

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

NORMSHIELD INC. d/b/a BLACK KITE INC.,

Petitioners,

v.

BITSIGHT TECHNOLOGIES, INC.,

Patent Owner.

Case No. IPR2024-01395

U.S. Patent No. 11,652,834

**PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 11,652,834**

TABLE OF CONTENTS

	<u>Page</u>
I. INTRODUCTION	1
II. STANDING	2
III. GROUNDS	2
IV. POSITA	3
V. BACKGROUND	3
A. Technology Background	3
1. Overview of the Domain Name System (DNS) (EX1003, ¶¶49-54)	3
2. Open Source Intelligence (“OSINT”) Tools (EX1003, ¶¶55-77)	5
B. Public Availability.....	21
1. McNab.....	21
2. Internet Archive Pages	22
3. Public Videos	23
C. ’834 patent (EX1003, ¶¶78-80).....	25
1. Prosecution History (EX1003, ¶¶81-82).....	26
VI. CLAIM CONSTRUCTION	26
VII. GROUNDS	26
A. Ground 1: Zhang and McNab.....	26
1. Motivation to Combine	27
2. Claim 1	33
(a) 1[pre]: A computer-implemented method for mapping Internet Protocol (IP) addresses to an entity, the method comprising:	33
(b) 1[a]: receiving a first domain name for the entity;.....	35
(c) 1[b]: sending, to a domain name system (DNS) server, a first passive DNS query to identify first name servers for the first domain name; 1[c]:	

	receiving, from the DNS server, a list of the first name servers for the first domain name;	35
(d)	1[d]: sending, for each of the first name servers, a second passive DNS query to identify second domain names for which the first name server is authoritative; 1[e]: receiving, for each of the first name servers, a list of the second domain names for which the first name server is authoritative;	40
(e)	1[f]: sending, for each of the second domain names, a third passive DNS query to identify host names for the hosts of the second domain name and IP addresses for the host names; 1[g]: receiving a list of the host names and the IP addresses for the host names; and	48
(f)	1[h]: mapping each IP address to an attribute for the entity.	55
3.	Claim 2: The method of claim 1, wherein the first domain name is received from a user interface.	57
4.	Claim 3: The method of claim 1, wherein the list of the host names comprises names of at least a subset of the first name servers.....	59
5.	Claim 4: The method of claim 1, wherein the attribute is a Classless Inter-Domain Routing block.	61
6.	Claim 5: The method of claim 1, wherein receiving the first domain name for the entity comprises: determining the first domain name for a website of the entity.	62
7.	Claim 6: The method of claim 1, wherein the method is executed at least two separate instances to track a change over time in the IP addresses for the host names.....	62
8.	Claim 7: The method of claim 6, further comprising: automatically updating the mapping of IP addresses upon identifying the change in the IP addresses.....	63
9.	Claim 8: The method of claim 1, further comprising: identifying a shared hosting service based on the mapped IP addresses.....	64

10.	Claim 9: The method of claim 1, further comprising: determining a security rating for the entity based on the mapped IP address.	64
11.	Claim 10: The method of claim 1, wherein the entity is a first entity, and wherein mapping each IP address to the attribute for the entity comprises: determining that the IP address is mapped to attributes for the first entity and a second entity.	65
12.	Claim 11: The method of claim 10, further comprising: determining a first security rating for the first entity and a second security rating for the second entity, the first security rating and the second security rating being related.	66
13.	Claims 12-22	66
VIII.	<i>FINTIV</i>	68
IX.	MANDATORY NOTICES AND FEES	68
A.	Real Party-In-Interest	68
B.	Related Matters.....	68
C.	Counsel and Service Information	68
D.	Payment of Fees	70

EXHIBIT LIST

Ex.	Description
1001	U.S. Patent No. 11,652,834 (“’834 patent”)
1002	Prosecution History for U.S. Patent No. 11,652,834
1003	Declaration of Dr. Kevin Almeroth
1004	<i>Curriculum Vitae</i> of Kevin Almeroth Ph.D.
1005	McNab, “Network Security Assessment,” O’Reilly Media, Inc. (2008) (prior art under at least AIA §102(a)(1))
1006	U.S. Pat. Pub. 2018/0041521 A1 (“Zhang”) (prior art under at least AIA §102(a)(2))
1007	U.S. Pat. No. 6,578,066
1008	Intentionally Left Blank
1009	Intentionally Left Blank
1010	Intentionally Left Blank
1011	Intentionally Left Blank
1012	Intentionally Left Blank
1013	Declaration of Nathaniel Frank-White
1014	Declaration of Sylvia Hall-Ellis, Ph.D.
1015	Intentionally Left Blank
1016	Intentionally Left Blank
1017	Moore & Valsmith, et al., “Tactical Exploitation,” available at https://web.archive.org/web/20100208161237/ https://www.blackhat.com/presentations/bh-usa-

	07/Moore_and_Valsmith/Whitepaper/bh-usa-07-moore_and_valsmith-WP.pdf
1018	Chris Gates, “New School Information Gathering” (2008) available at https://www.carnal0wnage.com/papers/17_Gates.pdf
1019	Chis Gates, “Toorcon X Gates: New School Information Gathering,” available at http://vimeo.com/2745624
1020	Chis Gates, “Toorcon X Gates: New School Information Gathering,” available at https://web.archive.org/web/20090302045813/https://vimeo.com/2745624
1021	Social-Engineer, LLC, “Social Engineering Using Maltego,” available at https://www.youtube.com/watch?v=qiv4-wy3mxo
1022	Intentionally Left Blank
1023	Intentionally Left Blank
1024	Intentionally Left Blank
1025	’834 Patent Claim Chart, <i>BitSight Technologies, Inc. v. NormShield Inc. d/b/a Black Kite Inc.</i> , Case No. 1:23-cv-12055-MJJ, D.I. 39-11 (Dec. 11, 2023)
1026-1028	Intentionally Left Blank
1029	August 29, 2024 Email from Melissa Nezhnik
1030	Metadata for YouTube video ID qiv4-wy3mxo
1031	Relevant screen captures of Exhibit 1021

LIST OF CHALLENGED CLAIMS

Claim	Limitation
1[pre]	A computer-implemented method for mapping Internet Protocol (IP) addresses to an entity, the method comprising:
1[a]	receiving a first domain name for the entity;
1[b]	sending, to a domain name system (DNS) server, a first passive DNS query to identify first name servers for the first domain name;
1[c]	receiving, from the DNS server, a list of the first name servers for the first domain name;
1[d]	sending, for each of the first name servers, a second passive DNS query to identify second domain names for which the first name server is authoritative;
1[e]	receiving, for each of the first name servers, a list of the second domain names for which the first name server is authoritative;
1[f]	sending, for each of the second domain names, a third passive DNS query to identify host names for the hosts of the second domain name and IP addresses for the host names;
1[g]	receiving a list of the host names and the IP addresses for the host names; and
1[h]	mapping each IP address to an attribute for the entity.
2	The method of claim 1, wherein the first domain name is received from a user interface.
3	The method of claim 1, wherein the list of the host names comprises names of at least a subset of the first name servers.

Claim	Limitation
4	The method of claim 1, wherein the attribute is a Classless Inter-Domain Routing block.
5	The method of claim 1, wherein receiving the first domain name for the entity comprises: determining the first domain name for a website of the entity.
6	The method of claim 1, wherein the method is executed at least two separate instances to track a change over time in the IP addresses for the host names.
7	The method of claim 6, further comprising: automatically updating the mapping of IP addresses upon identifying the change in the IP addresses.
8	The method of claim 1, further comprising: identifying a shared hosting service based on the mapped IP addresses.
9	The method of claim 1, further comprising: determining a security rating for the entity based on the mapped IP address.
10	The method of claim 1, wherein the entity is a first entity, and wherein mapping each IP address to the attribute for the entity comprises: determining that the IP address is mapped to attributes for the first entity and a second entity.

Claim	Limitation
11	The method of claim 10, further comprising: determining a first security rating for the first entity and a second security rating for the second entity, the first security rating and the second security rating being related.
12[pre]	A computer-implemented system for mapping Internet Protocol (IP) addresses to an entity, the system comprising:
12[a]	at least one memory storing computer-executable instructions; and
12[b]	at least one processor for executing the instructions stored on the memory, wherein execution of the instructions programs the at least one processor to perform operations comprising:
12[c]	receiving a first domain name for the entity;
12[d]	sending, to a domain name system (DNS) server, a first passive DNS query to identify first name servers for the first domain name;
12[e]	receiving, from the DNS server, a list of the first name servers for the first domain name;
12[f]	sending, for each of the first name servers, a second passive DNS query to identify second domain names for which the first name server is authoritative;
12[g]	receiving, for each of the first name servers, a list of the second domain names for which the first name server is authoritative;
12[h]	sending, for each of the second domain names, a third passive DNS query to identify host names for the hosts of the second domain name and IP addresses for the host names;

Claim	Limitation
12[i]	receiving a list of the host names and the IP addresses for the host names; and
12[j]	mapping each IP address to an attribute for the entity.
13	The system of claim 12, further comprising: a user interface coupled to the processor, wherein the processor is configured to receive the first domain name for the entity from the user interface.
14	The system of claim 12, wherein the list of the host names comprises names of at least a subset of the first name servers.
15	The system of claim 12, wherein the attribute is a Classless Inter-Domain Routing block.
16	The system of claim 12, wherein receiving the first domain name for the entity comprises: determining the first domain name for a website of the entity.
17	The system of claim 12, wherein the operations are executed at least two separate instances to track a change over time in the IP addresses for the host names.
18	The system of claim 17, wherein the operations further comprise: automatically updating the mapping of IP addresses upon identifying the change in the IP addresses.

Claim	Limitation
19	The system of claim 12, wherein the operations further comprise: identifying a shared hosting service based on the mapped IP addresses.
20	The system of claim 12, wherein the operations further comprise: determining a security rating for the entity based on the mapped IP address.
21	The system of claim 12, wherein the entity is a first entity, and wherein mapping each IP address to the attribute for the entity comprises: determining that the IP address is mapped to attributes for the first entity and a second entity.
22	The system of claim 21, wherein the operations further comprise: determining a first security rating for the first entity and a second security rating for the second entity, the first security rating and the second security rating being related.

I. INTRODUCTION

Domain Name System (“DNS”) servers are a fundamental, ubiquitous and long-standing technology developed during the advent of the Internet. At their core, they resolve a domain name into an IP address through the querying of DNS records. Network administrators have for decades used various DNS tools to query DNS records to enumerate the IP addresses of a network. These include well-known command line tools like “whois,” “nslookup,” “dig,” and “txdns.” By using such tools, network administrators can obtain information about a domain name, its authoritative name servers, and the hosts and IP addresses for that domain name.

In 2005, passive DNS was developed. With passive DNS, a server collects data from DNS queries, including data about DNS records being returned in response to such queries. By passively collecting this information, a significant amount of data about a domain name and name servers can be collected and stored, including information like the name servers of that domain name, other domains related to that domain name, and hosts of that domain.

The ’834 patent is a 2013 patent that uses these well-known technologies to enumerate IP addresses of an entity. Specifically, the ’834 patent uses three passive DNS queries, all of which had been known and frequently used by network administrators: (1) a query for obtaining a domain’s name servers, (2) a query for obtaining other domain names that share the name server, and (3) a query for

obtaining hosts and IP addresses of a domain name. Yet the '834 patent does not purport to have invented passive DNS or any of these DNS queries, nor does it purport to invent enumeration of IP addresses.

