PARTNER. ADVISOR. CATALYST.

IMPACT BRIEF

APRIL 2020

Joseph Krull, CISSP, IAM, CISA, CRISC, CIPP
+1.210.421.8233
jkrull@aitegroup.com

Ahead of the Curve: Accelerate Cyber Response With DNS Insights

The Domain Name Service (DNS) functions like a giant phone directory to power the internet. Humans have a hard time remembering multiple strings of random numbers but can more easily remember proper names and words. For example, it would be quite challenging to remember the internet protocol (IP) addresses 64.233.160.0 and 69.63.176.13, but it is much easier to remember www.google.com and www.facebook.com. DNS works diligently in the background to convert those names into IP addresses and facilitates routing of requests to the proper destination on the internet. For most internet users, the power and technology behind DNS are simply unknown; however, for savvy cybersecurity professionals, DNS data can be an important tool used for both active cyber defense and investigative/forensics purposes.

DNS-derived data can reveal a broad range of attacker activities, including registration of new domains to support attacks and creation of command and control (C2) servers for malware, phishing, and ransomware. Also, because DNS data can flow freely into and out of networks, attackers actively use DNS data streams to exfiltrate stolen data and avoid detection.

Until now, using DNS data for cyber defense was limited by three key factors: timeliness of data, the need to analyze extremely large data sets, and the requirement for expert analysis by traditionally overworked cyber analysts. This Impact Brief describes how DNS can be used as part of a comprehensive cyber defense program, recent developments related to accelerating the availability and analysis of DNS data feeds, and an example of how one organization uses DNS data to rapidly detect and block attacker activities.

INTRODUCTION

Every resource on the internet has an IP address, including web and mail servers. DNS translates domain names into IP addresses and functions to help users and services navigate to the correct resource based on a domain name to IP address lookup. This can be a numeric address for IP version 4 or, when enabled by the internet service provider, an alphanumeric address for IP version 6 (e.g., 2001:0db8:85a3:0000:0000:8a2e:0370:7334). It's impractical and inefficient to maintain one massive database of DNS addresses, so DNS is composed of a series of distributed directories that maintain dynamic domain-name-to-IP-address tables. DNS queries are handled by whatever DNS resource contains the matching domain-to-IP match, or the request is forwarded to an authoritative source that can provide the correct translation.

This behind-the-scenes network traffic is relatively unknown to most internet users, as the translations are usually made in milliseconds. Daily DNS lookups are measured in the trillions, and DNS administration is very dynamic as new domains are created and organizations change their hosting providers. DNS traffic normally flows through an organization's firewall over port 53 using the user datagram protocol (UDP) or sometimes the transmission control protocol (TCP).

DNS is a real-time system that supports domain-name-to-IP lookups. Each time that DNS activity occurs, DNS databases keep records of the IP address, the DNS record, and associated data, including the servers and domains. This historical data is commonly referred to as passive DNS, and analysis of this data can be used to keep track of previous DNS records and changes. For most IT professionals, this level of DNS awareness is likely enough for them to perform their normal daily functions.

Although DNS runs "in the background" and is often taken for granted, experienced cybersecurity professionals engaged in cyber defense, security operations, malware analysis, incident response, and forensics have realized that DNS data can provide valuable information related to pre-attack, attack, and post-attack activities against their organizations. Up until recently, harvesting DNS data for cyber defense has been challenging due to delays in receiving relevant DNS data, the need to work with extremely large data sets, and the advanced level of technical expertise needed to extract value from passive DNS data.

This Impact Brief identifies methods in which DNS data can be used by cybersecurity professionals, recent developments that have enhanced the usability of passive DNS, and a case study that shows how one organization uses accelerated DNS-based feeds to improve cyber defenses.

METHODOLOGY

Aite Group conducted research on the use of DNS data for cyber defense using publicly available resources complemented by conversations with a representative sampling of security architects, incident responders, security researchers, and security operations professionals. For the case study, Aite Group also interviewed a chief information security officer, security architect, and cyber operations manager at a large technology provider.

THE MARKET

Cybersecurity teams use a variety of sources to identify attacker activities in the pre-attack, attack, and post-attack phases. Threat sources vary widely in quality and often lack proper context to be an effective cybersecurity resource (Table A). Each of these sources is overshadowed by the speed at which attackers can plan, organize, and conduct cyberattacks as well as their ability to hide their tracks after a successful breach.

