: *Jai Vijayan is a seasoned technology reporter with over 20 years of experience in IT trade journalism. He specializes in writing on information security and data privacy topics. He was most recently a Senior Editor at Computerworld. He is a regular contributor to Dark Reading, CSO Online, and TechBeacon.*

DNS Name Server Hijack Attack Exposed Businesses, Government Agencies

Researchers found a "novel" class of DNS vulnerabilities in AWS Route53 and other DNS-as-a-service offerings that leaked sensitive information on corporate and government customers, with one simple registration step.

By Kelly Jackson Higgins, Editor-in-Chief, Dark Reading

Cloud security researchers from Wiz.io were poking around at Amazon Web Services' Route53 Domain Name Service (DNS) early last year when they suddenly realized that its self-service domain registration system let them set up a new hosted zone with the same name as the real AWS name server it was using. Within seconds, they watched in shock as their phony name server got flooded with DNS queries from other AWS customers' networks: external and internal IP addresses, computer names for finance, human resources, production servers, and organization names.

All told, they got traffic from more than 15,000 different AWS customers and a million endpoint devices, all after registering a phony AWS name server as *ns-852.awsdns-42.net*, the same name as an actual AWS name server.

"We were trying figure out how break DNS and we had no idea what traffic we were getting" at first, says Ami Luttwak, co-founder and CTO of Wiz.io as well as a former member of Microsoft's cloud security team. "In theory, if you register a name server name ... it shouldn't have any impact."

DNS services such as AWS Route53 let customers update their domain name and the name server to which their domains point for DNS queries. The researchers say they just created a new hosted zone inside *ns-852.awsdns-42.net* with the same

Previous Next

moniker and pointed it to their IP address. Then they received DNS queries from Route53 customers' devices to their rogue and same-named server.

The researchers were able to use that traffic to gather a treasure trove of information on Fortune 500 firms including a commodities-trading firm, 45 US government agencies, and 85 government agencies overseas. They gleaned from that traffic data details such as the physical locations of offices and employees at some of the organizations. "We understood then that we were on top of an unbelievable set of intelligence, just by tapping for a few hours into a small portion of the network," Luttwak says. "I called it a nation-state intelligence capability using a simple domain registration."

The researchers were, for instance, able to use the DNS query data to drill down into office locations and numbers of employees at the trading firm as well as that of a large credit union subsidiary with a branch office in Iran, and other organizations. AWS fixed the hole in mid-February 2021, shortly after the researchers alerted it in January.

All it took to close the vulnerability in AWS Route53 was placing the official AWS name-server name on a so-called "ignore" list, explains Shir Tamari, head of Wiz.io's security research team. "The problem was anyone could register the official name servers on the platform, so they put the list of their name servers on an 'ignore' list" so attackers can't register them anymore.

"It was a very quick and efficient fix," Tamari adds. Luttwak and Tamari presented their findings at Black Hat USA in Las Vegas.

"O.G." DNS Meets DNSaaS

The attack took advantage of a gray area in the DNS infrastructure: an unintended and unexpected consequence of the combination of traditional, old-school DNS technology on some Windows machines and today's cloud DNS service features. Traditional DNS client software is old — some of which was written 20 years ago — and not built for cloud-based enterprise infrastructures, but instead for trusted internal enterprise domains.

Endpoints reveal sensitive information when they query the DNS server, the researchers say, and much of this is a result of the complexity of DNS itself. "DNS clients perform non-standard queries, and DNS providers allow customers to enter their own DNS zones in their server," which creates a risky combination, Luttwak says. The clients reveal details via their Dynamic DNS updates that would be fine in an internal DNS infrastructure environment but when operating within a cloud-based DNS service could leak to other customers of that service provider.

"So, when an endpoint working from home … is no longer using an [internal] DNS resolver and is accessing the network from their DNS server," it updated the

Previous Next

researchers' rogue name server instead of its own, he explains. "It's a combination of the new world where you are able to do registration of shared domains, and in all of the algorithms put into Windows 20 years ago that [use] logic built for when there was no Internet problem — that wasn't for shared DNS servers. So, the endpoints register their locations" with the cloud-based name servers, he says.