The references used in this petition reflect the conventional nature of the '834 patent's DNS queries. Zhang is a patent filed by Palo Alto Networks that describes creating an association graph by enumerating IP addresses of domains associated with an entity controlling distribution of the malware. Zhang discloses collecting passive DNS information to collect DNS records, name servers IP addresses, and hostnames for these domains. McNab is an O'Reilly-published book called "Network Security Assessment" which describes various methods for enumerating IP addresses using the aforementioned DNS tools for querying DNS records. These include tools for enumerating name servers, other domain names, and hostnames. It would have been obvious to combine Zhang's passive DNS system with McNab's well-known use of DNS record queries to enumerate IP addresses and hostnames.

II. STANDING

Petitioners certify the '834 patent is available for IPR and they are not estopped.

III. GROUNDS

Petitioners present the below grounds.

Ground	Basis	Reference(s)	Claims
1	§103	Zhang, McNab	1-22

IV. POSITA

A person of ordinary skill in the art at the time of the alleged invention (“POSITA”) would have a Bachelor’s degree in Computer Science or similar field and at least two years of industry experience in network security. EX1003, ¶¶46-47. Petitioner further proposes that more education could compensate for less experience and vice versa. *Id.* Petitioner’s expert, Dr. Almeroth, surpassed this level of skill in the art at the earliest filing date of the ’834 patent. *Id.*

V. BACKGROUND

A. Technology Background

The ’834 patent is directed to using well-known domain name system (DNS) technologies to enumerate or discover IP addresses of assets within an entity’s network, also known as “network enumeration.” EX1003, ¶¶48-77. Below is a general background of DNS technologies and their use for network enumeration that were known for decades before the earliest filing date of the ’834 patent, September 9, 2013. *Id.*

1. Overview of the Domain Name System (DNS) (EX1003, ¶¶49-54)

Within a network, including the Internet, endpoints like servers and hosts are identified using IP addresses. The most common form of an IP address is known as “IPv4,” which is a series of four numbers, each number ranging from 0 to 255. Every node that is exposed to the Internet has a unique IP address, although it is typical for

organizations to create subnetworks within the larger Internet network wherein the nodes in the subnetworks are not directly exposed to the Internet and have IP addresses that are unique within the subnetwork.

Web servers, which host websites that Internet users browse, are also identified using an IP address. However, IP addresses are difficult for end-users to memorize. Thus, early on during the development of the Internet, the “Domain Name System” service was developed. The purpose of the DNS service is to map a domain name, such as “google.com,” to an IP address. When an Internet user enters “google.com” in their browser, the browser will query a DNS server to find out the IP address of “google.com” (which is currently 74.125.195.113). When a DNS server provides this functionality, it is called a “name server.” The records within a name server that map a domain name to an IP address are called “A” records. As discussed below, an A record is a type of DNS record, and there are other types of DNS records as well.

Internet users typically query the local name server that their ISP provides. However, every domain name has one or more “authoritative name servers” for which other name servers will rely on to get the authoritative A record and other DNS records for that domain. For example, if a user’s local name server does not have an A record for “google.com,” or there is an indication that the current A record has expired, the local name server will query the authoritative name server for

“google.com” to retrieve the A record. Authoritative name servers are often the authoritative name server for multiple domain names, especially domain names of the entity that controls the authoritative name server. A domain name’s authoritative name servers are listed in an NS (Name Server) record, another type of DNS record.

Name Servers also store a type of DNS record called a PTR record, which is the opposite of an A record in that it maps an IP address to a domain name.

Within a domain there could be various sub-domains, which are often called “host,” with each host having its own A record mapping its hostname to its own IP address, as well as corresponding NS and PTR records. For example, “drive.google.com” is a hostname of a host within the google.com domain, and its A record maps the hostname to the IP address 74.125.197.100.

2. Open Source Intelligence (“OSINT”) Tools (EX1003, ¶¶55-77)

Prior to 2013, ordinary artisans in the field were aware of numerous techniques for collecting information about a network and its resources using what are known as Open Source Intelligence (“OSINT”) tools, “open source” indicating that such information could be collected without access to or permission from the target organization.

Such OSINT tools frequently were used to query the DNS records described above in order to enumerate IP addresses for an entity, and particularly an entity’s domain name. The McNab reference, which is a book entitled “Network Security

Assessment” that was published by O’Reilly in 2008, describes several of these well-known tools. For example, McNab describes a tool called “whois,” which can look up various information about a domain name, including its authoritative name servers. EX1005, 22.

Example 3-1. Obtaining the domain WHOIS record for blah.com

\$ whois blah.com

```
Domain Name: BLAH.COM
Registrar: NETWORK SOLUTIONS, LLC.
Whois Server: whois.networksolutions.com
Referral URL: http://www.networksolutions.com
Name Server: NS1.BLAH.COM
Name Server: NS2.BLAH.COM
Name Server: NS3.BLAH.COM
Status: clientTransferProhibited
Updated Date: 04-oct-2006
Creation Date: 20-mar-1995
Expiration Date: 21-mar-2009
```

McNab describes another well-known tool called “nslookup” which “can identify ... hostnames” and “basic domain details” for a given domain, such as “cia.gov” in the example below. EX1005, 30-31. This query also returns the domain name’s IP address and authoritative name servers, along with other information.

Example 3-6. Using nslookup to enumerate basic domain details

```
$ nslookup
> set querytype=any
> cia.gov
Server:          213.228.193.145
Address:         213.228.193.145#53

Non-authoritative answer:
Name:   cia.gov
Address: 198.81.129.100
cia.gov nameserver = relay7.ucia.gov.
cia.gov nameserver = auth100.ns.uu.net.
cia.gov nameserver = relay1.ucia.gov.
cia.gov
      origin = relay1.ucia.gov
      mail addr = root.ucia.gov
      serial = 511250020
      refresh = 7200
      retry = 3600
      expire = 604800
      minimum = 86400
cia.gov mail exchanger = 10 mail1.ucia.gov.
cia.gov mail exchanger = 5 mail2.ucia.gov.

Authoritative answers can be found from:
relay7.ucia.gov internet address = 198.81.129.186
auth100.ns.uu.net      internet address = 198.6.1.202
relay1.ucia.gov internet address = 198.81.129.193
```

McNab also describes another well-known tool called “dig” which can perform a number of functions related to DNS queries. For example, McNab describes using the “dig” command to retrieve a list of hostnames for a specified domain name (“ucia.gov” in the example below) from a specified name server (“relay2.ucia.gov” in the example below), using a technique called a “DNS zone transfer.” EX1005, 33.

Example 3-8. Using dig to perform a DNS zone transfer

```
$ dig @relay2.ucia.gov ucia.gov axfr
```

```
; <<>> DiG 9.2.4 <<>> ucia.gov @relay2.ucia.gov axfr
;; global options: printcmd
ucia.gov. 3600 IN SOA relay1.ucia.gov. root.ucia.gov.
511210023 7200 900 604800 900
ucia.gov. 3600 IN NS relay1.ucia.gov.
ucia.gov. 3600 IN NS relay7.ucia.gov.
ucia.gov. 3600 IN NS auth100.ns.uu.net.
ucia.gov. 3600 IN MX 5 mail2.ucia.gov.
ain.ucia.gov. 3600 IN A 198.81.128.68
ain-relay.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
ain-relay-int.ucia.gov. 3600 IN CNAME ain-relay1-int.ucia.gov.
ain-relay1.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
ain-relay1-ext.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
ain-relay1-int.ucia.gov. 3600 IN A 192.168.64.2
ain-relay2.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
ain-relay2-ext.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
ain-relay2-int.ucia.gov. 3600 IN A 192.168.64.3
ain-relay7.ucia.gov. 3600 IN CNAME relay7.ucia.gov.
ain-relay7-ext.ucia.gov. 3600 IN CNAME relay7.ucia.gov.
ain-relay7-int.ucia.gov. 3600 IN A 192.168.64.67
ex-rtr.ucia.gov. 3600 IN CNAME ex-rtr-129.ucia.gov.
ex-rtr-129.ucia.gov. 3600 IN A 198.81.129.222
ex-rtr-129.ucia.gov. 3600 IN HINFO "Cisco 4000 Router" "NP-1E Board"
ex-rtr-191-a.ucia.gov. 3600 IN A 192.103.66.58
ex-rtr-191-b.ucia.gov. 3600 IN A 192.103.66.62
foia.ucia.gov. 3600 IN NS relay1.ucia.gov.
foia.ucia.gov. 3600 IN NS auth100.ns.uu.net.
mail1.ucia.gov. 3600 IN A 198.81.129.68
mail1out.ucia.gov. 3600 IN A 198.81.129.71
mail2.ucia.gov. 3600 IN A 198.81.129.148
mail2out.ucia.gov. 3600 IN A 198.81.129.146
relay.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
relay-int.ucia.gov. 3600 IN CNAME ain-relay1-int.ucia.gov.
relay1.ucia.gov. 3600 IN A 198.81.129.193
relay1-ext.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
relay1-int.ucia.gov. 3600 IN CNAME ain-relay1-int.ucia.gov.
relay2.ucia.gov. 3600 IN A 198.81.129.194
relay2-ext.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
relay2-int.ucia.gov. 3600 IN CNAME ain-relay2-int.ucia.gov.
relay2a.ucia.gov. 3600 IN A 198.81.129.200
relay2y.ucia.gov. 3600 IN A 198.81.129.68
relay2z.ucia.gov. 3600 IN A 198.81.129.69
relay7.ucia.gov. 3600 IN A 198.81.129.186
relay7-ext.ucia.gov. 3600 IN CNAME relay7.ucia.gov.
relay7a.ucia.gov. 3600 IN A 198.81.129.197
relay7b.ucia.gov. 3600 IN A 198.81.129.198
res.ucia.gov. 3600 IN A 198.81.129.116
wais.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
```

McNab explains that using a DNS zone transfer provides information such as:

(1) “Internal and external IP addresses for a number of systems,” (2) the hardware

of certain resources, such as “ex-rtr-129, telling us it is a Cisco 4000 series router,” and (3) the identification of various “IP address blocks,” such as the 192.103.66.0 or 198.81.128.0 blocks. EX1005, 33-34.

