



Interested in learning
more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

DNS: An Asset, Not a Liability

The Domain Name System, or DNS, is crucial to billions of Internet users daily, but it comes with issues that organizations must be aware of. Attackers are abusing DNS to conduct attacks that bring businesses to their knees. Fortunately, with the right detection and analysis mechanisms in place, security teams can turn DNS vulnerabilities into enterprise assets.

Copyright SANS Institute
Author Retains Full Rights

DNS: An Asset, Not a Liability

Written by **Matt Bromiley**

January 2018

Sponsored by:

Infoblox

Executive Summary

It's time to discuss one of the most fundamental parts of the Internet: the Domain Name System, or DNS. It is crucial to billions of Internet users daily, and it's what allows us to build a web presence around brand names instead of IP addresses, send email, use VoIP and basically connect to a network. However, DNS comes with its own issues that organizations must be aware of. Attackers are abusing DNS to redirect traffic to malicious sites, communicate with command and control (C&C) servers, steal data from organizations and conduct massive attacks that bring businesses to their knees. And unfortunately, many organizations simply are not prepared to mitigate, or even detect, the problems DNS might bring. Many organizations also are not aware that one of the earliest indicators of malicious activity can be DNS communications to known malicious sites. DNS can become the first line of defense against advanced threats.

Due to the criticality of DNS to maintain an Internet presence, access applications, connect to a network or simply send an email, everyone has the potential to be impacted by DNS vulnerabilities. Moreover, because of its importance in routing traffic, DNS cannot simply be disabled. Organizations must find a way to secure DNS to protect their data. Luckily, the winds are shifting to favor the organization: We are learning about effective ways to manage the DNS attack surface and even benefit from the capabilities DNS offers.



Many information security teams and vendors are aware of the importance of DNS, DHCP and IP address management (IPAM) data for diagnosing and resolving network and security problems rapidly. DNS is a natural control plane and the first line of defense for malware detection because of the scale and its logical place in a network's architecture. Incident response teams are building playbooks to include DNS, DHCP and IPAM data in their investigations in both threat hunting and incident response capacities. IPAM can provide visibility with contextual information on connected devices that can be useful in assessing risk. As new attack vectors are discovered, teams are finding ways to protect against them.

Typically, when network data is involved, organizations must learn how to work with large quantities of data at a time. DNS traffic is no different and, in fact, is likely one of the largest sources of network traffic an organization might encounter. However, with the right detection and analysis mechanisms in place, teams will be prepared to handle whatever might be thrown at them.

In this paper, we will examine the following:

- Types of DNS-based attacks and how they can impact an organization
- An example of DNS-based data exfiltration in a current variant of RAM-scraping malware
- Information security team best practices for collecting and utilizing DNS data to enhance their investigations
- Best practices for complementing existing processes with insights from DNS, DHCP and IP (DDI) data
- What's next on the horizon for DNS

DNS Attacks and an Organization's Security Posture

There's no denying it: DNS is essential to the overall functionality of the Internet. Without it, an organization's users would not be able to access internal or external sites without maintaining lists of complicated IP addresses. Even then, the routing might not work. An ISP will typically provide a DNS function, and most organizations provide internal and external DNS functionality. However, as essential as DNS is to networked operations, it is seldom included as a high-priority item within the security posture of many organizations.

Some basic network security devices, such as firewalls, ship with port 53 open, so that DNS can be used. Other ports that fall victim to default opening are 80 and 443—typically HTTP and HTTPS, respectively. Often, there is significantly more security interest in monitoring, profiling and analyzing HTTP and HTTPS traffic, thus they typically receive the analytical priority. Contrarily, DNS protection and security fall on the organization to design and manage.

Whether organizations manage their DNS or outsource it, they must first understand the risks posed by underestimated or unconsidered DNS. Quite simply, the risk of not including DNS in an organization's security posture can open it up to potentially damaging and disruptive attacks.

Following are a few common DNS attacks or misuses that organizations should be aware of:

DNS Cache Poisoning

DNS cache poisoning, also referred to as DNS “spoofing,” allows attackers to redirect traffic from legitimate to malicious websites. When a user visits a website, multiple DNS servers will likely help in providing the correct IP address to serve the desired content to the user. The DNS servers that provide the relevant routing data will typically cache results for a period of time. Caching DNS results allows for DNS servers to operate more efficiently and provide answers faster, instead of performing a lookup each time.