There's also the IPv6 factor: The researchers found some devices using the newer version of the Internet Protocol (IP) were exposed and thus accessible to an attacker. "Out of the millions of endpoints that sent us Dynamic DNS data, we noticed that internal IPv6 endpoints are accessible," notes Tamari. For that reason, users working from home or outside the office and running on IPv6 risk exposing their devices to the Internet.

Tamari says the researchers found that some 6% of IPv6 devices are exposed via HTTP, RDP (Remote Desktop Protocol), and SMB, for example.

Defending Your DNS

The researchers say they can't confirm whether any attackers employed this weakness in the DNS, but they sounded the alarm that it could also exist in other DNS providers' services. "It's important for all DNS providers" to ensure they're not leaving their customers exposed via this vulnerable DNS setup, Luttwak says.

The vuln is different from other flaws the research team has seen in cloud services. It's not a classic software bug: "The logic flows lead to unexpected results," he says. "They are hard to find, these new types of vulnerabilities. It's in the logic of how you build the [DNS] service."

DNS providers should use the DNS RFC's specifications for reserved domain names, validate domains, and verify ownership of domains, the researchers note.

Organizations also have options for protecting their DNS traffic from DNS hijacking: "There are specific things organizations can do to ensure that DynamicDNS doesn't go to a malicious server," Tamari says, such as firewalls, and tools that monitor DNS traffic to and from endpoints.

About the Author: *Kelly Jackson Higgins is the editor-in-chief at Dark Reading. She is an award-winning veteran technology and business journalist with more than two decades of experience in reporting and editing for various publications, including Network Computing, Secure Enterprise Magazine, Virginia Business magazine, and other major media properties.*