The dig tool can also query the PTR records of a name server, which can return other domain names for which that name server is authoritative. In the example below, the dig tool queries the name server “relay2.ucia.gov” for all PTR records that are within the 198.81.129.0 network block. EX1005, 34. This query returns other domain names like “intelligence.gov” and “nctc.gov” for which the name server “relay2.ucia.gov” is authoritative. *Id.*

Example 3-9. Using dig to perform a DNS zone transfer for a network block

```
$ dig @relay2.ucia.gov 129.81.198.in-addr.arpa axfr
```

```
; <<>> DiG 9.2.4 <<>> @relay2.ucia.gov 129.81.198.in-addr.arpa axfr
;; global options: printcmd
129.81.198.in-addr.arpa. 3600 IN SOA relay1.ucia.gov. root.ucia.gov.
509192750 7200 3600 604800 86400
129.81.198.in-addr.arpa. 3600 IN NS relay1.ucia.gov.
129.81.198.in-addr.arpa. 3600 IN NS relay7.ucia.gov.
129.81.198.in-addr.arpa. 3600 IN NS auth100.ns.uu.net.
015.129.81.198.in-addr.arpa. 3600 IN PTR wits01.nctc.gov.
100.129.81.198.in-addr.arpa. 3600 IN PTR cia.cia.gov.
101.129.81.198.in-addr.arpa. 3600 IN PTR www2.cia.gov.
103.129.81.198.in-addr.arpa. 3600 IN PTR www.intelligence.gov.
104.129.81.198.in-addr.arpa. 3600 IN PTR www.nctc.gov.
106.129.81.198.in-addr.arpa. 3600 IN PTR wits2.nctc.gov.
146.129.81.198.in-addr.arpa. 3600 IN PTR mail2out.ucia.gov.
148.129.81.198.in-addr.arpa. 3600 IN PTR mail2.ucia.gov.
186.129.81.198.in-addr.arpa. 3600 IN PTR relay7.ucia.gov.
193.129.81.198.in-addr.arpa. 3600 IN PTR relay1.ucia.gov.
194.129.81.198.in-addr.arpa. 3600 IN PTR relay2.cia.gov.
195.129.81.198.in-addr.arpa. 3600 IN PTR relay2a.cia.gov.
196.129.81.198.in-addr.arpa. 3600 IN PTR relay2b.cia.gov.
20.129.81.198.in-addr.arpa. 3600 IN PTR ddss.cia.gov.
21.129.81.198.in-addr.arpa. 3600 IN PTR ddssdata.cia.gov.
22.129.81.198.in-addr.arpa. 3600 IN PTR ddsstest.cia.gov.
```

McNab also describes using a tool called “TXDNS” to perform “Forward DNS Grinding,” which can “enumerate valid DNS address records and aliases

relating to the domain and hosts.” EX1005, 35. This is performed by using a dictionary to append a prefix to a domain name and then querying the name to determine if the host exists. In the below example, a “mail-dict.txt” dictionary is used to generate 98 possible hostnames for the domain name “bankofengland.co.uk,” and the tool returns 3 hostnames that actually exist: “mailhost.bankofengland.co.uk,” “mail2.bankofengland.co.uk,” and “mailhost.bankofengland.co.uk.” *Id.*, 35-36.

Example 3-11. Forward DNS grinding to identify mail servers

```
C:\tools> txdns -f mail-dict.txt bankofengland.co.uk
```

```
-----  
TXDNS (http://www.txdns.net) 2.0.0 running STAND-ALONE Mode  
-----
```

```
> mail.bankofengland.co.uk      - 217.33.207.254  
> mail2.bankofengland.co.uk     - 194.201.32.153  
> mailhost.bankofengland.co.uk  - 194.201.32.130  
-----
```

```
Resolved names: 3  
Failed queries: 95  
Total queries: 98  
-----
```


Several of the above tools require making multiple queries of DNS records, such as the above TXDNS example which performed at least 98 DNS queries. To increase the efficiency of frequent DNS queries, a technology known as “passive DNS” was created, wherein a passive DNS server and/or sensors (which are typically on the same machine as a regular DNS server) “collect[s] passively” DNS records “by listening to DNS traffic.” EX1006, ¶47. For a DNS server that handles a large enough number of requests (such as any ISP), the “passive DNS” server will be able to cache a large amount of information such as “IP mapping information,”

“Whois information,” “geographical IP information,” and “DNS name servers.” *Id.*, ¶79. The passive DNS server can directly store “DNS resource records” as well. *Id.*, ¶80. Such information can be stored in a database to allow for queries such as domains “mapp[ed] to the same IP,” or domains “hosted by the same name server,” or the hostnames (and their IP addresses) for a given domain. *Id.*, ¶97.


Network enumeration using DNS technologies was also described in the Maltego tool released in the late 2000s. For example, Exhibit 1018 is a presentation given in 2008 by a well-known security researcher named Chris Gates. *See also* EX1019. The presentation describes a “[b]lack box approach to info gathering” where “the customer wants you to find everything about them you can via ‘open source’ methods.” EX1018, 5. OSINT information gathering involves “no direct contact with victim’s servers OR no non-standard traffic directed toward victim.” *Id.*, 9.

One of the tools the presentation describes is Maltego, which “is a program that can be used to determine the relationships and real world links between: People ... Companies ... Web sites ... Internet infrastructure such as: Domains, DNS names, Netblocks, IP addresses” and it does this “using open source intelligent (OSINT).” *Id.*, 67. Maltego describes one method for automatically enumerating network assets given a domain name which uses a number of DNS queries. The described methodology is very similar to that described in McNab where for any


given domain name, (1) an NS server is retrieved, (2) domain names that share the same NS server are retrieved, and (3) hostnames and their IP addresses are discovered for any given domain name. First, given a domain name, Maltego will “[f]ind our ... NS servers,” such as the NS server ns1.baytsp.net for the domain baytsp.net. EX1018, 70.



Next, Maltego will look for domains that “are shared by the DNS servers,” in this case, ns1.baytsp.net and ns2.baytsp.net. *Id.*, 71.




For any of the domain names discovered, Maltego can automatically “[b]rute force our domain names and resolve them to IPs,” such as discovering hostnames such as “w2.baytsp.com” and “vpn.baytsp.com” and resolving their IP addresses. *Id.*, 72.




A video that was published on September 14, 2009 (Exhibit 1021) also describes how Maltego can enumerate IP addresses for a given domain name using the same 3-step process of retrieving an NS server, retrieving other domain names that share the NS server, and retrieving hostnames for each domain name.


For a given domain name, its NS record indicating its authoritative NS server can be retrieved. EX1021, 0:14.




This retrieves two NS servers for the domain name “shmoo.com”:
“krusty.shmoo.com” and “archimedes.shmoo.com.” *Id.*, 0:19. For each name
server, Maltego can retrieve “what other domains use these name servers.” *Id.*, 0:19.




Some of the results are shown on the right of the screen, such as “kusia.net” and “legacyhigh.com.” *Id.*, 0:41.



For each domain name discovered, Maltego can further find the hostnames for each domain. *Id.*, 5:16.





Maltego also resolves the IP address for each hostname. *Id.*, 5:34.



Beyond the well-known techniques used by Maltego, other documents describe the importance of network enumeration, *i.e.* mapping IP addresses of hostnames to an entity.

For example, in the same presentation given by Chris Gates, another tool called Fierce DNS is described. EX1018, 18. Fierce DNS “locate[s] likely targets both inside and outside a corporate network.” *Id.*, 18. Fierce DNS uses a number of DNS queries to find hosts on the same domain and IP blocks used by the domain. *Id.*, 20. The below output for “bytsp.com” shows the various hosts for that domain name. *Id.*, 21-22.





In a 2007 presentation given at the Blackhat conference provides an overview of “gathering and discovery techniques, with a concentration on third-party services and new tools” and then using such tools to “gain[] access to a target network.” EX1017 at 4. Thus, in Chapter 3, the paper states “[t]he first step to any security engagement is the initial discovery process,” and describes “discovery techniques, starting from the outside.” *Id.*, 7. The chapter goes to explain how to use tools to perform “network discovery” using “DNS zone transfers, Whois requests, and reverse DNS lookups.” *Id.*, 9. The chapter then provides descriptions of more advanced tools such as a “‘Domain Dossier’ service which combines the various

DNS and Whois requests into [a] single report, with the option to perform a quick port scan as well.” *Id.* There is also the “DomainTools.com web site” which “accepts an IP address or host name as an input and provides a list of all domains that reverse back to that IP.” *Id.*

All of the above techniques were well-known before the earliest filing date of the ’834 patent.

B. Public Availability

McNab and other exhibits discussed in this Petition were publicly available before September 9, 2013, as discussed below. EX1003, ¶¶85-98.

1. McNab

McNab is entitled “Network Security Assessment” and is a book published by O’Reilly Media in 2008. EX1005, 5. According to Dr. Hall-Ellis and her review of library records, McNab was publicly available at the Penrose Library at the University of Denver by November 10, 2008. EX1014, ¶44. She also notes that the Library of Congress created a catalog entry for the book by June 8, 2009. *Id.*, ¶45. Thus, she concludes that McNab was available by 2008, and further concludes that Exhibit 1005 is a true and correct copy as it was published in 2008. *Id.*, ¶¶42, 48.

Dr. Almeroth also agrees that McNab was available by 2008 based on his personal knowledge and as an expert in the field, because O’Reilly Media is “a reputable publisher such that copyright dates on their publications are a reliable and

trustworthy indicator of when the publication was published.” EX1003, ¶¶86-89.

Thus, he also concludes that McNab was publicly available by 2008 and that Exhibit 1005 is a true and correct copy as it was published in 2008. *Id.*

2. Internet Archive Pages

Exhibits 1017 and 1020 are webpages and documents from the Internet Archive’s Wayback Machine. According to Mr. Frank-White, the Wayback Machine archives webpages “by use of web archiving software that automatically stores copies of files available via the Internet, each file preserved as it existed at a particular point in time.” EX1013, ¶4. Each Internet Archive URL contains the date and time when the webpage was archived. *Id.*, ¶5. All of the Internet Archive exhibits are dated before September 24, 2010, and thus indicate that they were publicly available on the Internet before that date, and thus indicate that they were publicly available on the Internet before that date, and that the copies attached at Exhibits 1017 and 1020 are true and correct copies of the webpage on the date they were archived.¹

¹ While Mr. Frank-White’s description of the Internet Archive applies to any Internet Archive URL, his declaration at Exhibit 1013 does not provide “browser screenshots of the Internet Archive’s records of the archived files for the URLs” for Exhibits 1017 and 1020. EX1013, ¶¶6-7. Petitioner is seeking an additional

Dr. Almeroth, who is familiar with the Wayback Machine, agrees with Mr. Frank-White and agrees that the Internet Archive exhibits are archives of publicly available webpages at the time they were archived as indicated in the exhibits' URL, and that those exhibits are true and correct copies. EX1003, ¶¶90-91; *see Valve Corp. v. Ironburg Inventions, Ltd.*, 8 F.4th 1364, 1374-75 (Fed. Cir. 2021) (Wayback Machine webpages are “highly persuasive evidence of public accessibility”).

3. Public Videos

Exhibits 1019 and 1021 are printed webpages from Vimeo and YouTube, respectively. Vimeo and YouTube are video-hosting sites, and the URLs for Exhibits 1019 and 1021 link to videos hosted on those sites. Exhibit 1019 is a video of Chris Page presenting the slideshow at Exhibit 1018. Exhibit 1021 is a video providing a demonstration of the Maltego product. Exhibit 1031 are screenshots of relevant portions of the YouTube video at Exhibit 1021.

Dr. Almeroth opines that Vimeo's indication that Exhibit 1019's video was uploaded on January 6, 2009 is “generally a reliable indicator of when a video was publicly available.” EX1003, ¶¶92-95. In addition, Exhibit 1020 provides corroboration that the video was available in 2009, since it is a March 2, 2009

declaration from the Internet Archive for Exhibits 1017 and 1020, but had not yet received it at the time of this Petition's filing.

archived link of the Vimeo webpage. Thus, Dr. Almeroth opines that the video linked at Exhibit 1019 was publicly available by early 2009 is a true and correct copy. *Id.* Moreover, Ms. Melissa Nezhnik of Vimeo’s Trust and Safety group confirmed that the video at Exhibit 1019 was “uploaded on January 6, 2009” by the user “carnal0wnage.” Ex. 1029.

Dr. Almeroth also opines that YouTube’s date of September 14, 2009 for Exhibit 1021 is also generally reliable. EX1003, ¶¶96-98. Because the video has thousands of views, he also opines that it was publicly available by 2009, and that it was indexed by YouTube, like all uploaded videos are by default. *Id.* He also notes that the video itself shows dates from early 2009, providing corroborating evidence of the date of its public availability, and that it is a true and correct copy. *Id.* Metadata about the YouTube video indicates that it was “published[]At” September 14, 2009 and has a “view[]Count” of 60,168 views, thereby further indicating its public availability by 2009. EX1030.

