DNS Hijacking Abuses Trust In Core Internet Service

By Cisco Talos

Published: 2019-04-17 · Archived: 2026-04-05 12:51:30 UTC

Wednesday, April 17, 2019 11:00

Preface

This blog post discusses the technical details of a state-sponsored attack manipulating DNS systems. While this incident is limited to targeting primarily national security organizations in the Middle East and North Africa, and we do not want to overstate the consequences of this specific campaign, we are concerned that the success of this operation will lead to actors more broadly attacking the global DNS system. DNS is a foundational technology supporting the Internet. Manipulating that system has the potential to undermine the trust users have on the internet. That trust and the stability of the DNS system as a whole drives the global economy. Responsible nations should avoid targeting this system, work together to establish an accepted global norm that this system and the organizations that control it are off-limits, and cooperate in pursuing those actors who act irresponsibly by targeting this system.

Executive Summary

Cisco Talos has discovered a new cyber threat campaign that we are calling "Sea Turtle," which is targeting public and private entities, including national security organizations, located primarily in the Middle East and North Africa. The ongoing operation likely began as early as January 2017 and has continued through the first quarter of 2019. Our investigation revealed that at least 40 different organizations across 13 different countries were compromised during this campaign. We assess with high confidence that this activity is being carried out by an advanced, state-sponsored actor that seeks to obtain persistent access to sensitive networks and systems.

The actors behind this campaign have focused on using DNS hijacking as a mechanism for achieving their ultimate objectives. DNS hijacking occurs when the actor can illicitly modify DNS name records to point users to actor-controlled servers. The Department of Homeland Security (DHS) issued an alert about this activity on Jan. 24 2019, warning that an attacker could redirect user traffic and obtain valid encryption certificates for an organization's domain names.

In the Sea Turtle campaign, Talos was able to identify two distinct groups of victims. The first group, we identify as primary victims, includes national security organizations, ministries of foreign affairs, and prominent energy organizations. The threat actor targeted third-party entities that provide services to these primary entities to obtain access. Targets that fall into the secondary victim category include numerous DNS registrars, telecommunication companies, and internet service providers. One of the most notable aspects of this campaign was how they were able to perform DNS hijacking of their primary victims by first targeting these third-party entities.

We assess with high confidence that these operations are distinctly different and independent from the operations performed by DNSpionage, which we reported on in November 2018. The Sea Turtle campaign almost certainly poses a more severe threat than DNSpionage given the actor's methodology in targeting various DNS registrars and registries. The level of access we presume necessary to engage in DNS hijacking successfully indicates an ongoing, high degree of threat to organizations in the targeted regions. Due to the effectiveness of this approach, we encourage all organizations, globally, to ensure they have taken steps to minimize the possibility of malicious actors duplicating this attack methodology.

The threat actors behind the Sea Turtle campaign show clear signs of being highly capable and brazen in their endeavors. The actors are responsible for the first publicly confirmed case against an organizations that manages a root server zone, highlighting the attacker's sophistication. Notably, the threat actors have continued their attacks despite public reports documenting various aspects of their activity, suggesting they are unusually brazen and may be difficult to deter going forward. In most cases, threat actors typically stop or slow down their activities once their campaigns are publicly revealed.

This post provides the technical findings you would typically see in a Talos blog. We will also offer some commentary on the threat actor's tradecraft, including possible explanations about the actor's attack methodology and thought process. Finally, we will share the IOCs that we have observed thus far, although we are confident there are more that we have not seen.

Background on Domain Name Services and records management

The threat actors behind the Sea Turtle campaign were successful in compromising entities by manipulating and falsifying DNS records at various levels in the domain name space. This section provides a brief overview of where DNS records are managed and how they are accessed to help readers better understand how these events unfolded.

The first and most direct way to access an organization's DNS records is through the registrar with the registrant's credentials. These credentials are used to login to the DNS provider from the client-side, which is a registrar. If an attacker was able to compromise an organization's network administrator credentials, the attacker would be able to change that particular organization's DNS records at will.