For example, if 10 users want to visit Facebook.com, the DNS server can perform one lookup and answer the additional nine users with the same data. Aptly named, DNS cache poisoning aims to change the stored IP address of a legitimate site to that of a malicious site without the user’s knowledge.

Fixing Cache Poisoning with DNSSEC—Or Are We?

One method that Internet and security providers have been utilizing to strengthen DNS security is the use of DNS Security Extensions (DNSSEC). DNSSEC essentially works by adding cryptographic signatures to DNS records, which allows hosts to validate whether they are receiving the correct information. This method can help mitigate DNS cache poisoning—but it is not without inherent issues. Due to the addition of cryptographic data to a DNS response, the User Datagram Protocol (UDP) packets are larger. Larger packets obviously equal more data, which might be abused by attackers.

With that in mind, let’s look at specifically how attackers take advantage of data elements within DNS packets.

DDoS Attacks

DNS is also used by disruptive and hacktivist attackers to perform DDoS attacks. DDoS DNS attacks often hold records for the largest attacks, and they are typically successful in bringing down organizations for a period of time. One of the largest attacks on record was against the site belonging to journalist Brian Krebs and was measured at 363 Gbps.¹

Organizations should be aware of the following DNS-backed DDoS attacks:

- **DNS flood attacks.** Attackers attempt to overflow, or “flood,” a DNS server in attempts to consume server resources. A significant number of “legitimate” DNS requests can make the server(s) work overtime and be unable to facilitate additional requests. DNS floods can also consist of large numbers of DNS requests for non-existent domains. This type of flood attack, called a NXDOMAIN attack, also attempts to overflow server resources by looking for nonexistent domains. The overall goal of a flood attack is to consume resources and prevent the DNS server(s) from responding to other requests.

¹ “KrebsOnSecurity Hit With Record DDoS,” <https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/>

- **DNS reflection attacks.** DNS reflection attacks take advantage of the hundreds of thousands or millions of publicly accessible DNS servers currently facing the Internet. An attacker will make requests to the servers but will fake or “spoof” the return address (i.e. whom the servers should send the DNS response). By spoofing the return address to be that of a victim network, the attackers can redirect a significant amount of traffic to one network, system, website, etc. The goal of the reflective attack is to overload the victim system to the point where it cannot perform normal functions or crashes due to overload.
- **DNS amplification attacks.** DNS amplification attacks are a type of reflection attack that take advantage of the ability to store larger amounts of data within specially crafted packets. If using the Extension Mechanisms for DNS (EDNS0) or DNSSEC, attackers can increase the message size within a DNS packet. With clever packet crafting, attackers can send a DNS request message that might generate a response as much as 50 or 60 times larger. When combined with a reflection attack and redirecting the oversized packets to a victim network, a DNS amplification attack can bring down target networks.

DNS Tunneling

Some firewalls leave port 53 open in order to facilitate the DNS needs of an organization. Many security analysts know that 80 and 443 are often left open, and they usually have security features wrapped around them monitoring traffic over these ports. Some attack groups have taken this trio of commonly opened ports and developed the capability to tunnel traffic over port 53 to avoid detection. Using an open port 53, attackers can push protocols such as SSH or TCP traffic over DNS. Some malware families will even use DNS exclusively as a C&C channel. DNS tunneling poses a significant risk to organizations because port 53 is rarely monitored or captured due to the amount of traffic that flows over it.

Data Exfiltration Using DNS

The use of DNS for data exfiltration is another concern stemming from DNS tunneling. Whereas DNS tunneling might involve pushing a non-standard protocol or data through DNS packets to establish a bi-directional communication link, exfiltration via DNS can be performed unidirectionally. Malware variants, as well as attack groups, have been known to utilize DNS traffic for data exfiltration—again, primarily to evade monitoring and/or alerting from taking place on other ports. In the next section, we will examine a RAM-scraping malware variant that has been modified to utilize DNS for data exfiltration.

Contrarily, malware has also been used to receive malicious data. While attackers realize they can use DNS to push data from an environment, they also know it is a sneaky method to push data to an environment. A malware variant known as DNSMessenger, for example, utilizes DNS response records to download malicious payload data. This data is subsequently concatenated and executed on the infected system.

Although the aforementioned attacks are not inclusive of all DNS-based attacks, they represent enough of a concern that organizations should consider DNS in their security postures. Unfortunately, the impact of DNS attacks runs the gamut. Some examples:

- DNS cache poisoning can lead users to malicious websites, which can lead to credential theft, financial theft and/or additional compromise.
- DNS tunneling and exfiltration can offset even the best laid information security plans to protect a company's sensitive data by taking advantage of an often-ignored port.
- Domain generation algorithms (DGA) or fast fluxing techniques, which present ever-changing lists of compromised domains and/or IP addresses, may make it even more difficult to track down malicious sites.
- If all else fails, attackers can also use DNS to launch extremely disruptive attacks that can impact the normal course of business and cause serious financial harm.