Table A: The Market

Market trends	Market implications
An organization can subscribe to hundreds of open-source threat intelligence feeds.	Open-source threat feeds vary widely in quality and timeliness. Open-source feeds lack context and can introduce false positives. These feeds will likely not provide an organization with timely and actionable information to prepare for and deflect a cyberattack.
Multiple and overlapping sources of threat intelligence feeds can result in information overload.	An organization should be prudent and very selective when subscribing to threat intelligence feeds.
Several cybersecurity vendors provide commercial threat intelligence data on a subscription basis. It can be challenging to find a clear return on investment as quality and timeliness of the data do not always map to what can be a significant expense.	Organizations should examine the uniqueness and timeliness of commercial threat feeds against the relative cost. Whenever possible, high-cost commercial feeds should be relevant to the organization's specific business needs.
Threat data is also available from industry groups and government sources. Examples include the Financial Services Information Sharing and Analysis Center (FS-ISAC) and the U.S. Department of Homeland Security National Cyber Awareness System.	Threat data is curated and can derive from very sensitive sources. Some threat data must be de-identified for commercial reasons. These factors can result in substantial distribution delays due to complex release approval processes.
Parsing and analyzing threat feed data can be labor intensive and require highly experienced cyber professionals to achieve proper context for the organization. Most organizations lack the sophisticated tools to work with extremely large data sets.	Organizations starting to use threat feeds or expanding their use of these feeds should carefully consider the human element. Automation and orchestration are currently at early stages.
Recent developments in commercial threat feed distribution based on passive DNS and the application of artificial intelligence (AI) and machine learning have significantly increased the potential value of threat feeds.	This market segment is rapidly evolving and can enhance the value of passive DNS data to enhance cyber defense activities. Vendors such as managed security service providers (MSSPs) may be able to add enhanced threat feeds to their value-added services.

Source: Aite Group

METHODS TO USE PASSIVE DNS FOR ATTACK DETECTION

Attackers register new domains to support their nefarious activities and use DNS data streams to exfiltrate sensitive data from organizations. Table B details some of the attacker activities that can be detected using passive DNS.

Table B: Attack Types and Passive DNS Detection

Activity	Context
Attackers use domain generation algorithms (DGAs) in malware to create large numbers of domains. The attacker then registers one or more of the domains to allow the malware to communicate with a C2 server outside of the infected organization.	Traditionally, it was necessary to reverse-engineer or detonate the malware in a sandbox to examine the domain generation capabilities. Today, using analysis of passive DNS data, it is possible to identify behavioral patterns associated with DGAs. This method is significantly faster than malware analysis.
Attackers use DNS tunneling to hide communications with a C2 server or exfiltrate sensitive data in DNS queries and responses. As DNS tunneling is also used for legitimate purposes, such as product updates, attackers using DNS tunneling are difficult to detect.	Passive DNS data can be used to identify known domains that are associated with attacker activities. DNS data can also be used to examine queries and responses for indicators of tunnel activity. These methods can be challenging, as attackers will only use the domains for a short period of time to avoid detection.
Attackers register domains that may not be used for extended periods of time. These domains lie dormant and are then activated to support attacks, including phishing and ransomware campaigns. Likewise, recently deactivated domains that are quickly reactivated are usually suspicious.	Security professionals can use passive DNS data to track domains that have not been active for the preceding seven days. These include top-level domains, second-level domains, and fully qualified domain names. Sudden domain activation or reactivation of a domain can be an indicator of malicious activity, but early detection is important, as attackers use the domains for only a short time.
Attackers can hijack legitimate domains to redirect traffic to infrastructure that they control. They do this by manipulating underlying DNS records for domain updates to name servers or IP addresses.	It is possible to review name server or IP address changes during the preceding seven days in passive DNS data as an indicator of domain highjacking. This can be a time-consuming process for a cyber analyst unless advanced analytical tools are used.

Source: Aite Group

FUTURE USES FOR PASSIVE DNS DATA

Table B describes the methods to use passive DNS data to identify suspicious and actual attack activities. DNS data can be a rich source of valuable information for detecting, analyzing, and investigating cyberattacks, but the true value of the data exists if an organization is able to keep the attack from happening or limiting the impact from an attack. If passive DNS data is obtained quickly and the data provides enough context and reliability, an organization can proactively use firewalls, load balancers, and web application firewalls to block suspicious and known malicious domains and IP addresses. In the event of successful attacker infiltrations, DNS data can help security teams reduce the time needed to detect and remediate through expanded knowledge and insight into specific indicators of compromise.

Recent innovations regarding DNS-based insights have breathed new life into using DNS-derived data to proactively address attack activities. These innovations include collection of DNS data from many distributed global sensors, aggregating the data into easily consumable feeds, and using AI and machine learning to analyze extremely large data sets. Organizations can now add rich threat feeds to their arsenal of cybersecurity tools without the need to have highly experienced cyber professionals trying to derive context from an ocean of open-source and commercial threat data. Enhanced DNS data insights also offer MSSPs new capabilities to provide value-added services to their clients. Aite Group believes that these innovations will accelerate in 2020 and beyond to provide even more capabilities and value to passive DNS data.

CASE STUDY

The cybersecurity team at a large technology company with annual revenue in excess of US$1.2 billion has an innovation program to examine new methods to detect and reduce the impact of cyberattacks. The team had previously used passive DNS data and threat feeds to analyze attacker behavior but assessed that these sources were too cumbersome and not timely enough to stay ahead of the velocity of attacker activities. The team did not have enough staff to perform comprehensive bad actor research, particularly with regard to the high number of false positives they were seeing from other sources.