Analogous to the Federal Circuit’s recognition that Wayback Machine webpages are “highly persuasive evidence of public accessibility,” pages from YouTube and Vimeo are likewise persuasive evidence of public accessibility because they are “prominent or well-known among persons of ordinary skill in the art.” *M & K Holdings, Inc. v. Samsung Electronics Co.*, 985 F. 3d 1376, 1381 (Fed. Cir. 2021). In addition, the Board has held that the date provided by the YouTube

website along with testimonial support from an expert is enough evidentiary support to establish the video’s public availability. *Mission Integrated Technologies, LLC v. Clemente*, IPR2023-01285, Paper 16 at 22-23 (PTAB Feb. 23, 2024).

C. ’834 patent (EX1003, ¶¶78-80)

Although the “Background” of the ’834 patent describes the patent as “relate[d] to security risk management” where an “entity may evaluate the magnitude of the risks associated with the other entities to make decisions about its relationships with those other entities,” the independent claims of the ’834 patent are not strictly about security risk, but instead are broadly directed to mapping IP addresses of hosts to an entity, or what is known in the art as enumerating network assets. EX1001, Abstract. The independent claims recite three passive DNS queries to obtain IP addresses for an entity: (1) a first query that uses the entity’s domain name to receive a list of the domain name’s name servers, (2) for each name server, a second query that receives a list of other domain names for which that name server is authoritative, and (3) for each of the new domain names, a third query that receives a list of hostnames and IP addresses for the respective domain name. These IP addresses are then mapped to the entity.

The specification of the ’834 patent does not explain the purported novelty of its claims, nor does it claim to invent passive DNS. Nor does the ’834 patent claim that any of the recited queries are novel—as discussed in the previous section, the

recited queries were well-known and commonly used for enumerating the IP addresses of an entity's network.

1. Prosecution History (EX1003, ¶¶81-82)

During prosecution, the examiner did not make any Section 102 or 103 rejections. Instead, in the only rejection, the examiner rejected current claims 10 and 11 under Section 112 for lacking antecedent basis and all the claims under nonstatutory double patenting. The applicant corrected the issue and filed a terminal disclaimer, after which the application was allowed.

VI. CLAIM CONSTRUCTION

The claim terms should be given their plain and ordinary meaning under *Phillips*. EX1003, ¶¶83-84.

VII. GROUNDS

A. Ground 1: Zhang and McNab

Zhang discloses generating an “association graph” which involves enumerating IP addresses and hostnames of a malicious entity by querying passive DNS information. Zhang discloses that the “association graph” can be built by querying passive DNS information to determine name servers, other domains and hostnames of those other domains.

McNab provides specific disclosures for enumerating the IP addresses and hostnames of assets of an entity by using DNS queries, including tools for querying DNS information to receive a list of authoritative name servers for a domain name,

other domain names that share an authoritative name server, and hostnames and IP addresses for those other domain names. It would have been obvious to combine those teachings in McNab with the system and methods in Zhang. EX1003, ¶¶99-100

1. Motivation to Combine

A POSITA would have found it obvious to combine McNab's host and network enumeration techniques using DNS information with Zhang's teachings of generating an association graph for malicious entities by querying a passive DNS database containing stored DNS information about the malicious entity. EX1003, ¶¶101-11.

McNab is an O'Reilly-published book entitled "Network Security Assessment" and describes various tools and methodologies for assessing the security of a network. EX1005, Cover. In particular, McNab discloses a "Network Security Assessment Methodology" that begins with "Internet Host and Network Enumeration" to identify IP addresses of network assets. EX1005, 4-7. For example, McNab discloses various methods for querying DNS records to obtain IP addresses of computer within a domain. *Id.*, 30-36. These include the usage of well-known tools such as "whois," "nslookup," "dig," and "txdns." *Id.* As described in the Background section, McNab specifically describes using such tools to perform actions such as (1) obtaining a list of authoritative name servers for a specified

domain name, (2) obtaining a list of domain names for which a specified name server is authoritative, (3) obtaining a list of hostnames and IP addresses for a specified domain name, and (4) repeating this process as new domain names are discovered.

Id.

A POSITA would be motivated to incorporate such teachings because they were well-known, effective and conventional methods of mapping IP addresses to an entity by querying DNS records. EX1003, ¶103. Indeed, Zhang describes using DNS information to generate an association graph that determines how various malicious domains are related and associated to a malicious entity. EX1006, ¶97. Such associations are determined using DNS records to resolve a domain name's IP address and name server. *Id.*, ¶88. Zhang also describes using DNS records to determine whether multiple domains are “mapp[ed] to the same IP” or “are hosted by the same name server (NS)” or had information collected “from the same source and ... within a short time period.” EX1006, ¶97. Zhang describes using DNS records to find “new sources ... of malware” such as “hostnames, Ips, and/or domains. *Id.*, ¶67.

Given these teachings in Zhang, a POSITA would have been motivated to incorporate McNab's teachings as a mechanism for querying DNS records to obtain the desired information described in Zhang, such as information about a domain name's name servers, other related domains, and the hostnames of those domains.

EX1003, ¶104. Zhang describes collecting such information to build an association graph, and thus a POSITA would have been motivated to look to McNab for well-known techniques for querying DNS records to obtain the necessary information. *Id.* In fact, because Zhang describes collecting and storing DNS records in a passive DNS database, a POSITA would have known and been motivated to make use of any of a number of well-known techniques, including those taught by McNab, for querying such records to obtain the desired information.

A POSITA would have been further motivated to incorporate McNab's teachings because McNab itself teaches the importance of its methods for enumerating IP addresses. EX1003, ¶105. McNab as discussed above is directed to "Network Security Assessment" the first step of which involves "Internet Host and Network Enumeration" that involves "query[ing] open sources to identify hosts and networks of interest," including "internal IP addresses gathered from DNS servers." EX1005, 4-5. McNab teaches the importance of this methodology: "[t]o retain complete control over your networks and data," and to "take a proactive approach to security ... that starts with assessment to identify and categorize your risks." EX1005, 1. McNab also emphasizes the "Business Benefits" of network enumeration, which is "a business enabler." *Id.* McNab also explains that the book describes "assessment of technical security and *improving the integrity and resilience of IP networks.*" *Id.*, 2. A POSITA would therefore be motivated to make

use of McNab’s teachings in order to enumerate the IP addresses of a given domain associated with malware. EX1003, ¶105. With respect to McNab’s teachings of determining IP addresses of network assets, a POSITA would have been motivated by McNab’s teaching explains that they provide “useful data about the structure of the target network” and that such “[i]nitial reconnaissance is very important because it can uncover hosts that aren’t properly fortified against attack.” *Id.*, 5. As McNab points out, companies “often neglect hosts and networks that lay off the beaten track,” giving attackers access to “peripheral networks and hosts” that were not prioritized. *Id.*

Because Zhang’s desired outcome is an association graph that associates IP addresses with malicious entities, a POSITA would be motivated to use McNab’s teachings because they were known to effectively enumerate IP addresses for a domain name. EX1003, ¶106.

In addition, a POSITA would have been motivated to use McNab’s teachings to (1) obtain a list of name servers for a domain name, (2) for each of the listed name servers, obtain a list of other domain names for which it is authoritative, and (3) for each of the other domain names, obtain a list of hostnames and IP addresses. McNab describes each of these steps, and a POSITA would have been motivated to perform these well-known steps in this well-known order to obtain a comprehensive list of enumerated IP addresses associated with a particular entity in order to build the

association graph taught in Zhang, which is also intended to associate IP addresses with an entity. EX1003, ¶107. These steps were known to provide hostnames and IP addresses not just of a designated domain name, but also of other domain names which were related to the designated domain name because of a common name server. *Id.* Indeed, Zhang itself teaches associating domain names by determining if they have a common name server and/or IP address.

Moreover, as discussed above in the Technology Background, other tools that performed these three steps were known, including the Maltego tool described in Exhibits 1018 and 1021 as discussed in detail above. EX1003, ¶108. Indeed, such tools and techniques were presented by those skilled in the art as beneficial for discovering network assets. A POSITA would have been further motivated to combine Zhang and the methodologies described by McNab, because the latter is also used by the Maltego tool and a POSITA would have known of the capabilities of the Maltego tool in automatically and efficiently enumerating IP addresses of various domain names for an entity. *Id.*

Because Zhang teaches that it builds a graph that relates and associates different IP addresses with entities (and the domains they control) known to distribute or control malware, a POSITA would be motivated to perform the above steps using McNab's teachings to obtain a comprehensive list of IP addresses associated with those malicious domains—including hostnames for other domains

that share the same name server. EX1003, ¶109. A POSITA would have desired to discover and enumerate as many IP addresses and hostnames as possible in order to have a comprehensive set of IP addresses that were associated with a malicious domain. *Id.* This would allow Zhang's system to more effectively measure the reputation of a given domain as well as be as comprehensive as possible in determining which IP addresses were malicious. *Id.* A POSITA would have known that such comprehensiveness would improve security and provide a greater chance of avoiding malware. *Id.* Indeed, the very purpose of Zhang's system is to proactively determine when software came from a domain that was related to or associated with other malicious domains—thus, a POSITA would be motivated to use as many IP address enumeration techniques that were known. *Id.*

A POSITA would have reasonably expected success in combining Zhang with McNab's teachings. EX1003, ¶110. McNab provides detailed and sufficient information to perform queries of DNS records to obtain information such as a domain's name server, other domains that use said name server, and hostnames and IP addresses of a domain. *Id.* Because Zhang discloses storage of DNS records in a passive DNS database, it would have been well-within the capabilities of a POSITA to incorporate McNab's teachings of DNS record querying into Zhang by querying the DNS records in its passive DNS database to obtain the desired information for enumerating IP addresses. *Id.* Given the ubiquity of the DNS tools

described in McNab, it would have been well-within the capabilities of a POSITA to use Zhang’s passive DNS database to query DNS records in the particular way taught by McNab. *Id.*

Finally, Zhang and McNab are analogous art to the ’834 patent because all of these references are directed to the field of network security. EX1006, Abstract (“Techniques for malware domain detection using passive Domain Name Service (DNS) are disclosed.”); EX1005, xvi (“This book tackles one single area of information security in detail: that of undertaking IP-based network security assessment in a structured and logical way.”); EX1001 at Abstract (“[A] map is generated between (a) technical assets that contribute to security characteristics of respective entities.”); EX1003, ¶111.

2. Claim 1

- (a) **1[pre]: A computer-implemented method for mapping Internet Protocol (IP) addresses to an entity, the method comprising:**

Zhang discloses a method for *mapping IP addresses to an entity* associated with distributing malware by querying passive DNS information. EX1003, ¶112. Zhang generally discloses various methods for creating an “association graph,” which includes “identifying IP addresses that have been frequently involved in malicious activities.” EX1006, ¶46. In particular, the “association graph” is used “to identify malicious *entities*” by associating “Internet Protocol (IP) address

information and related domain information” with said entities, *i.e.* ***mapping IP addresses to an entity*** involved with distributing a particular malware. *Id.*, ¶¶47, 34; *see also* ¶45 (“What are needed are new and improved techniques that can leverage the information in DNS to ***identify malicious entities***, such as malware domains.”).