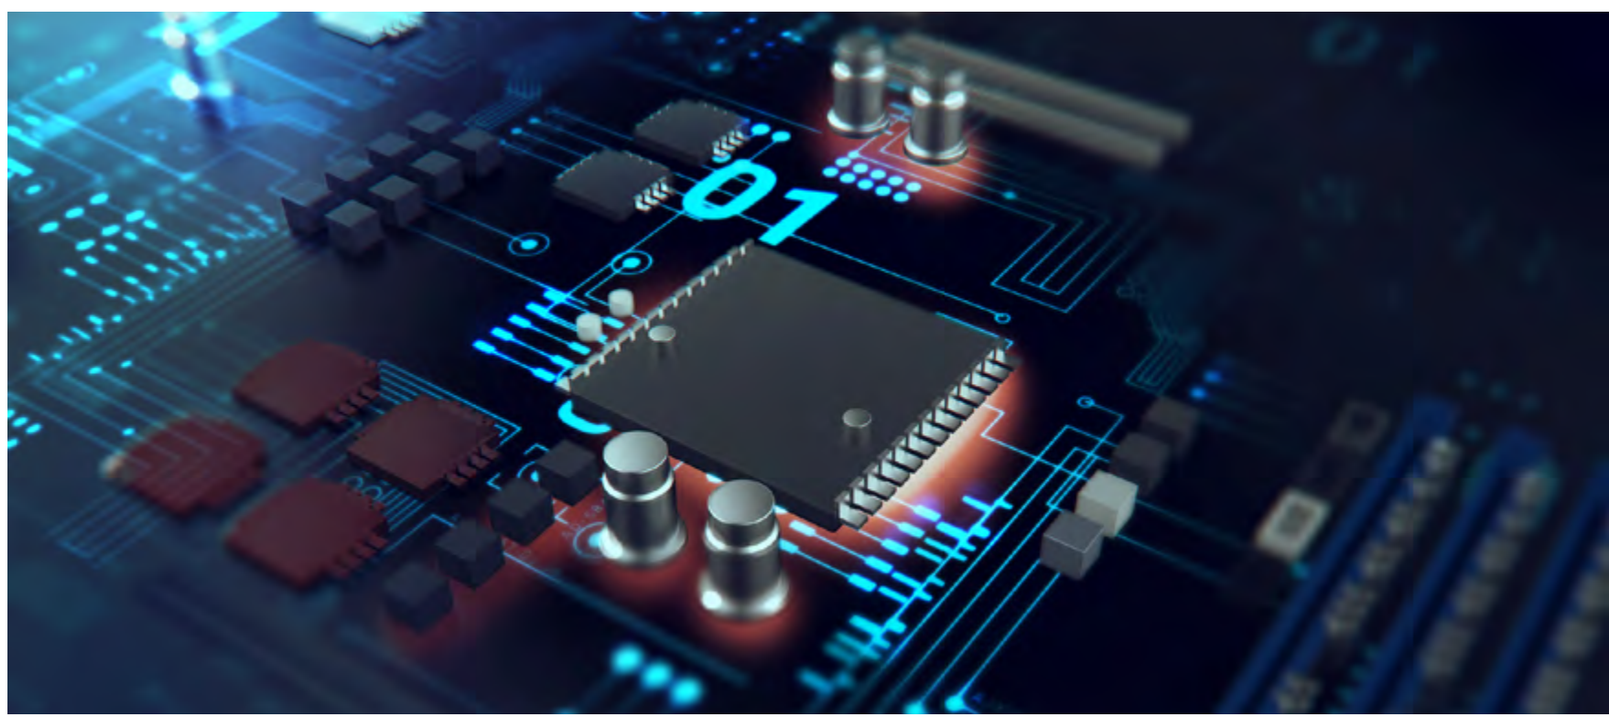

Researchers Connect Complex Specs to Software Vulnerabilities

Following their release of 70 different vulnerabilities in different implementations of TCP/IP stacks, two companies found a common link.

By Rob Lemos, Contributing Writer, Dark Reading

Six common mistakes in implementing network software led to scores of vulnerabilities, highlighting the impact that complex design requirements and ambiguous specifications can have on software security, according to two security researchers who spoke at 2021's Black Hat Asia conference.

Daniel dos Santos, research manager at network security firm Forescout, and Shlomi Oberman, CEO of security consultancy JSOF, discussed their teams' collaborative work on defining six anti-patterns, or common software mistakes, that have led to vulnerabilities in a variety of TCP/IP stacks. The two companies have disclosed a litany of vulnerabilities, many of them caused by one of the six anti-patterns.

Ambiguity and complexity in the DNS protocol caused the issues, dos Santos says.

"Because of the complexity of the DNS specification, vulnerability types that we have known about for 20 years are appearing in implementations of network stacks," he says. "The more complex the software or protocol gets, the more difficult the protocol is to implement, so we need to make them as least complex as possible, which is not always possible."

Last year, Forescout and JSOF disclosed nine vulnerabilities that affected four different TCP/IP stacks and could affect hundreds of millions of Internet of Things (IoT)

Previous Next

and network devices. The vulnerabilities were dubbed "NAME:WRECK" by the companies. Their research, called Project Memoria, also includes "Ripple20," a set of 19 vulnerabilities that affected the Treck TCP/IP stack, among others; "AMNESIA:33," a set of 33 vulnerabilities affecting four different open source TCP/IP stacks; and "NUMBER:JACK," a set of nine vulnerabilities affecting implementations of initial sequence numbers (ISN).

The companies evaluated 15 different networking software implementations, often called a "stack," and found seven had vulnerabilities due to errors in implementing a specific DNS feature.

"In Project Memoria, we learned that often the same mistake — anti-pattern — leads to similar vulnerabilities in different stacks," the researchers stated in a technical report.

The vulnerabilities occur in how different vendors implemented a DNS feature known as message compression. Because DNS responses often include the name of the specific domain several times, message compression allows software implementations to reduce the size of the DNS messages.

The researchers' discussion at Black Hat Asia focused on how complexity in the specification — specifically, the technical details to implement message compression — led to a variety of different vulnerabilities. In one case, one word — "may" versus "must" — resulted in different security issues due to a single anti-pattern.

"It is noteworthy that when a stack has a vulnerable DNS client, there are often several vulnerabilities together, but the message compression anti-pattern stands out because it commonly leads to potential [remote code execution], as it is often associated with pointer manipulation and memory operations," according to the report.

As part of their research, the companies released a technical report that discuss the six DNS anti-patterns, an open source script to identify vulnerable devices in the network, a set of queries to search for vulnerable network devices, and an amended request for comment (RFC) draft that contains more information for developers on how to avoid certain implementation mistakes.

"This research is further proof that DNS protocol complexity leads to several vulnerable implementations and that the community should act to fix a problem that we believe is more widespread of what we currently know," the companies stated in their report.