The team subscribed to UltraThreat Feeds, a new offering from Neustar, a provider of security, risk, communications, marketing, and registry solutions. As part of its Security Solutions division, Neustar has been providing DNS services for more than 22 years and has a large repository of DNS historical data. Neustar handles more than 100 billion daily DNS lookups from a large global network of platforms. Neustar also derives relevant data from its extensive OneID global identity services. UltraThreat Feeds are sourced from Neustar's own DNS and OneID data, which significantly increases the speed at which organizations can receive indicators of suspicious activities or actual attack behavior. Neustar has also applied AI and machine learning capabilities to the derived data, coupled with historical data, to rapidly apply context to the feeds and significantly reduce the need for human analysis of a mammoth data set.

The security team obtains a bundle of Neustar UltraThreat Feeds from a designated Amazon Simple Storage Service (S3). The feeds are refreshed in near-real time by Neustar approximately every 15 minutes. The team wrote an Amazon Web Services (AWS) script to automatically pull the threat feed data from the S3 bucket at regular intervals and export threat data to its security information and event management (SIEM) platform using a REST API. The team has initially focused on alerts related to DGAs and anonymized IP addresses.

Although the team members had only been using Neustar's feeds for about seven weeks at the time Aite Group spoke with them, the chief information security officer and the team reported a significant improvement in the value of passive DNS data. The team noted a significant improvement in timeliness over other threat feeds it had used in the past as well as a large decrease in false positives. The team also reported that less resources and time are needed to perform tuning than had been the case in the past. The threat data is being actively used by security operations to make faster decisions about proactive cyber defense actions. Relevant data is presented to the operations team via a Kibana dashboard. The team applies derived threat data to actual web server requests, indicators of C2 botnets, and activities that map to potential vulnerabilities derived from the open web application security project (OWASP) Top 10.

The team described UltraThreat Feeds as an advance in cyber technologies and indicated that it will be implementing additional capabilities based on the threat feeds.

RECOMMENDATIONS

Aite Group recommends that organizations take the following actions to use enhanced passive DNS data to augment their cybersecurity programs:

- If not already in use, consider adding DNS-based threat feeds to existing cyber defense efforts to have advanced warning of suspicious activities that can be a precursor to cyberattacks.
- If an organization subscribes to multiple open-source and commercial threat intelligence feeds, evaluate the relative value derived from each feed. Examine whether enhanced DNS data can replace redundant or poorly performing feeds.
- Security professionals should increase their knowledge of how DNS functions and the data that can be derived from both active and historical DNS data. Security professionals, particularly those involved in active cyber defense activities, should understand how domain registrations work and how attackers use DNS tunneling to exfiltrate sensitive data. The SANS Institute's Reading Room is a valuable source of advanced knowledge regarding DNS security.[1]
- Incident responders and forensic analysts can benefit from a large DNS historical repository and enhanced data analytics when examining cyber events and indicators of compromise, and supporting root cause analyses.
- Security vendors, particularly MSSPs, should evaluate whether timely threat feeds derived from enhanced DNS insights can be offered to their customers as a differentiated service.
- Organizations should consider using enhanced threat feeds based on passive DNS data and related automation and orchestration opportunities as these capabilities rapidly evolve.

[1] SANS Reading Room, The SANS Institute, https://www.sans.org/reading-room/, accessed April 2, 2020.

CONCLUSION

- For many years, passive DNS data has offered organizations the potential to identify suspicious activities related to domain registrations and the use of DNS tunneling as precursors to cyberattacks. However, the timeliness of data, the lack of context, the need for advanced skills, and a high degree of false positives made practical application of the data for proactive cyber defense very challenging.
- DNS-based threat feeds can reveal a wide range of attacker activities ranging from suspicious domain registrations to the use of anonymous proxies and bots.
- Organizations that use high-quality and timely DNS-based threat feeds can get additional value from traditional cyber defenses such as firewalls, load balancers, and web application firewalls to proactively block suspicious and known malicious domains and IP addresses.
- Enhanced DNS-based threat feeds can reduce the number of false detections and allow security professionals to concentrate on actionable information. In the event of successful attacker infiltrations, DNS-based threat feeds can help security teams reduce the time needed to detect and remediate through expanded knowledge of indicators of compromise.
- Recent innovations applied to DNS insights, such as the use of a large network of global sensors coupled with AI and machine learning, have simplified the use of DNS-derived data to make important cyber defense decisions. There will be additional enhancements related to automation and orchestration in 2020 and beyond.

ABOUT AITE GROUP

Aite Group is a global research and advisory firm delivering comprehensive, actionable advice on business, technology, and regulatory issues and their impact on the financial services industry. With expertise in banking, payments, insurance, wealth management, and the capital markets, we guide financial institutions, technology providers, and consulting firms worldwide. We partner with our clients, revealing their blind spots and delivering insights to make their businesses smarter and stronger. Visit us on the web and connect with us on Twitter and LinkedIn.

CONTACT

For more information on research and consulting services, please contact:

Aite Group Sales
+1.617.338.6050
sales@aitegroup.com

For all press and conference inquiries, please contact:

Aite Group PR
+1.617.398.5048
pr@aitegroup.com

For all other inquiries, please contact:

info@aitegroup.com

RELATED AITE GROUP RESEARCH

Top 10 Trends in Cybersecurity, 2020: More Ransomware, Evolving Strategies, and New Tools, January 2020.