The second way to access DNS records is through a DNS registrar, sometimes called registrar operators. A registrar sells domain names to the public and manages DNS records on behalf of the registrant through the domain registry. Records in the domain registry are accessed through the registry application using the Extensible Provisioning Protocol (EPP). EPP was detailed in the request for comment (RFC) 5730 as "a means of interaction between a registrar's applications and registry applications." If the attackers were able to obtain one of these EPP keys, they would be able to modify any DNS records that were managed by that particular registrar.

The third approach to gain access to DNS records is through one of the registries. These registries manage any known TLD, such as entire country code top-level domains (ccTLDs) and generic top-level domains (gTLDs). For example, Verisign manages all entities associated with the top-level domain (TLD) ".com." All the different registry information then converges into one of 12 different organization that manage different parts of the domain registry root. The domain registry root is stored on 13 "named authorities in the delegation data for the root zone," according to ICANN.

Finally, actors could target root zone servers to modify the records directly. It is important to note that there is no evidence during this campaign (or any other we are aware of) that the root zone servers were attacked or compromised. We highlight this as a potential avenue that attackers would consider. The root DNS servers issued a joint statement that stated, "There are no signs of lost integrity or compromise of the content of the root [server] zone…There are no signs of clients having received unexpected responses from root servers."

Assessed Sea Turtle DNS hijacking methodology

It is important to remember that the DNS hijacking is merely a means for the attackers to achieve their primary objective. Based on observed behaviors, we believe the actor ultimately intended to steal credentials to gain access to networks and systems of interest. To achieve their goals, the actors behind Sea Turtle:

1. Established a means to control the DNS records of the target.
2. Modified DNS records to point legitimate users of the target to actor-controlled servers.
3. Captured legitimate user credentials when users interacted with these actor-controlled servers.

The diagram below illustrates how we believe the actors behind the Sea Turtle campaign used DNS hijacking to achieve their end goals.