Teams must pay attention to the multiple steps within the attack lifecycle and make sure they consider DNS in their security posture going forward.

An attacker can abuse DNS at each end of the attack lifecycle. DNS can be used to perform recon, as well as remove data from the environment. Make sure DNS informs the organization's security posture.

DNS Abuse in the Wild

When determining whether to include mitigation of vulnerabilities in their security posture, some decision makers within organizations think, "That can never happen to me," or "Why would anyone want our data?" Others get convinced that some techniques are only proof of concept or utilized by highly advanced, state-sponsored attackers. Unfortunately, today's malware landscape shows that even commodity malware such as adware can utilize advanced techniques. Given the ease with which attackers can abuse DNS, organizations might be allowing data to walk out the door without a second glance.

There's an important concern to note for any type of data breach: Depending on where an organization conducts business geographically, the laws of breach identification and notification may differ. With increasing legislative awareness of data breaches, including upcoming regulations such as the General Data Protection Regulation (GDPR), we're quickly approaching the time where management can no longer afford to avoid security. Even worse, data breaches that go public may cause reputational damage or lost revenue, even before the root cause is identified. Let's examine an example of malware that has made the news in high-profile breaches, utilizing DNS as one of its exfiltration methods.

FrameworkPOS

Several high-profile public breaches in the past few years have targeted consumer credit card data. Large retailers and foodservice organizations have fallen victim to payment data breaches and have found themselves dealing with advanced attack groups that are able to utilize suites of malware with various capabilities. Sure enough, at least one point-of-sale (POS) malware family has adapted to DNS exfiltration: FrameworkPOS (see Figure 1).

FrameworkPOS is a variant that, like most POS-focused malware, takes advantage of in-memory card processing. When a credit card is swiped at a payment terminal, the data from the magnetic stripe can exist in memory, even for a brief moment, while the system performs the payment authorization. Most POS malware will scrape payment-specific processes looking for the track data and, once identified, will attempt to exfiltrate it from the environment for resale. This is where attackers tend to get crafty in their methods.

The authors of FrameworkPOS realized that DNS exfiltration was a method that would likely succeed in an environment where network connectivity is typically locked down and, per regulation, heavily monitored. Thus, the authors devised the following scheme:

- 1) When the malware identifies track data in memory, it separates the data using the “=” sign, which is a part of magnetic stripe data. This separator is indicative of “Track 2” card data.
- 2) The data, now separated into “left” and “right” sides, are separately encoded.
- 3) Those two blocks are combined into one DNS request that is crafted to be sent to an attacker-owned system. A sample encoded DNS request resembles the following:

```
<system_id>.<encoded_block_1>.<encoded_block_2>.<attacker_domain>
```

- a. The **system_id** is randomly generated by the malware.
 - b. The encoded blocks are chained together and separated with a period.
 - c. Finally, the attacker’s custom domain is inserted to ensure the record will arrive at the correct destination.
- 4) Once received, the credit card data is decoded, and the attackers can begin to utilize or sell card numbers.

FrameworkPOS is just one example of how DNS abuse is not limited to advanced, state-sponsored attackers or proof-of-concept malware. Attackers know their DNS traffic blends in with the rest of the environment, and they are banking on security analysts to miss the needles in the haystack.

However, the tides are turning. When it comes to DNS, attackers might no longer have the upper hand. Organizations are getting smarter and are turning DNS into an asset rather than a liability.