Zhang further discloses the use of “passive DNS (pDNS) information” to build the “association graph.” *Id.*, ¶47. Zhang explains that passive DNS information involves “collect[ing] passively (e.g., by listening to DNS traffic) from distributed sources (e.g., using passive DNS sensors).” *Id.*, ¶47. Such information includes “DNS resource records that include domain to IP mapping information 912, Whois information 914, geographical IP information 916 (e.g., geo-based locations of IP addresses), and various other passive DNS and public DNS information 918 (e.g., DNS name servers and/or other DNS information)” and is stored in a “pDNS database.” *Id.*, ¶79; EX1003, ¶113.

McNab also discloses various techniques for querying DNS information to ***map IP addresses to an entity***, and as discussed above, it would have been obvious to incorporate such techniques in Zhang by querying passive DNS information as taught by Zhang. EX1003, ¶114. McNab is a well-known book published by O’Reilly called “Network Security Assessment.” McNab provides a chapter called “Internet Host and Network Enumeration” that teaches how to “query open sources

to identify hosts and networks of interest,” including “WHOIS databases, and DNS name servers.” EX1005, 5. By using these techniques, information such as “internal *IP addresses* gathered from DNS servers” and “insight into the target organization’s *DNS structure* (including domain names, subdomains, and hostnames)” can be collected.” *Id.* Such information can be collected even starting with “a single DNS domain name” of the target entity. *Id.*, 4.

(b) 1[a]: receiving a first domain name for the entity;

Zhang discloses *receiving a first domain name for the entity* associated with distributing malware. Zhang discloses receiving a domain name from “a malware sample,” for example, “GHI.COM.” EX1006, ¶70. Zhang discloses that the domain name “GHI.COM” can be used to “identify ... new malware IPs based on comes from, and resolves to relationships as shown in which correlations between malware samples with domains and IPs can be determined.” *Id.*; *see also id.*, ¶88 (“‘foo.com’ can be resolved to one IP and has one name server (NS).”); EX1003, ¶115.

(c) 1[b]: sending, to a domain name system (DNS) server, a first passive DNS query to identify first name servers for the first domain name;


1[c]: receiving, from the DNS server, a list of the first name servers for the first domain name;

Zhang in view of McNab renders obvious *querying a passive DNS server* by sending a *first domain name to identify and receive the name servers for a first domain name*. In particular, Zhang discloses using “passive DNS information” to

collect information about a “first domain,” *i.e.* **first domain name**. EX1006, ¶49. Zhang teaches that such “passive DNS information” is stored on a “passive DNS (pDNS) database,” *i.e.* **a DNS server**, which stores and returns a list of “**DNS Name Servers**” for a particular domain. *Id.*, ¶79; *see also* ¶80 (“pDNS DB 910 ... stores the DNS resource records (RRs) (e.g., in the format of a tuple (of RR name, RR type, IP, **Name Server (NS)**, Time To Live (TTL)) and the source of records”). For example, Zhang teaches that a first domain name like “foo.com” “can be resolved to one IP and has one **name server (NS)**,” and a POSITA would have understood that additional name servers would be returned for that domain if they existed. *Id.*, ¶88; EX1003, ¶116.

A POSITA would have understood and considered it obvious that the “passive DNS” information collected by Zhang, and in particular, “DNS resource records” and “Whois information,” maps a particular domain name to a list of authoritative **name servers** for resolving that domain name. EX1006, ¶¶79-80; EX1003, ¶117. It was well-known that DNS records and WHOIS information includes the authoritative name server of a particular domain. EX1003, ¶117. This is, for example, taught in McNab, which discloses that “WHOIS Registrars” can be queried to determine “[d]etails of **authoritative name servers** for each given domain.” EX1005, 21. McNab provides an example where one can “submit a query of blah.com, revealing useful information regarding the domain ... and authoritative

DNS name servers.” Id. The below shows a query where the *domain name* blah.com returns a list of three authoritative *name servers*, NS1.BLAH.COM, NS2.BLAH.COM, and NS3.BLAH.COM:



Id., 22.

McNab also teaches using the “nslookup” tool to query DNS resource records for a *list of name servers for a first domain name*. For example, McNab discloses that “[u]tilities such as nslookup ... are used to issue DNS requests relating to domains” to “retrieve DNS records” which includes “Authoritative DNS server information, from *Name Server* (NS) resource records.” EX1005, 30. McNab provides an example (excerpted below) of a query using the domain name “cia.gov” which “identifies the authoritative DNS servers as relay1.ucia.gov, relay7.ucia.gov, and auth100.ns.uu.net.” *Id.*, 31; EX1003, ¶118.

Example 3-6. Using *nslookup* to enumerate basic domain details

```
$ nslookup
> set querytype=any first domain name
> cia.gov
Server:      213.228.193.145
Address:     213.228.193.145#53

Non-authoritative answer:
Name:   cia.gov
Address: 198.81.129.100
cia.gov nameserver = relay7.ucia.gov.
cia.gov nameserver = auth100.ns.uu.net.
cia.gov nameserver = relay1.ucia.gov.
cia.gov
      origin = relay1.ucia.gov
      mail addr = root.ucia.gov
      serial = 511250020
      refresh = 7200
      retry = 3600
      expire = 604800
      minimum = 86400
cia.gov mail exchanger = 10 mail1.ucia.gov.
cia.gov mail exchanger = 5 mail2.ucia.gov.


list of first name servers

Authoritative answers can be found from:
relay7.ucia.gov internet address = 198.81.129.186
auth100.ns.uu.net      internet address = 198.6.1.202
relay1.ucia.gov internet address = 198.81.129.193
```


As discussed above, a POSITA would have considered it obvious to incorporate such teachings in McNab for querying WHOIS and DNS resource records stored in Zhang's passive DNS database to receive a list of authoritative name servers for a specified domain name in order to create an association graph as taught in Zhang. EX1003, ¶119.

As further evidence that this limitation is obvious and that there is motivation to combine Zhang and McNab, the Maltego tool described in the Technical Background also discloses sending a DNS query to retrieve a domain name's name server. EX1003, ¶¶120-21. For example, in Mr. Gates's presentation, he describes

the first step of Maltego is to “[f]ind our ... NS servers,” such as the NS server ns1.baytsp.net for the domain baytsp.net. EX1018, 70.



Similarly, in the YouTube video demonstrating Maltego, the “name servers (NS) for this domain” are retrieved. EX1021, 0:14. This retrieves two NS servers for the domain name “shmoo.com”: “krusty.shmoo.com” and “archimedes.shmoo.com.” *Id.*, 0:19.



(d) 1[d]: sending, for each of the first name servers, a second passive DNS query to identify second domain names for which the first name server is authoritative;

1[e]: receiving, for each of the first name servers, a list of the second domain names for which the first name server is authoritative;

Zhang in view of McNab renders obvious *querying a passive DNS server* by sending a *name server* to *identify and receive a list of additional domain names for which the name server is authoritative*. EX1003, ¶122. In particular, Zhang teaches querying “passive DNS information” stored in a “passive DNS (pDNS) database,” *i.e. a passive DNS server*, which stores and returns a list of *domain names* for a specified *authoritative name server*. EX1006, ¶79. For example, Zhang discloses

using passive DNS information to determine that “foo.com” ***“has the same name server as ‘a.com’ and ‘b.com,’ i.e. a list of the second domain names. Id., ¶88.*** Zhang also discloses determining “a reputation for a name server” based on “all the domain names resolved by the name server,” ***i.e. a list of second domain names for which the name server is authoritative. Id., ¶92.*** Passive DNS information can also “facilitate identification of bad DNS name servers” which are “name servers that are determined to resolve ***malware domains,***” ***i.e. a list of second domain names for which the name server is authoritative. Id., ¶96.*** Thus, Zhang teaches querying passive DNS information to map name servers retrieved from the first passive DNS query with a list of domain names for which it is authoritative. EX1003, ¶122.

A POSITA would have understood and considered it obvious that the “passive DNS information” collected by Zhang, including “DNS resource records” and “Whois information” which has Name Server information for domain names, maps a particular name server to a list of domain names for which it is authoritative. EX1006, ¶¶79-80. It was well-known that DNS records and Whois information records Name Server information for domain names. EX1003, ¶123. This is, for example, taught in McNab, which teaches using the “dig” command to query records of a previously obtained name server, which will reveal other domain names for which that name server is authoritative, which McNab refers to as “PTR record

enumeration,” where “PTR” refers to a “reverse DNS pointer” record stored on a DNS name server. EX1005, 34.

McNab provides an example below where dig is used to retrieve PTR records for the name server “relay2.ucia.gov.” The query returns a list of other domain names for which “relay2.ucia.gov” is authoritative, such as “intelligence.gov” and “nctc.gov.” EX1005, 34-35. McNab further teaches that this query is executed for “all accessible name servers.” *Id.*, 35; EX1003, ¶124.

Example 3-9. Using dig to perform a DNS zone transfer for a network block

```
$ dig @relay2.ucia.gov 129.81.198.in-addr.arpa axfr
; <<>> DiG 9.2.4 <<>> @relay2.ucia.gov 129.81.198.in-addr.arpa axfr
;; global options: printcmd
129.81.198.in-addr.arpa. 3600 IN      SOA      relay1.ucia.gov. root.ucia.gov.
509192750 7200 3600 604800 86400
129.81.198.in-addr.arpa. 3600 IN      NS       relay1.ucia.gov.
129.81.198.in-addr.arpa. 3600 IN      NS       relay7.ucia.gov.
129.81.198.in-addr.arpa. 3600 IN      NS       auth100.ns.uu.net.
015.129.81.198.in-addr.arpa. 3600 IN     PTR      wits01.nctc.gov.
100.129.81.198.in-addr.arpa. 3600 IN     PTR      cia.cia.gov.
101.129.81.198.in-addr.arpa. 3600 IN     PTR      www2.cia.gov.
103.129.81.198.in-addr.arpa. 3600 IN     PTR      www.intelligence.gov.
104.129.81.198.in-addr.arpa. 3600 IN     PTR      www.nctc.gov.
106.129.81.198.in-addr.arpa. 3600 IN     PTR      wits2.nctc.gov.
146.129.81.198.in-addr.arpa. 3600 IN     PTR      mail2out.ucia.gov.
148.129.81.198.in-addr.arpa. 3600 IN     PTR      mail2.ucia.
186.129.81.198.in-addr.arpa. 3600 IN     PTR      relay7.ucia
193.129.81.198.in-addr.arpa. 3600 IN     PTR      relay1.ucia
194.129.81.198.in-addr.arpa. 3600 IN     PTR      relay2.cia.gov.
195.129.81.198.in-addr.arpa. 3600 IN     PTR      relay2a.cia.gov.
196.129.81.198.in-addr.arpa. 3600 IN     PTR      relay2b.cia.gov.
20.129.81.198.in-addr.arpa. 3600 IN     PTR      ddss.cia.gov.
21.129.81.198.in-addr.arpa. 3600 IN     PTR      ddssdata.cia.gov.
22.129.81.198.in-addr.arpa. 3600 IN     PTR      ddsstest.cia.gov.
```


A POSITA would understand that the query will return domain names *for which the designed name server is authoritative*. EX1003, ¶125. This is because the dig command queries DNS records (specifically PTR records) for the designated

name server, for example, “relay2.ucia.gov.” The DNS records for the designated name server *must* include records for which it is authoritative—since the “authoritative name server” is required to be the name server of last resort for retrieving any DNS records for which it is authoritative. *See, e.g.*, EX1007, 5:24-27 (“If there is insufficient information at the local name server, a request is made to the root to find the authoritative name server, and the information request is forwarded to that name server.”). Thus, McNab teaches querying DNS records of a name server at least some of which will be records for domain names for which the name server is authoritative. EX1003, ¶125. A POSITA would further understand that the authoritative name server of the returned domain names can be confirmed by querying WHOIS or DNS records (including as stored in Zhang’s passive DNS server) as discussed in limitation 1[b]. *Id.*

As discussed above, a POSITA would have considered it obvious to incorporate such teachings in McNab to query, for each name server received from the step in limitation 1[c], its DNS resource records stored in Zhang’s passive DNS database to receive a list of domain names for which the specified name server is authoritative in order to create an association graph as taught in Zhang. EX1003, ¶126.

In the alternative, this limitation is disclosed under Patent Owner’s infringement theory wherein Patent Owner accuses the use of WHOIS queries as the

claimed “second passive DNS query.” EX1025, 12. As shown in the annotated image below from Patent Owner’s infringement contentions, Patent Owner points to “other domains” that were obtained from “WHOIS” queries as the “second passive DNS query.” *Id.*; EX1003, ¶127.




As discussed above for limitation 1[b], a WHOIS query involves specifying a name server to receive, among other information, the authoritative name servers for that designated domain name. Thus, it appears that Patent Owner believes these


limitations can be met by sending WHOIS queries for specified one or more *second domain names*, and then using the results of that query to confirm that *the first name server is authoritative*. EX1003, ¶128. It appears Patent Owner contends that performing these WHOIS queries will “identify second domain names for which the first name server is authoritative” by confirming that the specified domain names have an authoritative name server that matches the first name server. *Id.*

Under this theory, Zhang discloses and in view of McNab renders obvious these two limitations because as discussed above for limitation 1[b], Zhang discloses querying and storing WHOIS records, and McNab discloses performing a WHOIS query to obtain authoritative name server information. EX1003, ¶129.


As further evidence that this limitation is obvious and that there is motivation to combine Zhang and McNab, the Maltego tool described in the Technical Background also discloses sending a DNS query to retrieve domain names that share an authoritative name server. EX1003, ¶¶130-32. For example, in Mr. Gates’s presentation, he describes the second step of Maltego is to “[s]ee what hostnames are shared by the DNS servers,” which results in many more domain names such as “ee-world.com” and “trademarkspider.com.” EX1018, 71.



Similarly, in the YouTube video demonstrating Maltego, after the name servers are retrieved, for each name server, Maltego can retrieve “what other domains use these name servers.” EX1021, 0:19.



This results in domain names such as “kusia.net” and “legacyhigh.com.” *Id.*, 0:41.



- (e) **1[f]: sending, for each of the second domain names, a third passive DNS query to identify host names for the hosts of the second domain name and IP addresses for the host names;**

1[g]: receiving a list of the host names and the IP addresses for the host names; and

Zhang in view of McNab renders obvious *querying a passive DNS server* by sending a *domain name* to *identify and receive a list of host names and IP addresses for hosts of that domain name*. EX1003, ¶133. In particular, Zhang teaches that the association graph includes information about “*IP, hostname/domain* (e.g., web domain), URI, border gateway protocol (BGP) information, and/or various other source and related information to determine associations between malware, malware families, and various sources of such malware and malware samples, and are not limited to just a specific IP relationship between sources of malware samples.” EX1006, ¶33.

In addition, Zhang teaches that “passive DNS information” is stored on a “passive DNS (pDNS) database,” *i.e. a DNS server*, which stores and returns “domain to IP mapping information” for a particular domain. *Id.*, ¶79; *see also* ¶80 (“pDNS DB 910 ... stores the DNS resource records (RRs) (e.g., in the format of a tuple (of RR name, RR type, IP, Name Server (NS), Time To Live (TTL)) and the source of records”). A POSITA would have understood and considered it obvious that the “passive DNS” information collected by Zhang, and in particular, “DNS

resource records” and “Whois information,” maps a particular domain name to a list of hostnames for that domain name and their IP addresses. EX1006, ¶¶79-80; EX1003, ¶134.

This is, for example, taught in McNab, which discloses a method for querying DNS records called “Forward DNS Grinding” which uses a specified domain name to “enumerate valid DNS address records and aliases relating to the domain and its hosts,” *i.e. hosts of the domain name and their IP addresses*. EX1005, 35-36; EX1003, ¶135. McNab provides an example of “Forward DNS Grinding” using a tool called TXDNS “against the bankofengland.co.uk domain, using a small dictionary of common mail server names.” *Id.* In the below example, the tool returned a list of “three hosts that appear to be SMTP mail servers” and their IP addresses: mail.bankofengland.co.uk, mail2.bankofengland.co.uk, and mailhost.bankofengland.co.uk. *Id.*

Example 3-11. Forward DNS grinding to identify mail servers

```
C:\tools> txdns -f mail-dict.txt bankofengland.co.uk
```

TXDNS (http://www.txdns.net) 2.0.0 running STAND-ALONE Mode

> mail.bankofengland.co.uk	- 217.33.207.254
> mail2.bankofengland.co.uk	- 194.201.32.153
> mailhost.bankofengland.co.uk	- 194.201.32.130

Resolved names: 3
Failed queries: 95
Total queries: 98

second domain name

list of hosts and IP addresses

In this example, the intent was to enumerate mail server hosts. EX1003, ¶136.

McNab provides another example using TXDNS that enumerates multiple types of hosts against the domain name “apple.com” to “reveal valid apple.com hostnames” and their IP addresses. EX1005, 85-86.

Example 5-6. Forward DNS grinding using txdns.exe under Windows

```
C:\tools> txdns -f smalllist.txt -s 17.254.0.50 apple.com
```

TXDNS (http://www.txdns.net) 2.0.0 running STAND-ALONE Mode

> ftp.apple.com	- 17.254.16.10
> guide.apple.com	- 17.254.12.37
> help.apple.com	- 17.254.3.26
> mercury.apple.com	- 17.250.248.40
> research.apple.com	- 17.255.4.30
> search.apple.com	- 17.254.0.160
> vpn.apple.com	- 17.252.68.41
> webmail.apple.com	- 17.254.13.52

Resolved names: 8
Failed queries: 348
Total queries: 356

second domain name

list of hosts and IP addresses

McNab also provides an example using the previously-mentioned “dig” tool to perform a zone transfer to enumerate various hosts for a specified domain name, in this case, “ucia.gov” (which is not the same as the *first domain name* discussed above, “cia.gov”). EX1005, 33-34. As shown below, the list includes IP addresses for some of the hosts. EX1003, ¶137.

Example 3-8. Using dig to perform a DNS zone

second domain
name

\$ dig @relay2.ucia.gov ucia.gov axfr


```
; <>> DiG 9.2.4 <>> ucia.gov @relay2.ucia.gov axfr
;; global options: printcmd
ucia.gov. 3600 IN SOA relay1.ucia.gov. root.ucia.gov.
511210023 7200 900 604800 900
ucia.gov. 3600 IN NS relay1.ucia.gov.
ucia.gov. 3600 IN NS relay7.ucia.gov.
ucia.gov. 3600 IN NS auth100.ns.uu.net.
ucia.gov. 3600 IN MX 5 mail2.ucia.gov.
ain.ucia.gov. 3600 IN A 198.81.128.68
ain-relay.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
ain-relay-int.ucia.gov. 3600 IN CNAME ain-relay1-int.ucia.gov.
ain-relay1.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
ain-relay1-ext.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
ain-relay1-int.ucia.gov. 3600 IN A 192.168.64.2
ain-relay2.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
ain-relay2-ext.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
ain-relay2-int.ucia.gov. 3600 IN A 192.168.64.3
ain-relay7.ucia.gov. 3600 IN CNAME relay7.ucia.gov.
ain-relay7-ext.ucia.gov. 3600 IN CNAME relay7.ucia.gov.
ain-relay7-int.ucia.gov. 3600 IN A 192.168.64.67
ex-rtr.ucia.gov. 3600 IN CNAME ex-rtr-129.ucia.gov.
ex-rtr-129.ucia.gov. 3600 IN A 198.81.129.222
ex-rtr-129.ucia.gov. 3600 IN HINFO "Cisco 4000 Router" "NP-1E Board"
ex-rtr-191-a.ucia.gov. 3600 IN A 192.103.66.58
ex-rtr-191-b.ucia.gov. 3600 IN A 192.103.66.62
foia.ucia.gov. 3600 IN NS relay1.ucia.gov.
foia.ucia.gov. 3600 IN NS auth100.ns.uu.net.
mail1.ucia.gov. 3600 IN A 198.81.129.68
mail1out.ucia.gov. 3600 IN A 198.81.129.71
mail2.ucia.gov. 3600 IN A 198.81.129.148
mail2out.ucia.gov. 3600 IN A 198.81.129.146
relay.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
relay-int.ucia.gov. 3600 IN CNAME ain-relay1-int.ucia.gov.
relay1.ucia.gov. 3600 IN A 198.81.129.193
relay1-ext.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
relay1-int.ucia.gov. 3600 IN CNAME ain-relay1-int.ucia.gov.
relay2.ucia.gov. 3600 IN A 198.81.129.194
relay2-ext.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
relay2-int.ucia.gov. 3600 IN CNAME ain-relay2-int.ucia.gov.
relay2a.ucia.gov. 3600 IN A 198.81.129.200
relay2y.ucia.gov. 3600 IN A 198.81.129.68
relay2z.ucia.gov. 3600 IN A 198.81.129.69
relay7.ucia.gov. 3600 IN A 198.81.129.186
relay7-ext.ucia.gov. 3600 IN CNAME relay7.ucia.gov.
relay7a.ucia.gov. 3600 IN A 198.81.129.197
relay7b.ucia.gov. 3600 IN A 198.81.129.198
res.ucia.gov. 3600 IN A 198.81.129.116
wais.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
```

list of hosts and
IP addresses

As discussed above, a POSITA would have considered it obvious to incorporate such teachings in McNab to query, for each domain name with a shared authoritative name server received from the step in limitation 1[e], its DNS resource



records stored in Zhang's passive DNS database to receive a list of hosts and IP addresses for that specified domain name in order to create an association graph as taught in Zhang. EX1003, ¶138.

As further evidence that this limitation is obvious and that there is motivation to combine Zhang and McNab, the Maltego tool described in the Technical Background also discloses sending a DNS query to retrieve IP addresses of hostnames for any domain names that are found. EX1003, ¶¶139-40. For example, in Mr. Gates's presentation, he describes a third step of Maltego which can automatically "[b]rute force our domain names and resolve them to IPs," such as discovering hostnames such as "w2.baytsp.com" and "vpn.baytsp.com" and resolving their IP addresses. EX1018, 72.



Similarly, in the YouTube video demonstrating Maltego, for each domain name discovered, Maltego can further find the hostnames for each domain. EX1021, 5:16. Maltego also resolves the IP address for each hostname. *Id.*, 5:34.

Petition for *Inter Partes* Review
U.S. Patent No. 11,652,834



(f) **1[h]: mapping each IP address to an attribute for the entity.**

As discussed above for limitation 1[pre], Zhang discloses a method for *mapping IP addresses to an entity* associated with distributing malware, *i.e. an attribute for the entity*, by querying passive DNS information. EX1003, ¶141. Zhang generally discloses various methods for creating a “association graph,” which includes “identifying IP addresses that have been frequently involved in malicious activities,” *i.e. an attribute for the entity*. EX1006, ¶46. In particular, the “association graph” is used “to identify malicious *entities*” by associating “Internet Protocol (IP) address information and related domain information” with said entities, *i.e. mapping IP addresses to an entity* involved with distributing a particular malware. *Id.*, ¶¶47, 34; *see also* ¶45 (“What are needed are new and improved techniques that can leverage the information in DNS to *identify malicious entities*, such as malware domains.”).

Zhang further discloses the use of “passive DNS (pDNS) information” to build the “association graph.” *Id.*, ¶47. Zhang explains that passive DNS information involves “collect[ing] passively (e.g., by listening to DNS traffic) from distributed sources (e.g., using passive DNS sensors).” *Id.*, ¶47. Such information includes “DNS resource records that include domain to IP mapping information 912, Whois information 914, geographical IP information 916 (e.g., geo-based locations of IP addresses), and various other passive DNS and public DNS information 918

(e.g., DNS name servers and/or other DNS information)” and is stored in a “pDNS database.” *Id.*, ¶79; EX1003, ¶142.

Zhang also describes domain “reputations,” which is also an *attribute of an entity* and make use of the association graph, *i.e. mapping IP addresses to an attribute of an entity*. For example, Zhang describes that “a domain reputation score can be generated using various techniques described herein (e.g., based on passive DNS information and using a malware *association graph* that associates malware samples with malware source information including domains), and if a DNS reputation score exceeds a threshold, then the domain can be identified as a malware domain (e.g., and malware samples from that domain can be identified as malware).” EX1006, ¶48. Because the reputation score uses the association graph, as the graph is built by enumerating IP addresses for an entity, those IP addresses are also mapped to the reputation of the entity. EX1003, ¶143.

As also discussed above, McNab also discloses various techniques for querying DNS information to *map IP addresses to an entity*, and as discussed above, it would have been obvious to incorporate such techniques in Zhang by querying passive DNS information as taught by Zhang. EX1003, ¶144. McNab provides a chapter called “Internet Host and Network Enumeration” that teaches how to “query open sources to identify hosts and networks of interest,” including “WHOIS databases, and DNS name servers.” EX1005, 5. By using these techniques,

information such as “internal *IP addresses* gathered from DNS servers” and “insight into the target organization’s *DNS structure* (including domain names, subdomains, and hostnames) can be collected,” the DNS structure being *an attribute of an entity*. *Id.* Such information can be collected even starting with “a single DNS domain name” of the target entity. *Id.*, 4. McNab further teaches that the IP addresses resulting from the “dig” tool discussed above for limitation 1[f] “identife[s]” “IP address blocks” for the IP addresses, IP address blocks being *an attribute of an entity*. *Id.*, 34.

McNab also describes enumerating IP address blocks and mapping them to an entity using an entity’s “Classless Inter-Domain Routing” (CIDR) blocks. *Id.*, 46. In particular, a POSITA would understand that (CIDR) refers to an IP address block identified using the notation A.B.C.D/xx, wherein “xx” indicates the number of bits used for the subnet mask. EX1003, ¶145. For example, McNab discloses “Enumerating subnet network ... with Nmap” by specifying the CIDR block “154.14.224.0/26,” which indicates that the subnet mask is 26 bits long. *Id.*

3. Claim 2: The method of claim 1, wherein the first domain name is received from a user interface.

As discussed above in limitation 1[a], Zhang discloses receiving a *first domain name* from a malware sample. EX1006, ¶70. Zhang further discloses that the malware sample can be received *from a user interface*, such as “disguised as genuine software and can come from an official company *website*.” *Id.*, ¶22. Zhang

also discloses that malware can be received “as an *e-mail attachment*,” which a POSITA would understand involves an email user interface to view email messages and their attachments. EX1003, ¶146.

Alternatively, McNab discloses a user interface for receiving a domain name, in particular, a command line user interface for which a domain name can be entered. *See, e.g.*, EX1005, 31. It would have been obvious to incorporate McNab with Zhang as discussed above. EX1003, ¶147.

It also would have been obvious to use a user interface to allow user entry of a domain name or a website link containing a domain name. EX1003, ¶148. For example, user interfaces for performing IP address enumeration were well-known, such as the user interface of the Maltego tool as described above in the Technical Background and in Exhibits 1018 and 1021. It would have been obvious to have a user interface in Zhang to allow a user to enter a suspicious domain name or website in order to generate the association graph described by Zhang. EX1003, ¶148. A POSITA would have been motivated to incorporate such a user interface in order to provide a better user experience and additional methods for inputting suspicious domain names beyond obtaining them from suspected malware. *Id.* For example, it was known that an e-mail may contain a suspicious domain name or website as a phishing link, and a POSITA would have been motivated by such phishing attacks to allow manual entry of a domain name or website into Zhang’s system to generate

an association graph, as well as determine whether the suspicious domain name was associated with previously-analyzed domain names with low reputations. *Id.* A POSITA would have reasonably expected this combination to succeed because developing user interfaces was straightforward and well-within the capabilities of a POSITA. *Id.*

4. Claim 3: The method of claim 1, wherein the list of the host names comprises names of at least a subset of the first name servers.

As discussed above for limitation 1[d] and 1[e], Zhang in view of McNab renders obvious obtaining a list of host names. In one of the examples described in McNab for obtaining a list of host names, one of the names includes “relay7.ucia.gov.” EX1005, 33-34; EX1003, ¶149.

Example 3-8. Using dig to perform a DNS zone transfer of the second domain name

\$ dig @relay2.ucia.gov ucia.gov axfr

```
; <>> DiG 9.2.4 <>> ucia.gov @relay2.ucia.gov axfr
;; global options: printcmd
ucia.gov. 3600 IN SOA relay1.ucia.gov. root.ucia.gov.
511210023 7200 900 604800 900
ucia.gov. 3600 IN NS relay1.ucia.gov.
ucia.gov. 3600 IN NS relay7.ucia.gov.
ucia.gov. 3600 IN NS auth100.ns.uu.net.
ucia.gov. 3600 IN MX 5 mail2.ucia.gov.
ain.ucia.gov. 3600 IN A 198.81.128.68
ain-relay.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
ain-relay-int.ucia.gov. 3600 IN CNAME ain-relay1-int.ucia.gov.
ain-relay1.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
ain-relay1-ext.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
ain-relay1-int.ucia.gov. 3600 IN A 192.168.64.2
ain-relay2.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
ain-relay2-ext.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
ain-relay2-int.ucia.gov. 3600 IN A 192.168.64.3
ain-relay7.ucia.gov. 3600 IN CNAME relay7.ucia.gov.
ain-relay7-ext.ucia.gov. 3600 IN CNAME relay7.ucia.gov.
ain-relay7-int.ucia.gov. 3600 IN A 192.168.64.67
ex-rtr.ucia.gov. 3600 IN CNAME ex-rtr-129.ucia.gov.
ex-rtr-129.ucia.gov. 3600 IN A 198.81.129.222
ex-rtr-129.ucia.gov. 3600 IN HINFO "Cisco 4000 Router" "NP-1E Board"
ex-rtr-191-a.ucia.gov. 3600 IN A 192.103.66.58
ex-rtr-191-b.ucia.gov. 3600 IN A 192.103.66.62
foia.ucia.gov. 3600 IN NS relay1.ucia.gov.
foia.ucia.gov. 3600 IN NS auth100.ns.uu.net.
mail1.ucia.gov. 3600 IN A 198.81.129.68
mail1out.ucia.gov. 3600 IN A 198.81.129.71
mail2.ucia.gov. 3600 IN A 198.81.129.148
mail2out.ucia.gov. 3600 IN A 198.81.129.146
relay.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
relay-int.ucia.gov. 3600 IN CNAME ain-relay1-int.ucia.gov.
relay1.ucia.gov. 3600 IN A 198.81.129.193
relay1-ext.ucia.gov. 3600 IN CNAME relay1.ucia.gov.
relay1-int.ucia.gov. 3600 IN CNAME ain-relay1-int.ucia.gov.
relay2.ucia.gov. 3600 IN A 198.81.129.194
relay2-ext.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
relay2-int.ucia.gov. 3600 IN CNAME ain-relay2-int.ucia.gov.
relay2a.ucia.gov. 3600 IN A 198.81.129.200
relay2y.ucia.gov. 3600 IN A 198.81.129.68
relay2z.ucia.gov. 3600 IN A 198.81.129.69
relay7.ucia.gov. 3600 IN A 198.81.129.186
relay7-ext.ucia.gov. 3600 IN CNAME relay7.ucia.gov.
relay7a.ucia.gov. 3600 IN A 198.81.129.197
relay7b.ucia.gov. 3600 IN A 198.81.129.198
res.ucia.gov. 3600 IN A 198.81.129.116
wais.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
```

list of hosts and
IP addresses


As discussed above in limitation 1[b], “relay7.ucia.gov” is one of the name servers for the domain “cia.gov.” EX1005, 31; EX1003, ¶150.

Example 3-6. Using nslookup to enumerate basic domain details

```
$ nslookup
> set querytype=any first domain name
> cia.gov
Server:      213.228.193.145
Address:     213.228.193.145#53

Non-authoritative answer:
Name:   cia.gov
Address: 198.81.129.100
cia.gov nameserver = relay7.ucia.gov.
cia.gov nameserver = auth100.ns.uu.net.
cia.gov nameserver = relay1.ucia.gov.
cia.gov
      origin = relay1.ucia.gov
      mail addr = root.ucia.gov
      serial = 511250020
      refresh = 7200
      retry = 3600
      expire = 604800
      minimum = 86400
cia.gov mail exchanger = 10 mail1.ucia.gov.
cia.gov mail exchanger = 5 mail2.ucia.gov.

Authoritative answers can be found from:
relay7.ucia.gov internet address = 198.81.129.186
auth100.ns.uu.net      internet address = 198.6.1.202
relay1.ucia.gov internet address = 198.81.129.193
```



5. Claim 4: The method of claim 1, wherein the attribute is a Classless Inter-Domain Routing block.

As discussed above for limitation 1[h], McNab discloses (and Zhang in view of McNab renders obvious) that the IP address information is mapped to an “IP address block.” EX1005, 34. McNab further discloses that IP address blocks can be “Classless Inter-Domain Routing” (CIDR) blocks. *Id.*, 46. In particular, a POSITA would understand that (CIDR) refers to an IP address block identified using the notation A.B.C.D/xx, wherein “xx” indicates the number of bits used for the

subnet mask. EX1003, ¶151. For example, McNab discloses “Enumerating subnet network ... with Nmap” by specifying the CIDR block “154.14.224.0/26,” which indicates that the subnet mask is 26 bits long. *Id.*

6. Claim 5: The method of claim 1, wherein receiving the first domain name for the entity comprises: determining the first domain name for a website of the entity.

Zhang discloses “determining that a domain (e.g. a *web domain*) is associated with malware based on an association with a cluster is performed.” EX1006, ¶77. In other words, Zhang discloses that the *first domain name* that is received from malware and for which an association graph is generated can be a “web domain,” *i.e. a website of the entity. Id.*; EX1003, ¶152.

Alternatively, as discussed above for Claim 2, it would have been obvious to allow Zhang’s user interface to allow for the manual entry of a website, which would include determining a domain name from said website. EX1003, ¶153.

7. Claim 6: The method of claim 1, wherein the method is executed at least two separate instances to track a change over time in the IP addresses for the host names.

Zhang discloses that “domain reputations are determined using passive DNS techniques” and are “evaluated by discovering relationships or associations with other domains, ... include[ing] common host IP, common name server, and/or common originating session.” EX1006, ¶48. Zhang further describes “re-evaluat[ing] the reputation of domains given new data to maintain and improve

malware domain detection accuracy.” *Id.* In other words, Zhang discloses generating a domain reputation based on the IP address of the host names in that domain, and further discloses re-evaluating the reputation given new data for the inputs of the reputation score, including new data regarding the IP addresses of the host names. EX1003, ¶154. The update of the reputation score is therefore ***tracking a change over time in the IP addresses for the host names***. Thus, Zhang discloses this limitation.

In addition, McNab discloses that its “Network Enumeration” techniques, which as discussed for limitation 1[f] include deriving IP addresses for hostnames, are repeated on a cyclic basis. EX1005, 8. Indeed, McNab discloses that “an analyst may identify previously unknown IP address blocks, which can then be fed back into the network enumeration process to identify further network components.” *Id.*, 9. Thus, McNab discloses executing the network enumeration process multiple times to track changes in IP addresses, including for the hostnames of a domain. EX1003, ¶155. As discussed above, it would have been obvious to include these teachings in Zhang’s system.

8. **Claim 7: The method of claim 6, further comprising: automatically updating the mapping of IP addresses upon identifying the change in the IP addresses.**

As discussed above for Claim 6, both Zhang and McNab disclose updating enumerations of IP addresses upon identifying a change in the IP address. EX1003, ¶156.

9. Claim 8: The method of claim 1, further comprising: identifying a shared hosting service based on the mapped IP addresses.

Zhang discloses that “[s]everal types of relations are used to connect two domains ... based on the information in pDNS” such as “if domain A and domain B are mapped to the same IP, then these domains can be deemed related.” EX1006, ¶86. Deeming a domain as related is *identifying a shared hosting service* because it identifies an IP address that is being shared for multiple domains that host malware. EX1003, ¶157. Zhang further describes detecting that “an IP [is] used for domain hosting” using “a filtering approach to identify IP/name server that are used for web hosting/domain parking/CDN.” EX1006, ¶87. Domain hosting, web hosting and CDN are all examples of a *shared hosting service* that is identified based on the mapped IP addresses. EX1003, ¶157.

10. Claim 9: The method of claim 1, further comprising: determining a security rating for the entity based on the mapped IP address.

Zhang discloses a *security rating* in the form of a *domain reputation*. Specifically, Zhang discloses that “domain reputations are determined using passive DNS techniques” and are “evaluated by discovering relationships or associations

with other domains, ... include[ing] common host IP, common name server, and/or common originating session.” EX1006, ¶48; EX1003, ¶158.

11. Claim 10: The method of claim 1, wherein the entity is a first entity, and wherein mapping each IP address to the attribute for the entity comprises: determining that the IP address is mapped to attributes for the first entity and a second entity.

Zhang discloses that “passive DNS information can be analyzed and correlated with other information, such as a malware association graph, to identify malicious *entities* [*i.e. a first and second entity*], such as malware domains.” EX1006, ¶47. Zhang also discloses that “malware domain detection using passive DNS further includes providing a framework that monitors *different entities* in DNS and evaluates the reputations of such entities.” *Id.*, ¶48. Specifically, “the framework can improve the detection accuracy of malicious domains, and the framework can also monitor the malicious activities *based on passive DNS information* (e.g., passive DNS data).” *Id.*

As discussed above in claim 1, the building of the “association graph” involves using passive DNS information to map each IP address to an entity or attribute of an entity. EX1003, ¶¶159-60. In view of the above teachings, Zhang discloses that this process involves using this “passive DNS information” to map each address to one or more entities. *Id.* A POSITA would understand that this could involve noting that certain domain names in the list of second domain names are associated with an entity different from the entity of the first domain name, and

therefore the hostnames and IP addresses of those hostnames would be mapped to the new entity. *Id.* A POSITA would also consider it obvious to do this in order to give effect to Zhang’s teaching of “monitor[ing] different entities” using “passive DNS information.” EX1006, ¶58.

12. Claim 11: The method of claim 10, further comprising: determining a first security rating for the first entity and a second security rating for the second entity, the first security rating and the second security rating being related.

As discussed above in claim 10, Zhang discloses that “malware domain detection using passive DNS further includes providing a framework that monitors *different entities* in DNS and evaluates *the reputations* of such entities.” *Id.*, ¶48. Specifically, “the framework can improve the detection accuracy of malicious domains, and the framework can also monitor the malicious activities *based on passive DNS information* (e.g., passive DNS data).” *Id.* Thus, the reputations of such entities are a *security rating*, since they reflect the association of the domain with malware. EX1003, ¶161. Zhang further describes “determin[ing] relationships and associations among various entities in DNS,” and therefore the reputations of the different entities are *determined to be related*. *Id.*, ¶78.

13. Claims 12-22

Claims 12 through 22 mirror the above claims as set forth in the following chart. Zhang-McNab render these claims obvious for the same reasons as above. EX1003, ¶¶162-81.

Claim Limitation	Corresponding Claim Limitation(s)
12[pre]	1[pre]
12[c]	1[a]
12[d]	1[b]
12[e]	1[c]
12[f]	1[d]
12[g]	1[e]
12[h]	1[f]
12[i]	1[g]
12[j]	1[h]
13	2
14	3
15	4
16	5
17	6
18	7
19	8
20	9
21	10
22	11

Limitations 12[a] and 12[b] further recite “memory storing computer-executable instructions” and a “processor for executing the instructions.” Zhang discloses that its methods are stored as instructions in memory and executed by a processor. EX1003, ¶163. Specifically, Zhang discloses that its “invention can be implemented ... as a processor configured to execute instructions stored on and/or provided by a memory coupled to the processor.” EX1006, ¶19.

VIII. *FINTIV*

Fintiv weighs against denial. A stay has not been requested, trial has not been scheduled, discovery has not yet started, and the Court has not even set forth any case schedule. In addition, the Petition challenges more claims than Patent Owner asserted in its complaint. Moreover, the Petition presents compelling grounds of unpatentability. Thus, all of the *Fintiv* factors weigh against denial.

IX. MANDATORY NOTICES AND FEES

A. Real Party-In-Interest

The real party-in-interest is Normshield Inc. d/b/a Black Kite Inc.

B. Related Matters

To the best of Petitioner’s knowledge, the ’834 patent is asserted in the district court case *BitSight Technologies, Inc. v. NormShield Inc.*, 1:23-cv-12055 (D. Mass.)

C. Counsel and Service Information

Petitioner provides the following counsel and service information.

Petitioner consents to electronic service the email addresses listed in the table

below. Pursuant to 37 C.F.R. § 42.10(b), Power of Attorney accompany this
Petition.

LEAD COUNSEL	BACK-UP COUNSEL
James M. Glass (Reg. No. 46729) jimglass@quinnemanuel.com QUINN EMANUEL URQUHART & SULLIVAN, LLP 51 Madison Avenue, 22 nd Floor New York, NY 10010 Tel: (212) 849-7000 Fax: (212) 849-7100	David Elihu (Reg. No. 65,717) davideliu@quinnemanuel.com QUINN EMANUEL URQUHART & SULLIVAN, LLP 865 S. Figueroa St., 10th Floor Los Angeles, California 90017 Tel : (213) 443-3000 Fax : (213) 443-3000 Quincy Lu (Reg. No. 76,954) quincylu@quinnemanuel.com QUINN EMANUEL URQUHART & SULLIVAN, LLP 1109 First Avenue, Suite 210 Seattle, WA 98101 Tel: (206) 905-7000 Fax: (206) 905-7100 Brett Watkins (Reg. No. 60,458) brettwatkins@quinnemanuel.com QUINN EMANUEL URQUHART & SULLIVAN, LLP 700 Louisiana St., Suite 3900 Houston, TX 77002 Tel: (713) 221-7000 Fax: (713) 221-7100 Andrew Bramhall (<i>pro hac vice</i> forthcoming) andrewbramhall@quinnemanuel.com QUINN EMANUEL URQUHART

	& SULLIVAN, LLP 555 Twin Dolphin Dr., 5th Floor Redwood Shores, California 94065 Tel: (650) 801-5000 Fax: (650) 801-5100
--	--

D. Payment of Fees

The undersigned authorizes the Office to charge the fee required for this Petition for *Inter Partes* Review to Deposit Account No. 50-5708. Any additional fees that might be due are also authorized.

Date: September 6, 2024

Respectfully submitted,

By: /s/ James M. Glass
James M. Glass

Counsel for Petitioner

**CERTIFICATE OF COMPLIANCE WITH
TYPE-VOLUME LIMITATION, TYPEFACE REQUIREMENTS,
AND TYPE STYLE REQUIREMENTS**

1. This Petition complies with the type-volume limitation of 14,000 words, comprising 10,538 words, as counted using the Microsoft Word software that was used to prepare this paper, excluding the parts exempted by 37 C.F.R.

§ 42.24(a).

2. This Petition complies with the general format requirements of 37 C.F.R. § 42.6(a) and has been prepared using Microsoft® Word 2016 in 14-point Times New Roman.

Date: September 6, 2024

By: /s/ James M. Glass

James M. Glass

QUINN EMANUEL URQUHART
& SULLIVAN, LLP

51 Madison Avenue, 22nd Floor

New York, NY 10010

Tel: (212) 849-7000

Fax: (212) 849-7100

Counsel for Petitioner

CERTIFICATE OF SERVICE

The undersigned hereby certifies that on September 6, 2024, true and correct copies of the foregoing document and supporting materials were served in its entirety on the Patent Owner at the following address of record as listed on PAIR via Priority Mail Express® or Express Mail:

GOODWIN PROCTER LLP
Patent Administrator
100 Northern Avenue
Boston, MA 02210
United States

Official Correspondence Address

The undersigned hereby certifies that on September 6, 2024, true and correct copies of the foregoing document and supporting materials were served in its entirety on the Patent Owner at the following address of record via Priority Mail Express® or Express Mail:

Hunter Malboeuf
Goodwin Procter, LLP
100 Northern Avenue
Boston, MA 02210
United States

Courtesy copies were also sent via electronic mail to Patent Owner's counsel of record in the related district court proceeding:

Douglas J. Kline
Goodwin Procter, LLP
100 Northern Avenue
Boston, MA 02210
617-570-1000
Email: dkline@goodwinlaw.com
ATTORNEY TO BE NOTICED

Kelly Grosshuesch
Goodwin Procter LLP
1900 N Street, N.W.

Washington, DC 20036
202-346-4000
Fax: 202-346-4444
Email: KGrosshuesch@goodwinlaw.com
ATTORNEY TO BE NOTICED

Naomi Birbach
Goodwin Procter LLP
The New York Times Building
620 Eighth Avenue
New York City, NY 10018
212-459-7374
Email: nbirbach@goodwinlaw.com
ATTORNEY TO BE NOTICED

Robert D. Carroll
Goodwin Procter, LLP
100 Northern Avenue
Boston, MA 02210
617-570-1000
Fax: 617-523-1231
Email: rcarroll@goodwinlaw.com
ATTORNEY TO BE NOTICED

Timothy Keegan
Goodwin Procter LLP
620 Eighth Avenue
New York, NY 10018
212-459-7022
Email: tkeegan@goodwinlaw.com
ATTORNEY TO BE NOTICED

Date: September 6, 2024

Petition for *Inter Partes* Review

U.S. Patent No. 11,652,834

By: /s/ *James M. Glass*

James Glass

QUINN EMANUEL URQUHART

& SULLIVAN, LLP

51 Madison Avenue, 22nd Floor

New York, NY 10010

Tel: (212) 849-7000

Fax: (212) 849-7100

Counsel for Petitioner