Redirection Attack Methodology Diagram

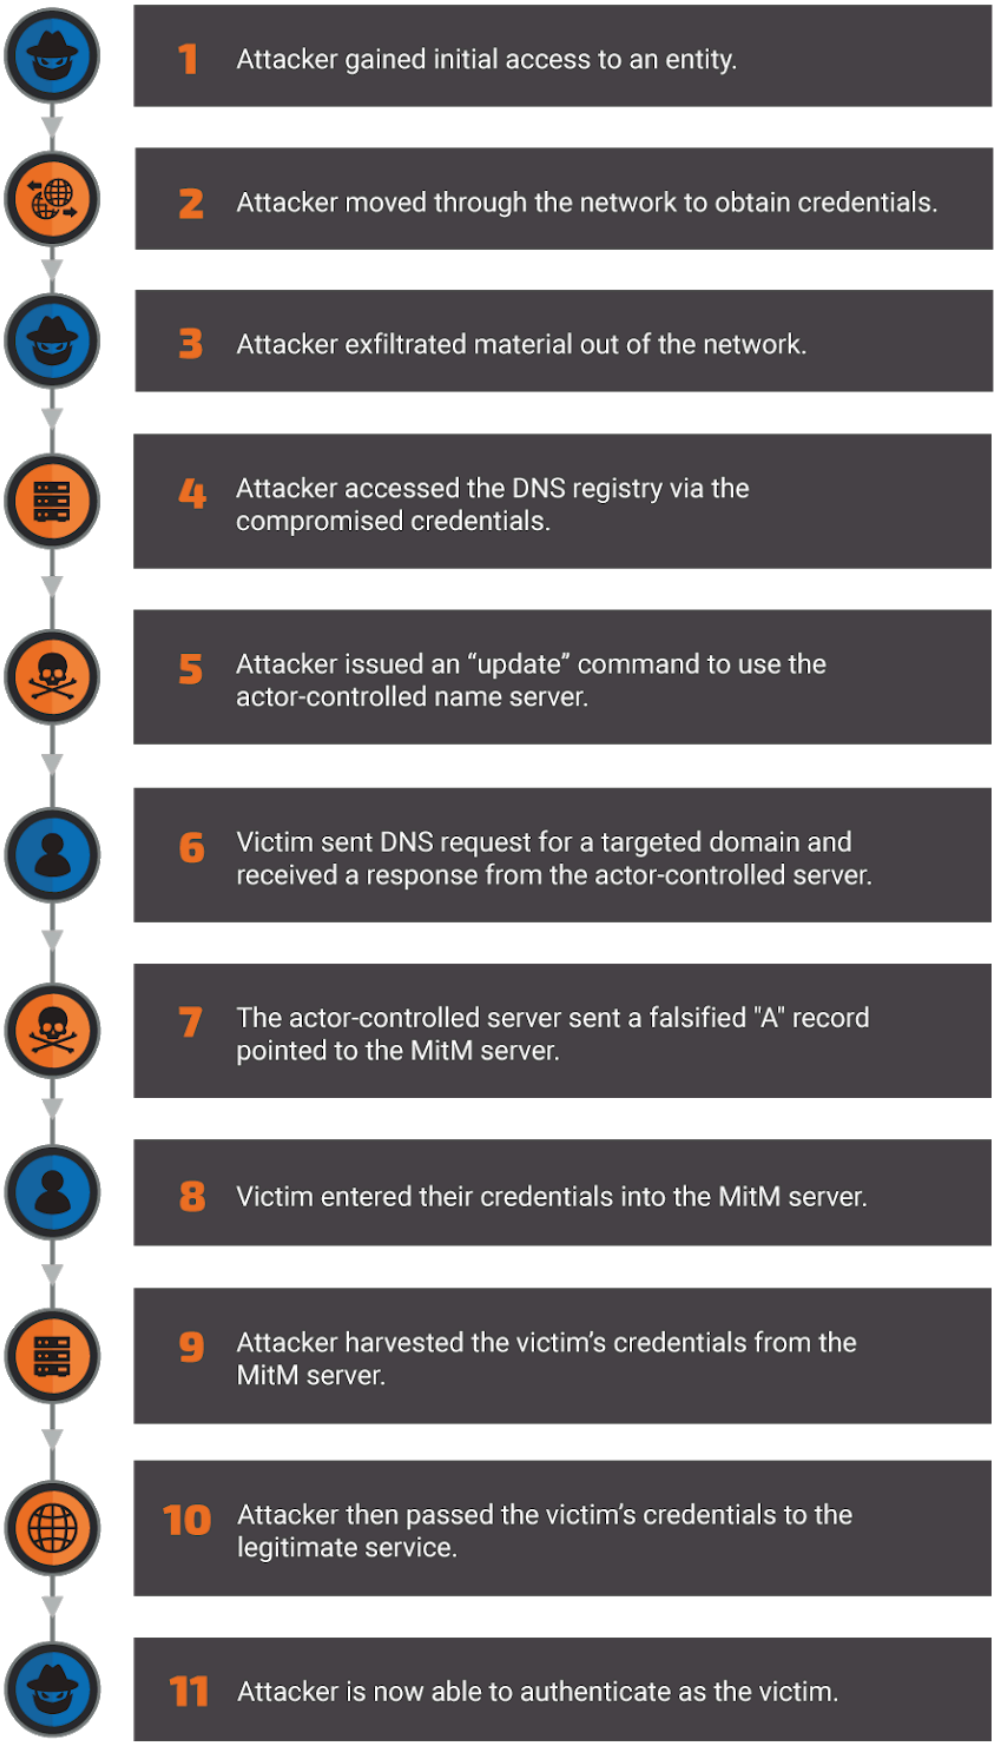
1 Attacker gained initial access to an entity.
2 Attacker moved through the network to obtain credentials.
3 Attacker exfiltrated material out of the network.
4 Attacker accessed the DNS registry via the compromised credentials.
5 Attacker issued an “update” command to use the actor-controlled name server.
6 Victim sent DNS request for a targeted domain and received a response from the actor-controlled server.
7 The actor-controlled server sent a falsified "A" record pointed to the MitM server.
8 Victim entered their credentials into the MitM server.
9 Attacker harvested the victim’s credentials from the MitM server.
10 Attacker then passed the victim’s credentials to the legitimate service.
11 Attacker is now able to authenticate as the victim.