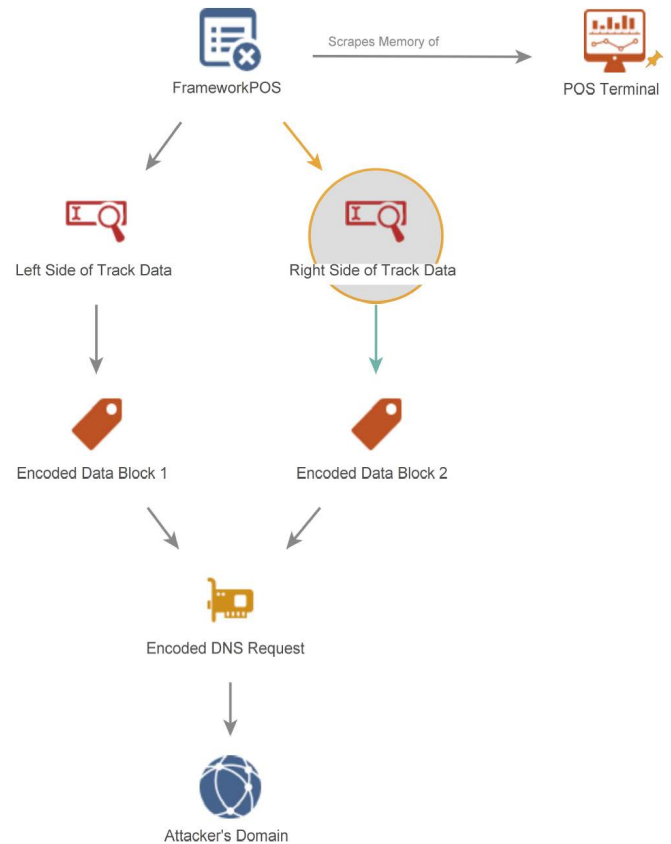


Figure 1. Diagram of FrameworkPOS

Turning a Liability into an Asset

After learning about the impacts that DNS exposure can have on an environment, the natural next step is to wonder how to mitigate these risks. Some of the attacks examined in previous sections are tough to mitigate within an organization. DDoS attacks, for example, direct so much traffic that even high-end networking devices may eventually fall. To help withstand such external attacks, it may be worth seeking out a third party that can sit in the middle and assist in DNS traffic mitigation. Mitigated DDoS attacks may still result in some Internet downtime, but your network devices will thank you for the assistance.

Other DNS abuses, such as tunneling or data exfiltration, typically occur from within an organization. When the problem starts from within, then we have mitigation and detection options. However, we can't simply "turn DNS off" – it's essential to networking. How do we work around mitigating an attack surface that we cannot turn off? Smart information security teams are taking the next step and turning DNS into a data point that can be pivoted off of and included in investigations.

Turning DNS into an asset can greatly increase an organization's effectiveness to respond to incidents.

Capturing DNS Traffic for Analysis

As previously discussed, security appliances are not always inclusive of DNS traffic due to the sheer volume of data. However, even with basic security appliances, DNS activity can be captured. Many organizations employ network flow capture solutions, which of course can be used to identify metadata about traffic communicating over port 53. For organizations that utilize full packet capture within their organizations, they will obviously have DNS traffic for the capture period. Note that network flow capture solutions and packets already traverse an organization's network—the security team simply needs to grab them.

For organizations looking to implement specific DNS capture tools or increase their DNS visibility, there are also security appliances specifically designed for DNS capture and analysis. Before implementing any device, information security managers must make sure their organization is aware of the options and tuning required for achieving the level of visibility desired. Many organizations simply plug in devices and expect results. Managers must ensure their team has time to assess, work with and evaluate the solution. Regardless of the method chosen, the first step to utilizing DNS traffic within an organization is to ensure it has access to it.

Let's examine a few techniques that can be used to turn DNS into a security asset.

Look for Anomalous DNS Traffic

One common technique to utilize DNS data to identify suspicious activity is to identify anomalous DNS queries. Detecting anomalous DNS queries can exist in multiple forms:

- **Know the organization.** Security analysts will often have an idea of what is normal and abnormal based on daily activity. Look for significant activity outside of normal activity hours.

List of DDI Analysis Techniques:

- Look for anomalous DNS requests
- Ensure continuous visibility with DHCP data
- Combine with additional log sources
- Scope infections using known indicators
- Integrate DNS, DHCP and IPAM data into existing security tools

- **Suspicious top-level domains (TLD).** If security teams have visibility into domains being requested via DNS, they should look for suspicious domains that might be outside of the organization’s typical requests. Obviously, most organizations perform constant lookups for .com, .net, .org and other popular top-level domains. Security teams should look for the outliers, such as .top or .onion, both of which are often associated with malicious activity.
- **High byte counts.** As previously mentioned, attackers can abuse DNS record sizes to include larger-than-normal data. This can also be a sign of tunneling or data exfiltration—look for DNS requests with abnormally high byte counts.
- **Individual anomalous traffic.** DNS can also provide an important clue about individual anomalous traffic. Imagine a rogue user who is visiting irregular cloud storage sites with the intention of exporting sensitive data from the environment. Because these sites are not likely visited by many within the organization, it’s possible that the DNS results are not cached and will need to be resolved. Information security teams can look for “new” DNS records that tie back to a single system and build a profile as to what the user was attempting to do.