About the Author: *Rob Lemos is a contributing writer at Dark Reading. He is a veteran technology journalist of more than 20 years and a former research engineer. He has written for more than two dozen publications, including MIT's Technology Review, Popular Science, and Wired News.*

Newer Generic Top-Level Domains a Security 'Nuisance'

Ten years of passive DNS data shows classic TLDs such as .com and .net dominate newer TLDs in popularity and use.

By Jai Vijayan, Contributing Writer, Dark Reading

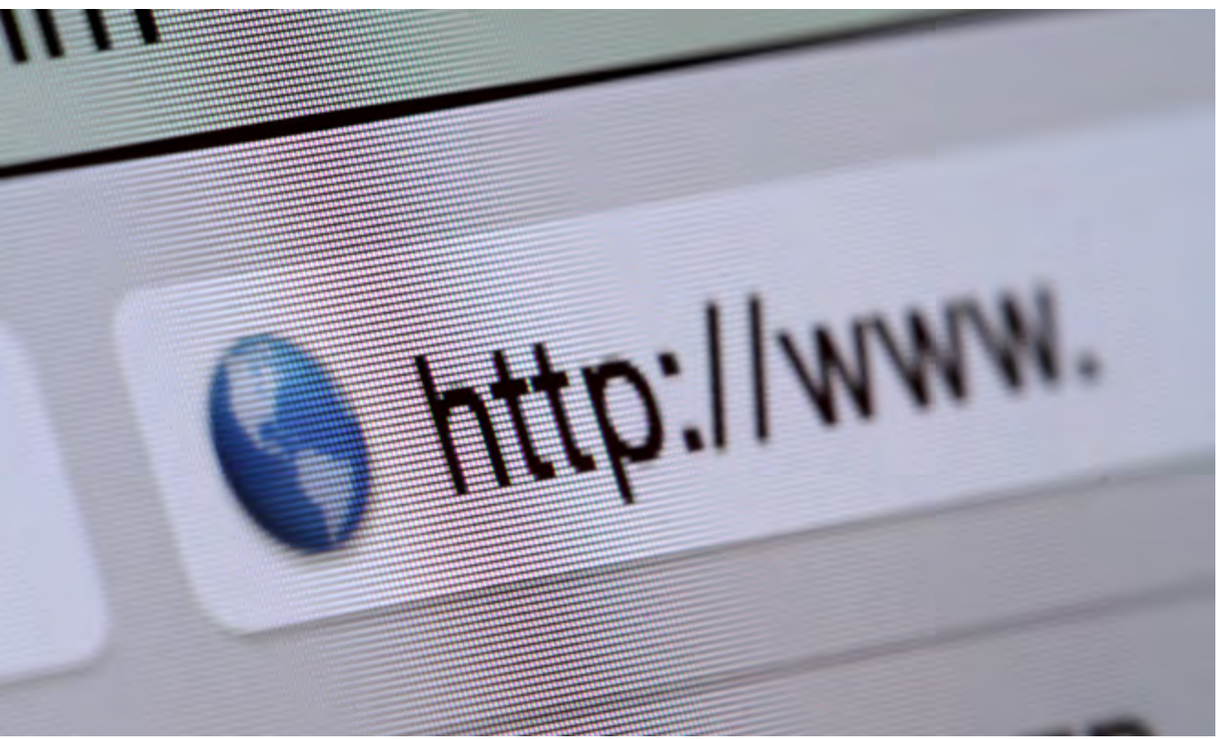

A study into the use and popularity of the Internet's top-level domains (TLDs) over a 10-year period shows that many newer TLDs may present more of a security nuisance for organizations than anything else.

That's according to Farsight Security, which released a 182-page snapshot of top-level domain traffic associated with each of 1,576 TLDs recognized by the Internet Assigned Numbers Authority (IANA). The company's findings are based on passive DNS data from 2010 to 2019 and do not include DNSSEC-related records.

The dataset includes traffic associated with generic top-level domains, such as .com, .net, and .org; country-code TLDs, such as .uk, .ca, and .de; new generic TLDs, such as .aarp, .nba, and .abc; and internationalized domain names, or TLDs with non-Latin characters.

One of the main goals of the study was to get a general sense of how broadly popular — or not — various TLDs have become over the past 10 years. While .com is generally perceived as — and actually is — the largest TLD, there's less information on the uptake of other TLDs after IANA began recognizing a lot more of them in recent years, Farsight notes in its report.

"One aspect of this to ask if it was a valuable extension of the namespace or

Previous Next

pointless nuisance" to add more TLDs in recent years, says Ben April, chief technology officer at Farsight Security. The data around TLD use suggests that the latter might well be the case, he says.

"Overall, the new TLDs aren't thriving," April says. Many have a user population and some even show signs of growth. Even so, there is no evidence of the broad migration to sector-specific TLDs that many had expected initially. "We don't see entire sectors — for example, banks — dropping .com as a primary TLD and refocusing on .bank."

From a security perspective, one concern with the growth in the number of TLDs over the past few years is that attackers have more opportunities for spoofing domains for phishing, cyber squatting, and other malicious activities. For instance, by registering a popular brand's domain name on a newer generic TLD and sending phishing emails from there, an attacker might have more success in getting victims to part with credentials and other sensitive information. In a 2019 Proofpoint study, nearly 96% of organizations found an exact match of their brand-owned domain on other TLDs.

Concerns over the threat have prompted interest in so-called defensive registrations where organization register their domains, sometimes in varied grammatical formats, on different TLDs just to prevent others from doing it for malicious purposes.

Varying Risks

What Farsight's data provides is a way for organizations to identify TLDs that present the biggest risk to their brand, April says. "When evaluating risks to your brand, the size of your target surface is directly proportional to the number of TLDs relevant to your brand," he says. "You need to evaluate each TLD to determine the level of risk it presents. This report gives you data to compare how much risk each new TLD represents."

April says the data shows that while some TLDs are likely to be worthy of concern for specific organizations, others can be safely ignored.

For example, TLDs such as .aero and .gov that have access limitations present less of a risk as registrants need to prove their identity. "If you were an airline, you don't have to worry about an attacker registering myairline.areo," April notes. Open TLDs present more of a risk, but even here that risk varies with the relevance of the TLD. "For example, if I were a retailer, I would consider TLDs like .bargains, .blackfriday, .boutique, and .shop more of a risk than TLDs like .university, .travel, and .webcam," he says.

Organizations concerned about brand abuse on new TLDs should also do substring matching, April advises. This is where a TLD may contain part of an organization's domain or brand name. "For example, if you have bobsyoga.com, you might also want to evaluate the value of a defensive registration for bobs.yoga," April says.

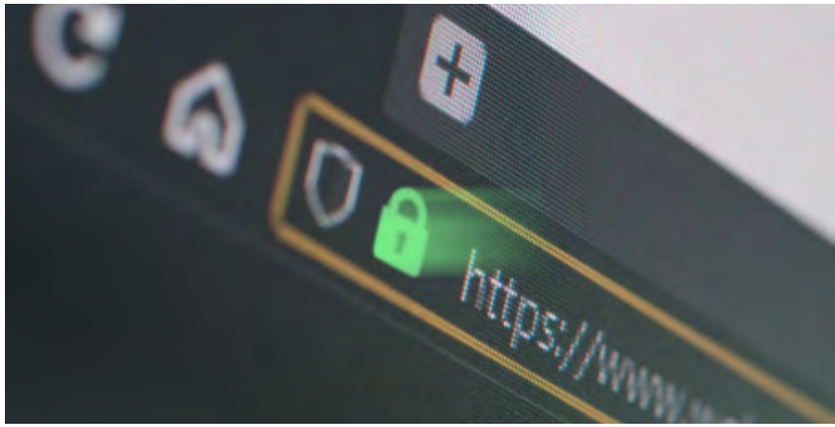

Another issue that organizations need to consider is whether TLDs have enough of a critical user base to justify accepting email from them. Decisions would need to be made on a case-by-case basis and after a careful evaluation of each TLD. "The decision to reject mail from an entire TLD is not one to be taken lightly," April says. "Organizations with a low risk tolerance and an identifiable customer/vendor base can use the data in this report to eliminate TLDs that add exposure to their security operations without also adding value."

About the Author: *Jai Vijayan is a seasoned technology reporter with over 20 years of experience in IT trade journalism. He specializes in writing on information security and data privacy topics. He was most recently a Senior Editor at Computerworld. He is a regular contributor to Dark Reading, CSO Online, and TechBeacon.*

INFOBLOX PERSPECTIVES

See More, Secure More, with DNS

Security teams should be treating DNS as one of the key tools for getting full visibility over what they have in their network and what is happening.

By Krupa Srivatsan, Director of Product Marketing, Infoblox

To say there were a lot of cyberattacks last year is a major understatement. The Kaseya supply chain attack, ransomware attacks on Oil & Gas, exploits using the Log4j vulnerability and more recently, nation-state threats due to the Russia-Ukraine crisis, are all proof of a volatile cyber landscape. As attacks become increasingly prevalent, and enterprises get more distributed, businesses are faced with the challenge of keeping networks and data safe as they modernize their technology stack.

Defense-in-Depth Is Great, but What Are You Missing?

Next-gen firewalls, email security, endpoint security, and Web gateways are crucial defense-in-depth tools to protect against various threat vectors. However, there are still gaps and open doors that allow attackers to infiltrate networks, propagate laterally, and steal data. The Domain Name System is critical network infrastructure that's needed for online connectivity. In fact, most malware, including ransomware, also relies on DNS for connecting to their command-and-control (C2) servers to download encryption software and other malicious tools onto the compromised device. The fact that many security tools don't inspect DNS is unfortunate as it provides attackers a "free ride" to evade existing security defenses and carry out their campaigns.

DNS also can be misused to infiltrate malware or exfiltrate data. This exfiltration of data via DNS (or DNS tunneling) can happen by breaking up sensitive data into small chunks and embedding them into DNS traffic going to the Internet. One example of a threat group that has used DNS tunneling extensively is OilRig. In the SUNBURST supply chain attack, the malware uses DNS to exfiltrate data about the victim, threat researchers have found.

When DNS servers are empowered with threat intelligence about known C2 and other malicious destinations, the back channel used by malware can be closed. Your DNS becomes your first line of defense against malware. Because the servers will not resolve lookups to these malicious sites, users are prevented from going to a malicious website and communications to the C2 servers from compromised devices are blocked. Coupling threat intelligence with DNS query analytics will add another layer of defense by detecting zero-day DNS-based threats.

Visibility and Asset Discovery for Efficient Security Operations

Knowing what's on the network is the first step toward protecting what's on the network. In today's world of physical data centers, multicloud deployments, direct-to-Internet branches, and IT/OT systems, full visibility seems almost a utopian concept. But there is one network asset you already own that comes close to providing you with that full visibility. You guessed right: DNS, along with DHCP and IP address management (IPAM). Together, they are called DDI, which knows at any given point what assets are on your network. The DHCP server can identify system characteristics, such as type of device and OS version, while the IPAM metadata provides information on network location and username (if integrated with Active Directory). This information becomes super critical when doing incident investigation because now you have all the forensic information to quickly understand the severity and scope of the incident. Moreover, DNS provides a valuable audit trail telling you exactly what resources a given asset has been accessing in the last hour, day, week, or month. Using DDI provides real-time discovery, whereas traditional scanning techniques need to be scheduled, may not be up to date, and can be disruptive to network operations.

Taking an Ecosystem Approach

Security teams know too well that they don't need another siloed tool in their networks that doesn't share data and lacks interoperability. Integrating security and networking tools provide a pathway for sharing data and speeding up response times. Since DNS, DHCP, and IPAM hold a gold mine of forensic (and contextual) information, integrating security information and event management (SIEM) and security orchestration, automation, and response (SOAR) tools with the DDI platform can help security operations center teams gain better understanding of the threats during analysis. If DNS security is implemented for detecting and blocking malware activity, as discussed earlier, it can also automatically trigger a vulnerability scan on the device when an event occurs via ecosystem integrations. This minimizes dwell time for threats because users are not waiting for the next scan window to detect the compromised asset.

Things to Consider When Implementing DNS Security

DNS security is best implemented on a DNS server. This ensures that protection is implemented closest to the endpoints and asset visibility is not lost. With mobility and "work from anywhere" becoming a bigger part of our lives, hybrid DNS security that provides consistent protection on and off-premises is also something to consider.