Operational tradecraft

Initial access

The threat actors behind the Sea Turtle campaign gained initial access either by exploiting known vulnerabilities or by sending spear-phishing emails. Talos believes that the threat actors have exploited multiple known CVEs to either gain initial access or to move laterally within an affected organization. Based on our research, we know the actor utilizes the following known exploits:

- CVE-2009-1151: PHP code injection vulnerability affecting phpMyAdmin
- CVE-2014-6271: RCE affecting GNU bash system, specifically the SMTP (this was part of the Shellshock CVEs)
- CVE-2017-3881: RCE by unauthenticated user with elevated privileges Cisco switches
- CVE-2017-6736: Remote Code Exploit (RCE) for Cisco integrated Service Router 2811
- CVE-2017-12617: RCE affecting Apache web servers running Tomcat
- CVE-2018-0296: Directory traversal allowing unauthorized access to Cisco Adaptive Security Appliances (ASAs) and firewalls
- CVE-2018-7600: RCE for Website built with Drupal, aka "Drupalgeddon"

As of early 2019, the only evidence of the spear-phishing threat vector came from a compromised organization's public disclosure. On January 4, Packet Clearing House, which is not an Internet exchange point but rather is an NGO which provides support to Internet exchange points and the core of the domain name system, provided confirmation of this aspect of the actors' tactics when it publicly revealed its internal DNS had been briefly hijacked as a consequence of the compromise at its domain registrar.

As with any initial access involving a sophisticated actor, we believe this list of CVEs to be incomplete. The actor in question can leverage known vulnerabilities as they encounter a new threat surface. This list only represents the observed behavior of the actor, not their complete capabilities.

Globalized DNS hijacking activity as an infection vector

During a typical incident, the actor would modify the NS records for the targeted organization, pointing users to a malicious DNS server that provided actor-controlled responses to all DNS queries. The amount of time that the targeted DNS record was hijacked can range from a couple of minutes to a couple of days. This type of activity could give an attacker the ability to redirect any victim who queried for that particular domain around the world. Other cybersecurity firms previously reported some aspects of this activity. Once the actor-controlled name server was queried for the targeted domain, it would respond with a falsified "A" record that would provide the IP address of the actor-controlled MitM node instead of the IP address of the legitimate service. In some instances, the threat actors modified the time-to-live (TTL) value to one second. This was likely done to minimize the risk of any records remaining in the DNS cache of the victim machine.

During 2019, we observe the following name servers being used in support of the Sea Turtle campaign:

Domain	Active Timeframe
ns1[.]intersecdns[.]com	March - April 2019
ns2[.]intersecdns[.]com	March - April 2019
ns1[.]lcjcomputing[.]com	January 2019
ns2[.]lcjcomputing[.]com	January 2019

Credential harvesting: Man-in-the-middle servers

Once the threat actors accessed a domain's DNS records, the next step was to set up a man-in-the-middle (MitM) framework on an actor-controlled server.

The next step for the actor was to build MitM servers that impersonated legitimate services to capture user credentials. Once these credentials were captured, the user would then be passed to the legitimate service. to evade detection, the actors performed "certificate impersonation," a technique in which the attacker obtained a certificate authority-signed X.509 certificate from another provider for the same domain imitating the one already used by the targeted organization. For example, if a DigiCert certificate protected a website, the threat actors would obtain a certificate for the same domain but from another provider, such as Let's Encrypt or Comodo. This tactic would make detecting the MitM attack more difficult, as a user's web browser would still display the expected "SSL padlock" in the URL bar.

When the victim entered their password into the attacker's spoofed webpage, the actor would capture these credentials for future use. The only indication a victim received was a brief lag between when the user entered their information and when they obtained access to the service. This would also leave almost no evidence for network defenders to discover, as legitimate network credentials were used to access the accounts.

In addition to the MitM server IP addresses published in previous reports, Talos identified 16 additional servers leveraged by the actor during the observed attacks. The complete list of known malicious IP addresses are in the Indicators of Compromise (IOC) section below.

Credential harvesting with compromised SSL certificates

Once the threat actors appeared to have access to the network, they stole the organization's SSL certificate. The attackers would then use the certificate on actor-controlled servers to perform additional MitM operations to harvest additional credentials. This allowed the actors to expand their access into the targeted organization's network. The stolen certificates were typically only used for less than one day, likely as an operational security

measure. Using stolen certificates for an extended period would increase the likelihood of detection. In some cases, the victims were redirected to these actor-controlled servers displaying the stolen certificate.

One notable aspect of the campaign was the actors' ability to impersonate VPN applications, such as Cisco Adaptive Security Appliance (ASA) products, to perform MitM attacks. At this time, we do not believe that the attackers found a new ASA exploit. Rather, they likely abused the trust relationship associated with the ASA's SSL certificate to harvest VPN credentials to gain remote access to the victim's network. This MitM capability would allow the threat actors to harvest additional VPN credentials.

As an example, DNS records indicate that a targeted domain resolved to an actor-controlled MitM server. The following day, Talos identified an SSL certificate with the subject common name of "ASA Temporary Self Signed Certificate" associated with the aforementioned IP address. This certificate was observed on both the actor-controlled IP address and on an IP address correlated with the victim organization.

In another case, the attackers were able to compromise NetNod, a non-profit, independent internet infrastructure organization based in Sweden. NetNod acknowledged the compromise in a public statement on February 5, 2019. Using this access, the threat actors were able to manipulate the DNS records for sa1[.]dnsnode[.]net. This redirection allowed the attackers to harvest credentials of administrators who manage domains with the TLD of Saudi Arabia (.sa). It is likely that there are additional Saudi Arabia-based victims from this attack.

In one of the more recent campaigns on March 27, 2019, the threat actors targeted the Sweden-based consulting firm Cafax. On Cafax's public webpage, the company states that one of their consultants actively manages the i[.]root-server[.]net zone. NetNod managed this particular DNS server zone. We assess with high confidence that this organization was targeted in an attempt to re-establish access to the NetNod network, which was previously compromised by this threat actor.

Primary and secondary victims

We identified 40 different organizations that have been targeted during this campaign. The victim organizations appear to be broadly grouped into two different categories. The first group of victims, which we refer to as primary victims, were almost entirely located in the Middle East and North Africa. Some examples of organizations that were compromised include:

- Ministries of foreign affairs
- Military organizations
- Intelligence agencies
- Prominent energy organizations

The second cluster of victim organizations were likely compromised to help enable access to these primary targets. These organizations were located around the world; however, they were mostly concentrated in the Middle East and North Africa. Some examples of organizations that were compromised include:

- Telecommunications organizations
- Internet service providers
- Information technology firms
- Registrars
- One registry

Notably, the threat actors were able to gain access to registrars that manage ccTLDs for Amnic, which is listed as the technical contact on IANA for the ccTLD .am. Obtaining access to this ccTLD registrars would have allowed attackers to hijack any domain that used those ccTLDs.

How is this tradecraft different?

The threat actors behind the Sea Turtle campaign have proven to be highly capable, as they have been able to perform operations for over two years and have been undeterred by public reports documenting various aspects of their activity. This cyber threat campaign represents the first known case of a domain name registry organization that was compromised for cyber espionage operations.

In order to distinguish this activity from the previous reporting on other attackers, such as those affiliated with DNSpionage, below is a list of traits that are unique to the threat actors behind the Sea Turtle campaign:

- These actors perform DNS hijacking through the use of actor-controlled name servers.
- These actors have been more aggressive in their pursuit targeting DNS registries and a number of registrars, including those that manage ccTLDs.
- These actors use Let's Encrypts, Comodo, Sectigo, and self-signed certificates in their MitM servers to gain the initial round of credentials.
- Once they have access to the network, they steal the organization's legitimate SSL certificate and use it on actor-controlled servers.

Why was it so successful?

We believe that the Sea Turtle campaign continues to be highly successful for several reasons. First, the actors employ a unique approach to gain access to the targeted networks. Most traditional security products such as IDS and IPS systems are not designed to monitor and log DNS requests. The threat actors were able to achieve this level of success because the DNS domain space system added security into the equation as an afterthought. Had

more ccTLDs implemented security features such as registrar locks, attackers would be unable to redirect the targeted domains.

The threat actors also used an interesting techniques called certificate impersonation. This technique was successful in part because the SSL certificates were created to provide confidentiality, not integrity. The attackers stole organizations' SSL certificates associated with security appliances such as ASA to obtain VPN credentials, allowing the actors to gain access to the targeted network.

The threat actors were able to maintain long term persistent access to many of these networks by utilizing compromised credentials.

We will continue to monitor Sea Turtle and work with our partners to understand the threat as it continues to evolve to ensure that our customers remain protected and the public is informed.

Mitigation strategy

In order to best protect against this type of attack, we compiled a list of potential actions. Talos suggests using a registry lock service, which will require an out-of-band message before any changes can occur to an organization's DNS record. If your registrar does not offer a registry lock service, we recommend implementing multi-factor authentication, such as DUO, to access your organization's DNS records. If you suspect you were targeted by this type of activity intrusion, we recommend instituting a network-wide password reset, preferably from a computer on a trusted network. Lastly, we recommend applying patches, especially on internet-facing machines. Network administrators can monitor passive DNS record on their domains, to check for abnormalities.

Coverage

CVE-2009-1151: PHP code injection vulnerability affecting phpMyAdmin
SID: 2281

CVE-2014-6271: RCE affecting GNU bash system, specific the SMTP (this was part of the Shellshock CVEs)
SID: 31975 - 31978, 31985, 32038, 32039, 32041 - 32043, 32069, 32335, 32336

CVE-2017-3881: RCE for Cisco switches
SID: 41909 - 41910

CVE-2017-6736: Remote Code Exploit (RCE) for Cisco integrated Service Router 2811
SID: 43424 - 43432

CVE-2017-12617: RCE affecting Apache web servers running Tomcat
SID: 44531

CVE-2018-0296: Directory traversal to gain unauthorized access to Cisco Adaptive Security Appliances (ASAs) and Firewalls
SID: 46897

CVE-2018-7600: RCE for Website built with Drupal aka "Drupalgeddon"
SID: 46316

Indicators of Compromise

The threat actors utilized leased IP addresses from organizations that offer virtual private server (VPS) services. These VPS providers have since resold many of these IP addresses to various benign customers. To help network defenders, we have included the IP address, as well as the month(s) that the IP address was associated with the threat actor.

IP address	Month	Year	Country of targets
199.247.3.191	November	2018	Albania, Iraq
37.139.11.155	November	2018	Albania, UAE
185.15.247.140	January	2018	Albania
206.221.184.133	November	2018	Egypt
188.166.119.57	November	2018	Egypt
185.42.137.89	November	2018	Albania
82.196.8.43	October	2018	Iraq
159.89.101.204	December - January	2018-2019	Turkey, Sweden, Syria, Armenia, US
146.185.145.202	March	2018	Armenia
178.62.218.244	December - January	2018-2019	UAE, Cyprus
139.162.144.139	December	2018	Jordan

142.54.179.69	January - February	2017	Jordan
193.37.213.61	December	2018	Cyprus
108.61.123.149	February	2019	Cyprus
212.32.235.160	September	2018	Iraq
198.211.120.186	September	2018	Iraq
146.185.143.158	September	2018	Iraq
146.185.133.141	October	2018	Libya
185.203.116.116	May	2018	UAE
95.179.150.92	November	2018	UAE
174.138.0.113	September	2018	UAE
128.199.50.175	September	2018	UAE
139.59.134.216	July - December	2018	United States, Lebanon
45.77.137.65	March - April	2019	Syria, Sweden
142.54.164.189	March - April	2019	Syria
199.247.17.221	March	2019	Sweden

The following list contains the threat actor name server domains and their IP address.

Domain	Active Timeframe	IP address
ns1[.]intersecdns[.]com	March - April 2019	95.179.150.101
ns2[.]intersecdns[.]com	March - April 2019	95.179.150.101
ns1[.]lcjcomputing[.]com	January 2019	95.179.150.101
ns2[.]lcjcomputing[.]com	January 2019	95.179.150.101