Combine with Additional Log Sources

As with nearly any other log source, DNS data does not exist in a vacuum. It is usually the product of multiple things occurring at one time on one system multiplied by the size of the entire enterprise. As such, DNS data should not be analyzed individually either. Combining DNS data with additional log sources provides a powerful analytical punch that can be used to make information security teams exponentially more efficient.

Consider an environment that identified multiple malicious DNS requests at the perimeter coming from a handful of IP addresses within a DHCP wireless network. At first thought, tracing the DNS traffic back to the specific hosts might seem daunting, especially depending on the DHCP lease times. However, when coupled with DHCP logs, analysts can quickly identify which system had that IP at the time of the DNS request and ensure that the correct actions are taken.

The coupling of DHCP and DNS logs is just an example—think of the various network sources within a network and how they would be enhanced when combined with DNS data. Another successful combination is that of DNS data and IPAM logs. IPAM software combines DNS and DHCP data to effectively track resources and endpoints within a network. The implementation of IPAM software often increases efficiency for many organizations because they don’t have to spend resources tracking down hosts that have changing IP addresses. As organizations move into the IPv6 space, which sees a much larger pool of potential addresses, IPAM and DNS activity will be a critical pair to maintain visibility within an environment.

Scope Infections Using Known Indicators

One of the best uses for DNS data is to utilize known indicators to scope infections or breaches within an organization. Armed with threat intelligence—from both external sources and internal investigations—incident responders can go directly to DNS logs to see who else has made requests to suspicious domains. This analysis technique is extremely effective in detecting “rapid-spread” malware variants, such as worms, viruses and ransomware. Many ransomware variants will quickly mirror themselves onto multiple systems, hoping to cause enough damage that the organization will pay. Using DNS to find a scope of infected hosts can help an organization quickly nullify the impact of the malware.

What’s Next for DNS?

It should be evident by now that DNS is not a liability to an organization, but an important part of the Internet that can be harnessed as an effective defensive tool. With that in mind, security organizations must consider what’s next for DNS. The information security industry is likely to continue seeing a breakneck competition between how fast attackers can weaponize DNS and how quickly organizations can defend against said attacks.

On the horizon, information security teams are likely to see increased integration of DNS traffic within enterprise detection mechanisms. As more analysts realize the power of their own DNS, DHCP and IPAM data, we will also continue to see DNS become an integral part of incident response. Integration of DNS data is not an impossible task; APIs can be leveraged to ingest DNS data into most security solutions, such as endpoint security, vulnerability scanners, network access control (NAC) and SIEM. When coupled with incident and network context, DNS integration can lead to faster detection of compromise, prioritization based on risk and potential prevention of additional infection. If malware is discovered to be using DNS as a communication method, the teams can utilize that knowledge to help contain and quarantine the infection.

From a structure point of view, many hope that DNSSEC will continue to pave the way for secure DNS resolutions—even with the inherent larger data size. Although larger data sizes can lead to increased amplification attacks, we cannot simply dismiss security; instead, we must find a way to make it work.

Another method to make defensive use of DNS is via the implementation of DNS security as a first line of defense prior to routing traffic through a secure web gateway. DNS can be utilized as an initial security layer, subsequently directing traffic to the proxy. By placing DNS at the beginning, threats utilizing DNS can be detected faster and potentially neutralized, thereby reducing load on the web proxy.

Conclusion

As essential as DNS is to the Internet and networked devices, we must also be aware of its potential for abuse and the integral role it can play in improving an organization's security posture. Organizations that are not considering the risks of DNS as part of their security posture can find themselves victims to DNS-based attacks or malware that utilizes DNS as either a tunneling or exfiltration vehicle. Furthermore, by excluding DNS, organizations might be tipping off attackers that they have exposure elsewhere. DNS cannot simply be "turned off," so we must find a way to make it useful.

To get the most out of DNS, DHCP and IPAM, mature information security programs are complementing their existing processes and tools with DNS data points. DNS is a natural control plane and the first line of defense for malware detection because of the scale and its logical place in a network architecture. Organizations find that they can use DNS, DHCP and IPAM to help their teams be more efficient, scope incidents faster and uncover malware before it causes significant damage to the organization. Luckily, as DNS traffic typically follows a standard format, we can use DNS packets that fall outside of those norms to identify suspicious activity.

Despite the benefits some organizations are finding with DNS, DHCP and IPAM data, it is not yet time to celebrate victory over the attackers. Attackers are growing more sophisticated daily, and it's only a matter of time before a new attack comes along that takes advantage of a system used by the entire Internet. However, with the right detection mechanisms in place and a strong information security team to help minimize incident impact, organizations can rest at ease knowing their data is protected.

About the Author

Matt Bromiley is a SANS Digital Forensics and Incident Response instructor and a GIAC Advisory Board member. He is also a consulting director at a major incident response and forensic analysis company, bringing together experience in digital forensics, incident response/triage and log analytics. His skills include disk, database, memory and network forensics, as well as network security monitoring. Matt has worked with clients of all types and sizes, from multinational conglomerates to small, regional shops. He is passionate about learning, teaching and working on open source tools.

Sponsor

SANS would like to thank this paper's sponsor:





Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

CyberThreat Summit 2018	London, GB	Feb 27, 2018 - Feb 28, 2018	Live Event
SANS London March 2018	London, GB	Mar 05, 2018 - Mar 10, 2018	Live Event
SANS Secure Osaka 2018	Osaka, JP	Mar 12, 2018 - Mar 17, 2018	Live Event
SANS San Francisco Spring 2018	San Francisco, CAUS	Mar 12, 2018 - Mar 17, 2018	Live Event
SANS Paris March 2018	Paris, FR	Mar 12, 2018 - Mar 17, 2018	Live Event
SANS Secure Singapore 2018	Singapore, SG	Mar 12, 2018 - Mar 24, 2018	Live Event
SANS Northern VA Spring - Tysons 2018	McLean, VAUS	Mar 17, 2018 - Mar 24, 2018	Live Event
ICS Security Summit & Training 2018	Orlando, FLUS	Mar 18, 2018 - Mar 26, 2018	Live Event
SANS Munich March 2018	Munich, DE	Mar 19, 2018 - Mar 24, 2018	Live Event
SEC487: Open-Source Intel Beta One	McLean, VAUS	Mar 19, 2018 - Mar 24, 2018	Live Event
SANS Pen Test Austin 2018	Austin, TXUS	Mar 19, 2018 - Mar 24, 2018	Live Event
SANS Secure Canberra 2018	Canberra, AU	Mar 19, 2018 - Mar 24, 2018	Live Event
SANS Boston Spring 2018	Boston, MAUS	Mar 25, 2018 - Mar 30, 2018	Live Event
SANS 2018	Orlando, FLUS	Apr 03, 2018 - Apr 10, 2018	Live Event
SANS Abu Dhabi 2018	Abu Dhabi, AE	Apr 07, 2018 - Apr 12, 2018	Live Event
Pre-RSA® Conference Training	San Francisco, CAUS	Apr 11, 2018 - Apr 16, 2018	Live Event
SANS Zurich 2018	Zurich, CH	Apr 16, 2018 - Apr 21, 2018	Live Event
SANS London April 2018	London, GB	Apr 16, 2018 - Apr 21, 2018	Live Event
SANS Baltimore Spring 2018	Baltimore, MDUS	Apr 21, 2018 - Apr 28, 2018	Live Event
SANS Seattle Spring 2018	Seattle, WAUS	Apr 23, 2018 - Apr 28, 2018	Live Event
Blue Team Summit & Training 2018	Louisville, KYUS	Apr 23, 2018 - Apr 30, 2018	Live Event
SANS Riyadh April 2018	Riyadh, SA	Apr 28, 2018 - May 03, 2018	Live Event
SANS Doha 2018	Doha, QA	Apr 28, 2018 - May 03, 2018	Live Event
SANS SEC460: Enterprise Threat Beta Two	Crystal City, VAUS	Apr 30, 2018 - May 05, 2018	Live Event
Automotive Cybersecurity Summit & Training 2018	Chicago, ILUS	May 01, 2018 - May 08, 2018	Live Event
SANS SEC504 in Thai 2018	Bangkok, TH	May 07, 2018 - May 12, 2018	Live Event
SANS Security West 2018	San Diego, CAUS	May 11, 2018 - May 18, 2018	Live Event
SANS Melbourne 2018	Melbourne, AU	May 14, 2018 - May 26, 2018	Live Event
SANS Northern VA Reston Spring 2018	Reston, VAUS	May 20, 2018 - May 25, 2018	Live Event
SANS New York City Winter 2018	OnlineNYUS	Feb 26, 2018 - Mar 03, 2018	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced