

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. IPR2024-01394

U.S. Patent No. 9,438,615

**PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 9,438,615**

TABLE OF CONTENTS

	<u>Page</u>
I. INTRODUCTION	1
II. STANDING	2
III. GROUNDS	2
IV. POSITA	2
V. BACKGROUND	3
1. Overview of the Domain Name System (DNS) (EX1003, ¶¶49-54)	3
2. Open Source Intelligence (“OSINT”) Tools (EX1003, ¶¶55-77)	4
B. Public Availability and Authenticity	17
1. McNab	17
2. Internet Archive Pages	18
3. Public Availability of Maltego Videos	19
C. ’615 patent (EX1003, ¶¶78-80)	21
1. Claim Mapping	22
2. Prosecution History	23
VI. CLAIM CONSTRUCTION	24
VII. GROUND 1: LONGO, MCNAB	24
A. Motivation to Combine	24
B. Claim 1	27
1. 1[pre], 1[a]: A method comprising: generating a map between (a) technical assets that contribute to security characteristics of respective entities and (b) the identities of the entities that are associated with the respective technical assets, at least part of the generating of the map being done automatically, ...	27
2. 1[b]: generating of the map comprising online discovery of information about the technical assets; and	29

3.	1[c]: enabling a user to assist in the generating of the map by presenting to the user through a user interface (a) data about the technical assets of entities and (b) an interactive tool for associating the technical assets with the identities of the entities,.....	30
4.	1[d]: the technical assets comprising network-related information,.....	31
5.	1[e]: and in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or Regional Internet Registry.	31
C.	Claim 2	32
D.	Claim 3: The method of claim 2 in which the network-related information comprises at least one of Internet Protocol addresses, ...	32
E.	Claim 4	32
F.	Claim 9: The method of claim 1 comprising presenting to the user non-technical-asset information through the user interface. Claim 10: The method of claim 9 in which the non-technical-asset information comprises information about the entities. Claim 11: The method of claim 10 in which the information about the entities comprises at least one of descriptions of entities	32
G.	Claim 12: The method of claim 1 in which the tool enables the user to assign a technical asset to an entity.	33
H.	Claim 13: The method of claim 1 comprising invoking external crowd-sourced services to aid in generating the map.	33
I.	Claim 14: The method of claim 1 comprising enabling a separate review and approval of entities proposed to be included in the map.	33
J.	Claim 16: The method of claim 1 in which the part of the generating of the map that is done automatically comprises ... associating technical assets with entities	34
K.	Claim 33-37, 42-46, 48, 50-54, 59-63, 65.....	34
VIII.	GROUND 2: LONGO, MCNAB, MALTEGO.....	34
A.	Public Availability of Maltego Documents.....	35

B.	Motivation to Combine	35
C.	1[c].....	36
D.	Claims 9-11	42
E.	Claim 17: The method of claim 1 comprising generating graphs of relationships among entities based on their associations with technical assets.	43
F.	Remaining Limitations From Claim 1 and Claims 2-4, 12-14, and 16	44
G.	Claim 33-37, 42-46, 48-49, 50-54, 59-63, 65-66, 84-88, 93-98, 100	44
IX.	GROUND 3: LONGO, MCNAB, AND ZHANG (WITH OR WITHOUT MALTEGO).....	45
A.	Motivation to Combine	45
B.	Claim 5: The method of claim 4 in which the information about the technical assets is discovered through passive DNS queries.	50
C.	Claim 6: The method of claim 5 comprising identifying from the passive DNS queries associations between domain names and network addresses.....	51
D.	Claim 7: The method of claim 6 comprising tracking changes over time in the network addresses that are associated with a given domain name.	51
E.	Claim 8	51
1.	8[pre]: The method of claim 6 in which identifying from the passive DNS queries associations between domain names and network addresses comprises:.....	51
2.	8[a]: receiving a first domain name for an entity from the entities;	52
3.	8[b]: sending a first passive DNS query to identify first name servers for the first domain name;.....	52
	8[c]: receiving a list of the first name servers for the first domain;	52
4.	8[d]: sending a second passive DNS query, for each of the first name servers, to identify second domain names for which the name server is authoritative;	57

8[e]: receiving a list, for each of the first name servers, of the second domain names for which the name server is authoritative;	57
5. 8[f]: sending a third passive DNS query, for each of the second domain names, to identify host names for the hosts of the second domain name and Internet Protocol addresses for the host names;.....	65
8[g]: receiving a list of the host names and the Internet Protocol addresses for the host names; and.....	65
6. 8[h]: mapping each of the Internet Protocol addresses to an attribute for the entity.....	71
F. Claim 13	72
G. Claim 17	72
H. Claims 18-29, 31-32, 33-46, 48-49, 55-58, 63, 66, 84-98, 100	73
X. GROUND 4: LONGO, MCNAB, AND TIPPETT (WITH OR WITHOUT MALTEGO).....	73
A. Motivation to Combine	73
B. Claim 15: The method of claim 1 comprising providing the map to an application for joining to event data or scoring the security state of the entities.	75
C. Claim 47, 64, 67-71, 76-83, 99	76
XI. GROUND 5: LONGO, MCNAB, ZHANG, AND TIPPETT (WITH OR WITHOUT MALTEGO).....	76
XII. GROUND 6: SHULL	77
A. Claim 1	77
1. 1[pre], 1[a]	77
2. 1[b]	78
3. 1[c]	79
4. 1[d]	80
5. 1[e]	80
B. Claim 2	80
C. Claim 3	80

D.	Claim 4	81
E.	Claims 9-11	81
F.	Claim 12	81
G.	Claim 13	81
H.	Claim 14	82
I.	Claim 15	82
J.	Claim 16	82
K.	Claims 33-37, 42-48, 50-54, 59-65, 67-71, 76-82	83
XIII.	GROUND 7: SHULL, MCNAB, ZHANG	83
A.	Motivation to Combine	83
B.	Claims 5-8, 13, 17	87
C.	Claims 18-32, 33-49, 55-58, 63, 66, 72-75, 80, 83, 84-100.....	87
XIV.	<i>FINTIV</i>	87
XV.	MANDATORY NOTICES AND FEES	88
A.	Real Party-In-Interest	88
B.	Related Matters.....	88
C.	Counsel and Service Information	88
D.	Payment of Fees	89

EXHIBIT LIST

Ex.	Description
1001	U.S. Patent No. 9,438,615 (“‘615 patent”)
1002	Prosecution History for U.S. Patent No. 9,438,615
1003	Declaration of Kevin Almeroth, Ph.D.
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 § 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	Stoneburgner, et al., “Risk Management Guide for Information Technology Systems,” NIST, available at https://www.archives.gov/files/era/recompete/sp800-30.pdf
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	Chris Gates, “Toorcon X Gates: New School Information Gathering,” available at http://vimeo.com/2745624
1020	Chris 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	“Maltego User Guide” available at https://web.archive.org/web/20120606172056/http://ctas.paterva.com/view/Userguide
1026	U.S. Pat. Pub. No. 2011/0106920 A1 (“Longo”) (prior art under at least AIA §102(a)(1))
1027	’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)
1028	U.S. Pat. Pub. No. 2005/0278786 A1 (“Tippett”) (prior art under at least AIA §102(a)(1))
1029	August 29, 2024 Email from Melissa Nezhnik
1030	Metadata for YouTube video ID qiv4-wy3mxo
1031	Relevant screen captures of Exhibit 1021

1032	U.S. Pat. Pub. No. 2006/0230039 A1 (“Shull”) (prior art under at least AIA §102(a)(1))
------	--

LIST OF CHALLENGED CLAIMS

Claim	Limitation
1[pre]	A method comprising:
1[a]	generating a map between (a) technical assets that contribute to security characteristics of respective entities and (b) the identities of the entities that are associated with the respective technical assets, at least part of the generating of the map being done automatically,
1[b]	generating of the map comprising online discovery of information about the technical assets; and
1[c]	enabling a user to assist in the generating of the map by presenting to the user through a user interface (a) data about the technical assets of entities and (b) an interactive tool for associating the technical assets with the identities of the entities,
1[d]	the technical assets comprising network-related information,
1[e]	and in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or Regional Internet Registry.
2	The method of claim 1 in which the technical assets comprise network-related information.
3	The method of claim 2 in which the network-related information comprises at least one of Internet Protocol addresses, blocks of Internet Protocol addresses, mail server configurations, domain names, social media handles, third-party data hosting, internet service providers, Domain Name System services, Autonomous System numbers, and Border Gateway Protocol advertisements.

Claim	Limitation
4	The method of claim 2 in which generating the map comprises online discovery of information about the technical assets.
5	The method of claim 4 in which the information about the technical assets is discovered through passive DNS queries.
6	The method of claim 5 comprising identifying from the passive DNS queries associations between domain names and network addresses.
7	The method of claim 6 comprising tracking changes over time in the network addresses that are associated with a given domain name.
8[pre]	The method of claim 6 in which identifying from the passive DNS queries associations between domain names and network addresses comprises:
8[a]	receiving a first domain name for an entity from the entities;
8[b]	sending a first passive DNS query to identify first name servers for the first domain name;
8[c]	receiving a list of the first name servers for the first domain;
8[d]	sending a second passive DNS query, for each of the first name servers, to identify second domain names for which the name server is authoritative;
8[e]	receiving a list, for each of the first name servers, of the second domain names for which the name server is authoritative;

Claim	Limitation
8[f]	sending a third passive DNS query, for each of the second domain names, to identify host names for the hosts of the second domain name and Internet Protocol addresses for the host names;
8[g]	receiving a list of the host names and the Internet Protocol addresses for the host names; and
8[h]	mapping each of the Internet Protocol addresses to an attribute for the entity.
9	The method of claim 1 comprising presenting to the user non-technical-asset information through the user interface.
10	The method of claim 9 in which the non-technical-asset information comprises information about the entities.
11	The method of claim 10 in which the information about the entities comprises at least one of descriptions of entities, industry classifications, and employee counts.
12	The method of claim 1 in which the tool enables the user to assign a technical asset to an entity.
13	The method of claim 1 comprising invoking external crowd-sourced services to aid in generating the map.
14	The method of claim 1 comprising enabling a separate review and approval of entities proposed to be included in the map.

Claim	Limitation
15	The method of claim 1 comprising providing the map to an application for joining to event data or scoring the security state of the entities.
16	The method of claim 1 in which the part of the generating of the map that is done automatically comprises at least one of collecting information about technical assets and about entities, associating technical assets with entities, and approving proposed portions of the map.
17	The method of claim 1 comprising generating graphs of relationships among entities based on their associations with technical assets.
18 <pre>[pre]</pre>	A method comprising:
18[a]	generating a map between (a) technical assets that contribute to security characteristics of respective entities and (b) the identities of the entities that are associated with the respective technical assets, at least part of the generating of the map being done automatically,
18[b]	the technical assets comprising network-related information,
18[c]	the generating of the map comprising online discovery of information about the technical assets,
18[d]	the information about the technical assets being discovered through passive DNS queries; and
18[e]	enabling a user to assist in the generating of the map by presenting to the user through a user interface (a) data about the

Claim	Limitation
	technical assets of entities and (b) an interactive tool for associating the technical assets with the identities of the entities.
19	The method of any of claim 18 in which the network-related information comprises at least one of Internet Protocol addresses, blocks of Internet Protocol addresses, mail server configurations, domain names, social media handles, third-party data hosting, internet service providers, Domain Name System services, Autonomous System numbers, and Border Gateway Protocol advertisements.
20	The method of claim 18 in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or a Regional Internet Registry.
21	The method of claim 18 comprising identifying from the passive DNS queries associations between domain names and network addresses.
22	The method of claim 21 comprising tracking changes over time in the network addresses that are associated with a given domain name.
23[pre]	The method of claim 21 in which identifying from the passive DNS queries associations between domain names and network addresses comprises:
23[a]	receiving a first domain name for an entity from the entities;

Claim	Limitation
23[b]	sending a first passive DNS query to identify first name servers for the first domain name;
23[c]	receiving a list of the first name servers for the first domain;
23[d]	sending a second passive DNS query, for each of the first name servers, to identify second domain names for which the name server is authoritative;
23[e]	receiving a list, for each of the first name servers, of the second domain names for which the name server is authoritative;
23[f]	sending a third passive DNS query, for each of the second domain names, to identify host names for the hosts of the second domain name and Internet Protocol addresses for the host names;
23[g]	receiving a list of the host names and the Internet Protocol addresses for the host names; and
23[h]	mapping each of the Internet Protocol addresses to an attribute for the entity.
24	The method of claim 18 comprising presenting to the user non-technical-asset information through the user interface.
25	The method of claim 24 in which the non-technical-asset information comprises information about the entities.
26	The method of claim 25 in which the information about the entities comprises at least one of descriptions of entities, industry classifications, and employee counts.

Claim	Limitation
27	The method of claim 18 in which the tool enables the user to assign a technical asset to an entity.
28	The method of claim 18 comprising invoking external crowd-sourced services to aid in generating the map.
29	The method of claim 18 comprising enabling a separate review and approval of entities proposed to be included in the map.
30	The method of claim 18 comprising providing the map to an application for joining to event data or scoring the security state of the entities.
31	The method of claim 18 in which the part of the generating of the map that is done automatically comprises at least one of collecting information about technical assets and about entities, associating technical assets with entities, and approving proposed portions of the map.
32	The method of claim 18 comprising generating graphs of relationships among entities based on their associations with technical assets.
33[pre]	A method comprising:
33[a]	generating a map between (a) technical assets that contribute to security characteristics of respective entities and (b) the identities of the entities that are associated with the respective technical

Claim	Limitation
	assets, at least part of the generating of the map being done automatically,
33[b]	invoking external crowd-sourced services to aid in generating the map; and
33[c]	enabling a user to assist in the generating of the map by presenting to the user through a user interface (a) data about the technical assets of entities and (b) an interactive tool for associating the technical assets with the identities of the entities.
34	The method of claim 33 in which the technical assets comprise network-related information.
35	The method of any of claim 34 in which the network-related information comprises at least one of Internet Protocol addresses, blocks of Internet Protocol addresses, mail server configurations, domain names, social media handles, third-party data hosting, internet service providers, Domain Name System services, Autonomous System numbers, and Border Gateway Protocol advertisements.
36	The method of claim 34 in which generating the map comprises online discovery of information about the technical assets.
37	The method of claim 36 in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or a Regional Internet Registry.

Claim	Limitation
38	The method of claim 36 in which the information about the technical assets is discovered through passive DNS queries.
39	The method of claim 38 comprising identifying from the passive DNS queries associations between domain names and network addresses.
40	The method of claim 39 comprising tracking changes over time in the network addresses that are associated with a given domain name.
41[pre]	The method of claim 39 in which identifying from the passive DNS queries associations between domain names and network addresses comprises:
41[a]	receiving a first domain name for an entity from the entities;
41[b]	sending a first passive DNS query to identify first name servers for the first domain name;
41[c]	receiving a list of the first name servers for the first domain;
41[d]	sending a second passive DNS query, for each of the first name servers, to identify second domain names for which the name server is authoritative;
41[e]	receiving a list, for each of the first name servers, of the second domain names for which the name server is authoritative;
41[f]	sending a third passive DNS query, for each of the second domain names, to identify host names for the hosts of the second domain name and Internet Protocol addresses for the host names;

Claim	Limitation
41[g]	receiving a list of the host names and the Internet Protocol addresses for the host names; and
41[h]	mapping each of the Internet Protocol addresses to an attribute for the entity.
42	The method of claim 33 comprising presenting to the user non-technical-asset information through the user interface.
43	The method of claim 42 in which the non-technical-asset information comprises information about the entities.
44	The method of claim 43 in which the information about the entities comprises at least one of descriptions of entities, industry classifications, and employee counts.
45	The method of claim 33 in which the tool enables the user to assign a technical asset to an entity.
46	The method of claim 33 comprising enabling a separate review and approval of entities proposed to be included in the map.
47	The method of claim 33 comprising providing the map to an application for joining to event data or scoring the security state of the entities.

Claim	Limitation
48	The method of claim 33 in which the part of the generating of the map that is done automatically comprises at least one of collecting information about technical assets and about entities, associating technical assets with entities, and approving proposed portions of the map.
49	The method of claim 33 comprising generating graphs of relationships among entities based on their associations with technical assets.
50[pre]	A method comprising:
50[a]	generating a map between (a) technical assets that contribute to security characteristics of respective entities and (b) the identities of the entities that are associated with the respective technical assets, at least part of the generating of the map being done automatically,
50[b]	enabling a separate review and approval of entities proposed to be included in the map, and
50[c]	enabling a user to assist in the generating of the map by presenting to the user through a user interface (a) data about the technical assets of entities and (b) an interactive tool for associating the technical assets with the identities of the entities.
51	The method of claim 50 in which the technical assets comprise network-related information.
52	The method of any of claim 51 in which the network-related information comprises at least one of Internet Protocol addresses, blocks of Internet Protocol addresses, mail server

Claim	Limitation
	configurations, domain names, social media handles, third-party data hosting, internet service providers, Domain Name System services, Autonomous System numbers, and Border Gateway Protocol advertisements.
53	The method of claim 50 in which generating the map comprises online discovery of information about the technical assets.
54	The method of claim 53 in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or a Regional Internet Registry.
55	The method of claim 53 in which the information about the technical assets is discovered through passive DNS queries.
56	The method of claim 55 comprising identifying from the passive DNS queries associations between domain names and network addresses.
57	The method of claim 56 comprising tracking changes over time in the network that are associated with a given domain name.
58 <pre>[pre]</pre>	The method of claim 56 in which identifying from the passive DNS queries associations between domain names and network addresses comprises:
58 <pre>[a]</pre>	receiving a first domain name for an entity from the entities;

Claim	Limitation
58[b]	sending a first passive DNS query to identify first name servers for the first domain name;
58[c]	receiving a list of the first name servers for the first domain;
58[d]	sending a second passive DNS query, for each of the first name servers, to identify second domain names for which the name server is authoritative;
58[e]	receiving a list, for each of the first name servers, of the second domain names for which the name server is authoritative;
58[f]	sending a third passive DNS query, for each of the second domain names, to identify host names for the hosts of the second domain name and Internet Protocol addresses for the host names;
58[g]	receiving a list of the host names and the Internet Protocol addresses for the host names; and
58[h]	mapping each of the Internet Protocol addresses to an attribute for the entity.
59	The method of claim 50 comprising presenting to the user non-technical-asset information through the user interface.
60	The method of claim 59 in which the non-technical-asset information comprises information about the entities.
61	The method of claim 60 in which the information about the entities comprises at least one of descriptions of entities, industry classifications, and employee counts.

Claim	Limitation
62	The method of claim 50 in which the tool enables the user to assign a technical asset to an entity.
63	The method of claim 50 comprising invoking external crowd-sourced services to aid in generating the map.
64	The method of claim 50 comprising providing the map to an application for joining to event data or scoring the security state of the entities.
65	The method of claim 50 in which the part of the generating of the map that is done automatically comprises at least one of collecting information about technical assets and about entities, associating technical assets with entities, and approving proposed portions of the map.
66	The method of claim 50 comprising generating graphs of relationships among entities based on their associations with technical assets.
67[pre]	A method comprising:
67[a]	generating a map between (a) technical assets that contribute to security characteristics of respective entities and (b) the identities of the entities that are associated with the respective technical assets, at least part of the generating of the map being done automatically,
67[b]	providing the map to an application for joining to event data or scoring the security state of the entities; and

Claim	Limitation
67[c]	enabling a user to assist in the generating of the map by presenting to the user through a user interface (a) data about the technical assets of entities and (b) an interactive tool for associating the technical assets with the identities of the entities.
68	The method of claim 67 in which the technical assets comprise network-related information.
69	The method of any of claim 68 in which the network-related information comprises at least one of Internet Protocol addresses, blocks of Internet Protocol addresses, mail server configurations, domain names, social media handles, third-party data hosting, internet service providers, Domain Name System services, Autonomous System numbers, and Border Gateway Protocol advertisements.
70	The method of claim 68 in which generating the map comprises online discovery of information about the technical assets.
71	The method of claim 70 in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or a Regional Internet Registry.
72	The method of claim 70 in which the information about the technical assets is discovered through passive DNS queries.

Claim	Limitation
73	The method of claim 72 comprising identifying from the passive DNS queries associations between domain names and network addresses.
74	The method of claim 73 comprising tracking changes over time in the network addresses that are associated with a given domain name.
75[pre]	The method of claim 73 in which identifying from the passive DNS queries associations between domain names and network addresses comprises:
75[a]	receiving a first domain name for an entity from the entities;
75[b]	sending a first passive DNS query to identify first name servers for the first domain name;
75[c]	receiving a list of the first name servers for the first domain;
75[d]	sending a second passive DNS query, for each of the first name servers, to identify second domain names for which the name server is authoritative;
75[e]	receiving a list, for each of the first name servers, of the second domain names for which the name server is authoritative;
75[f]	sending a third passive DNS query, for each of the second domain names, to identify host names for the hosts of the second domain name and Internet Protocol addresses for the host names;
75[g]	receiving a list of the host names and the Internet Protocol addresses for the host names; and
75[h]	mapping each of the Internet Protocol addresses to an attribute for the entity.

Claim	Limitation
76	The method of claim 67 comprising presenting to the user non-technical-asset information through the user interface.
77	The method of claim 76 in which the non-technical-asset information comprises information about the entities.
78	The method of claim 77 in which the information about the entities comprises at least one of descriptions of entities, industry classifications, and employee counts.
79	The method of claim 67 in which the tool enables the user to assign a technical asset to an entity.
80	The method of claim 67 comprising invoking external crowd-sourced services to aid in generating the map.
81	The method of claim 67 comprising enabling a separate review and approval of entities proposed to be included in the map.
82	The method of claim 67 in which the part of the generating of the map that is done automatically comprises at least one of collecting information about technical assets and about entities, associating technical assets with entities, and approving proposed portions of the map.

Claim	Limitation
83	The method of claim 67 comprising generating graphs of relationships among entities based on their associations with technical assets.
84[pre]	A method comprising:
84[a]	generating a map between (a) technical assets that contribute to security characteristics of respective entities and (b) the identities of the entities that are associated with the respective technical assets, at least part of the generating of the map being done automatically,
84[b]	generating graphs of relationships among entities based on their associations with technical assets; and
84[c]	enabling a user to assist in the generating of the map by presenting to the user through a user interface (a) data about the technical assets of entities and (b) an interactive tool for associating the technical assets with the identities of the entities.
85	The method of claim 84 in which the technical assets comprise network-related information.
86	The method of any of claim 85 in which the network-related information comprises at least one of Internet Protocol addresses, blocks of Internet Protocol addresses, mail server configurations, domain names, social media handles, third-party data hosting, internet service providers, Domain Name System services, Autonomous System numbers, and Border Gateway Protocol advertisements.

Claim	Limitation
87	The method of claim 85 in which generating the map comprises online discovery of information about the technical assets.
88	The method of claim 85 in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or a Regional Internet Registry.
89	The method of claim 87 in which the information about the technical assets is discovered through passive DNS queries.
90	The method of claim 89 comprising identifying from the passive DNS queries associations between domain names and network addresses.
91	The method of claim 90 comprising tracking changes over time in the network addresses that are associated with a given domain name.
92[pre]	The method of claim 90 in which identifying from the passive DNS queries associations between domain names and network addresses comprises:
92[a]	receiving a first domain name for an entity from the entities;
92[b]	sending a first passive DNS query to identify first name servers for the first domain name;
92[c]	receiving a list of the first name servers for the first domain;

Claim	Limitation
92[d]	sending a second passive DNS query, for each of the first name servers, to identify second domain names for which the name server is authoritative;
92[e]	receiving a list, for each of the first name servers, of the second domain names for which the name server is authoritative;
92[f]	sending a third passive DNS query, for each of the second domain names, to identify host names for the hosts of the second domain name and Internet Protocol addresses for the host names;
92[g]	receiving a list of the host names and the Internet Protocol addresses for the host names; and
92[h]	mapping each of the Internet Protocol addresses to an attribute for the entity.
93	The method of claim 84 comprising presenting to the user non-technical-asset information through the user interface.
94	The method of claim 93 in which the non-technical-asset information comprises information about the entities.
95	The method of claim 94 in which the information about the entities comprises at least one of descriptions of entities, industry classifications, and employee counts.
96	The method of claim 84 in which the tool enables the user to assign a technical asset to an entity.

Claim	Limitation
97	The method of claim 84 comprising invoking external crowd-sourced services to aid in generating the map.
98	The method of claim 84 comprising enabling a separate review and approval of entities proposed to be included in the map.
99	The method of claim 84 comprising providing the map to an application for joining to event data or scoring the security state of the entities.
100	The method of claim 84 in which the part of the generating of the map that is done automatically comprises at least one of collecting information about technical assets and about entities, associating technical assets with entities, and approving proposed portions of the map.

I. INTRODUCTION

The '615 patent relates to a conventional method for mapping network assets, such as IP addresses, to an entity, and providing a user interface to assist in that mapping. But by 2013, the earliest priority date, such mapping tools were well-known. Indeed, although the '615 patent is purportedly about “[s]ecurity risk management” (EX1001, Title), very few of the 100 claims are directed to security.

Petitioner brings grounds that rely on a patent application (“Longo”) that is entitled “[m]apping network addresses to organizations,” and therefore directed to exactly the same subject matter as the '615 patent, including a user interface for assigning network assets to entities. In particular, Longo describes generating probable associations between candidate IP addresses and entities, scoring those associations, and then storing them in a database, with a user interface that allows manual associations. In combination with an O'Reilly-published book called “Network Security Assessment” (“McNab”), which describes various methods for enumerating IP address assets for an entity's network, these references render obvious the independent claims of the '615 patent.

The grounds also present an alternative based on another patent application (“Shull”) which similarly describes mapping IP addresses to entities and storing those associations in a data store. Shull describes a pop-up window that allows an administrators to confirm the associations.

II. STANDING

Petitioner certifies the '615 patent is available for IPR and they are not estopped.

III. GROUNDS

Petitioner presents the below grounds against Claims 1-100.

Ground	Basis	Reference(s)	Claims
1	§103	Longo, McNab	1-4, 9-14, 16, 33-37, 42-46, 48, 50-54, 59-63, 65
2	§103	Longo, McNab, Maltego	1-4, 9-14, 16-17, 33-37, 42-46, 48-49, 50-54, 59-63, 65-66, 84-88, 93-98, 100
3	§103	Longo, McNab, Zhang (optionally Maltego)	5-8, 13, 17, 18-29, 31-32, 33-46, 48-49, 55-58, 63, 66, 84-98, 100
4	§103	Longo, McNab, Tippet (optionally Maltego)	15, 47, 64, 67-71, 76-83, 99
5	§103	Longo, McNab, Zhang, Tippet (optionally Maltego)	30, 47, 72-75, 80, 83, 99
6	§103	Shull	1-4, 9-16, 33-37, 42-48, 50-54, 59-65, 67-71, 76-82
7	§103	Shull, McNab, Zhang	5-8, 13, 17, 18-32, 33-49, 55-58, 63, 66, 72-75, 80, 83, 84-100

IV. POSITA

A 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. More education could compensate for less experience and vice versa. *Id.*

Petitioner's expert surpassed this level of skill by September 2013. *Id.*

V. BACKGROUND

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. Web servers, which host websites that Internet users browse, are also identified using an IP address. However, IP addresses are difficult to memorize. Thus, 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.” 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.

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 PTR records, which map an IP address to a domain name.

Within a domain there could be various "hosts," with A records mapping the 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 discovering network assets using Open Source Intelligence ("OSINT") tools.

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. One tool is "whois," which returns 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
```

Another tool is 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.
```

The dig tool can also query the PTR records of a name server, which returns 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. To increase the efficiency of frequent DNS queries, a technology known as “passive DNS” was created, wherein a passive 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. One 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.


Maltego “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; EX1031, 3.




This retrieves two NS servers for the domain name “shmoos.com”: “krusty.shmoos.com” and “archimedes.shmoos.com.” *Id.*, 0:19; EX1031, 4. For each name server, Maltego can retrieve “what other domains use these name servers.” *Id.*




Some of the results are shown on the right of the screen, such as “kusia.net” and “legacyhigh.com.” EX1018, 0:41; EX1031, 6.



For each domain name discovered, Maltego can further find the hostnames for each domain. EX1018, 5:16; EX1031, 30.



Maltego also resolves the IP address for each hostname. EX1018, 5:34; EX1031, 32.



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. EX1003, ¶¶73-76.

B. Public Availability and Authenticity

McNab and other exhibits discussed in this Petition were publicly available before September 24, 2010, 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 indexed and made available to those of skill in the art 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, 1020, and 1025 are webpages and documents obtained 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 that the webpage was archived. *Id.*, ¶5. All of the Internet Archive exhibits are dated before September 13, 2009, and thus indicate that they were publicly available on the Internet before that date, and that the copies attached

at Exhibits 1017, 1025, and 1025 are true and correct copies of the webpage on the date they were archived.¹

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 Availability of Maltego 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 a video hosted on those sites. Exhibit 1019 is a video of Chris Page presenting the slideshow at Exhibit 1018. Exhibit 1021 is a video

¹ While Mr. Frank-White's description of the Internet Archive applies to any Internet Archive URL, his declaration at Exhibit 1013 only provides “browser screenshots of the Internet Archive's records of the archived files for the URLs” for webpages that are not exhibits in this Petition. EX1013, ¶¶6-7. Petitioner is seeking an additional declaration from the Internet Archive for Exhibits 1017, 1020, and 1025 but had not yet received it at the time of this Petition's filing.

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 archived link of the Vimeo webpage. Thus, Dr. Almeroth opines that the video linked at Exhibit 1019 was publicly available by early 2009 and is a true and correct copy. In addition, 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 published date of September 14, 2009 for Exhibit 1021 is also reliable, because YouTube is prominent and well-known among ordinary artisans. 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. 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. 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. ’615 patent (EX1003, ¶¶78-80)

The ’615 patent to Gladstone *et al.* was filed September 9, 2013 and does not claim priority to any earlier application.

Although the ’615 patent is purportedly directed to “security risk management” and in particular “a security state of [an] entity,” the claims of the ’615 patent are instead largely directed to merely mapping technical assets (such as IP addresses) to entities by querying online sources. The independent claims of the ’615 patent also provide for a user interface where a user can associate a technical asset with an entity.

1. Claim Mapping

Although the '615 patent has 100 claims, there is a large amount of overlap among them. The 100 claims can be categorized into six different sets of one independent claim and its dependent claims, based on the independent claims 1, 18, 33, 50, 67, and 84. As shown in the below chart, each set recites the same terms though in differing order. EX1003, ¶¶99-100.

	Independent Claim Set					
Claim Language	1	18	33	50	67	84
“generating a map between ...”	1[a]	18[a]	33[a]	50[a]	67[a]	84[a]
“online discovery ...”	1[b], 4	18[c]	36	53	70	87
“enabling a user to assist ...”	1[c]	18[e]	33[c]	50[c]	67[c]	84[c]
“... network-related information”	1[d], 2	18[b]	34	51	68	85
“information ... from [IANA] or [RIR]”	1[e]	20	37	54	71	88
“information comprises ... [IP] addresses ...”	3	19	35	52	69	86
“discovered through passive DNS queries”	5	18[d]	38	55	72	89
“associations between domain names and network addresses”	6	21	39	56	73	90
“tracking changes over time in the network addresses”	7	22	40	57	74	91
sending 3 passive DNS queries	8	23	41	58	75	92

	Independent Claim Set					
Claim Language	1	18	33	50	67	84
“presenting ... non-technical-asset ...”	9	24	42	59	76	93
“... information about the entities”	10	25	43	60	77	94
“descriptions of entities, industry classifications, and employee counts”	11	26	44	61	78	95
“... user to assign a technical asset to an entity”	12	27	45	62	79	96
“invoking external crowd-sourced services ...”	13	28	33[b]	63	80	97
“separate review and approval of entities ...”	14	29	46	50[b]	81	98
“... scoring the security state of the entities”	15	30	47	64	67[b]	99
“... automatically ... collecting information about technical assets and about entities, ...”	16	31	48	65	82	100
“generating graphs ...”	17	32	49	66	83	84[b]

2. Prosecution History

During prosecution, the examiner rejected the sole independent claim of the original set of claims, but allowed the dependent claims. The applicant then

amended the dependent claims to be rewritten as independent claims, resulting in allowance. EX1003, ¶¶81-82.

VI. CLAIM CONSTRUCTION

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

VII. GROUND 1: LONGO, MCNAB

Longo describes a system and user interface for mapping the IP addresses of network assets with organizations. McNab describes various methods for discovering network assets (and their IP addresses) of organizations, including using information from third-party systems like DNS servers and Regional Internet Registries. It would have been obvious to combine the two.

A. Motivation to Combine

A POSITA would be motivated to incorporate McNab's teachings of enumerating assets of an entity by discovering their IP addresses. EX1003, ¶¶101-08. Such teachings were known to provide useful and necessary information to assess the security risk of an organization's network. *Id.*

McNab 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, and then performing "Bulk Network Scanning and Probing" and "Investigation of Vulnerabilities." EX1005, 4-

7. For example, McNab discloses methods such as “DNS Querying” and a “zone transfer” to obtain IP addresses of computer within a domain. EX1005, 30-35. McNab also describes querying Regional Internet Registries such as ARIN and RIPE to obtain WHOIS information about network assets. EX1005, 23-24.

A POSITA would have been motivated to incorporate these teachings from McNab to provide a well-known and reliable method of determining the IP addresses of entities. Longo already describes that knowing which IP addresses belong to which entities has “valuable uses.” EX1026, ¶5. Indeed, the very goal of Longo is to store “actual associations between the candidate address and the organizational entity.” *Id.*, Abstract. A POSITA would have been motivated to give effect to these teachings in Longo by using the techniques disclosed in McNab to identify IP addresses for a particular entity. EX1003, ¶¶103-04. In addition, as discussed above in the background, the use of these tools to enumerate network assets (including enumerating their IP addresses) using public information sources was well-known and could be accomplished using automated tools such as Maltego (which itself uses well-known DNS queries). *Id.* Network enumeration and security risk assessment were also well-known and ubiquitous methods for helping secure an organization’s networks. *Id.* A POSITA would be motivated in view of this knowledge to associate IP addresses with entities with the assistance of tools described in McNab and generally known to a POSITA, such as DNS query tools like Maltego. *Id.*

In addition, a POSITA would have been motivated to use McNab’s teachings because of the well-known benefits of determining IP addresses of network assets to perform vulnerability scans on those assets. EX1003, ¶105. McNab itself teaches the importance of its methodologies: “[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 security assessment, 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.

With respect to McNab’s teachings of determining IP addresses of network assets, a POSITA would have been motivated by McNab’s teaching of collecting “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.*

A POSITA would have reasonably expected success in combining Longo with McNab. McNab provides detailed and sufficient information to perform asset enumeration (such as using DNS querying and requesting a zone transfer). It would

have been well-within the capabilities of a POSITA to incorporate such mechanisms into Longo as another source of potential IP address-entity associations. EX1003, ¶107. Indeed, Longo already discloses that “public address registr[ies]” such as RIRs are used to identify potential IP addresses associated with an organization, and therefore a POSITA would find it readily achievable to use other sources of information as well. EX1026, ¶15.

Finally, Longo and McNab are analogous art to the ’615 patent because all three references are directed to the field of determining network assets of an entity. EX1026, Abstract; EX1005, 5; EX1001, Abstract; EX1003, ¶108.

B. Claim 1

1. **1[pre], 1[a]: A method comprising: generating a map between (a) technical assets that contribute to security characteristics of respective entities and (b) the identities of the entities that are associated with the respective technical assets, at least part of the generating of the map being done automatically,**

“generating a map between (a) technical assets ... and (b) the identities of the entities that are associated with the respective technical assets”: Longo describes a method for mapping IP addresses (*i.e. technical assets*) with “organizational entities” that the IP addresses belong to. EX1026, Abstract; EX1003, ¶109. In particular, Longo describes a method where “[n]etwork addresses (e.g., IP addresses) are considered candidates for association with organizational entities ..., based on some evidence of a relationship between them.” *Id.* For

example, Longo describes “a correlation engine 135 that ... identifies which addresses [*i.e. technical assets*] genuine[ly] belong to which organizational *entities*,” which are then stored in “a network address database 138 that stores the *address/organization associations* identified by the correlation engine 135,” *i.e. generating a map between assets and entities*. *Id.*, ¶15. Longo further describes that the correlation engine is coupled to an “organization database 134” which stores information such as “a name of the organization entity,” *i.e. identities of the entities*.

“technical assets that contribute to security characteristics of respective entities”: Longo discloses *entities’ technical assets* in the form of an organization’s “IP addresses,” and a POSITA would understand (and McNab makes explicit) that such IP addresses *contribute to security characteristics of respective entities*. EX1003, ¶110. McNab explains that the first step of “Network Security Assessment Methodology” is to “identify IP networks and hosts of interests,” including obtaining “[k]ey pieces of information ... gathered through initial reconnaissance” such as “details of Internet-based network blocks, internal *IP addresses* gathered from DNS servers, insight into the target organization’s DNS structure (including domain names, subdomains, and hostnames) ...” EX1005, 5. The reason IP addresses *contribute to security characteristics* of an entity’s network is because every IP address is a potential target for attackers. *See id.*, 2-3 (“[O]ppportunistic attackers will continue to scour the public IP address space seeking vulnerable hosts.”). Thus,

McNab teaches that enumerating the IP addresses of an entity's network is, from a security perspective, "very important because it can uncover hosts that aren't properly fortified against attack." *Id.*, 5. Such teachings in McNab generally describe well-known characteristics of an entity's network addresses that a POSITA would understand applies to Longo's network addresses. EX1003, ¶110. In the alternative, as discussed above, it would have been obvious to combine Longo with these teachings in McNab.

"at least part of the generating of the map being done automatically":

Longo's methods for mapping IP addresses to entities are performed at least partly automatically. EX1003, ¶111. For example, Longo describes an "association server 130," which includes the correlation engine discussed above, that "receives and *processes* interaction information from which IP addresses ... can be identified as candidates for association with organizational entities." EX1026, ¶¶14-15. Once the associations are identified, the "network address database 138" then stores those associations, *i.e. generates the map. Id.* A POSITA would understand that since these processes do not require human interaction, they are *done automatically*. In addition, Longo discloses a "scoring module 137" which "computes a quality score for each relationship of an address with an organization" using various weighs, wherein "the weights are determined ... *automatically.*" *Id.*, ¶45.

2. **1[b]: generating of the map comprising online discovery of information about the technical assets; and**

Longo discloses that generating the address-entity associations (*i.e. generating the map*) includes obtaining information about the IP addresses from a “public address registry” such as “ARIN or RIPE,” which store “blocks of IP addresses with which [individual businesses] are associated.” EX1026, ¶17. ARIN and RIPE are known as “Regional Internet Registries (RIRs)” “under which all public Internet-based address spaces [for a region] fall.” EX1005, 23-24. ARIN and RIPE “can be queried to glean useful information” such as “details of IP network blocks.” *Id.* Since they are publicly accessible, such information constitutes *online discovery of information*. EX1003, ¶112.

3. 1[c]: enabling a user to assist in the generating of the map by presenting to the user through a user interface (a) data about the technical assets of entities and (b) an interactive tool for associating the technical assets with the identities of the entities,

Longo describes generating IP address-entity associations (*i.e. generating a map*) through a user interface where users provide “interaction information from which IP addresses ... can be identified as candidates for association with organizational entities,” specifically, selecting an organization on a web form which is then associated with the user’s IP address. EX1026, ¶14. For example, a user may “view a web page form for which the application server 130 ... provides all or part of the functionality,” *i.e. presenting to the user through a user interface. Id.*, ¶19. Through this web page, a user is presented with *data about the technical assets*

of entities, such as an a “welcome control ... based on the IP address through which the user accesses the page..” *Id.* The user interface also presents to the user “an organization selection control” that “auto-complete[s] organization names entered on the form,” *i.e. an interactive tool.* *Id.* When the form is submitted, the form “provide[s] the organization name [*i.e. identities of an entity*] and the IP address [*i.e. technical asset*] through which the page is accessed, back to the interaction module,” which stores this data as an “address/organization relationship[] in the interaction database 132,” *i.e. associating the technical asset with the identities of the entities.* *Id.*; ¶15 (“[C]orrelation engine 135 ... processes the interaction information and identifies which addresses genuine[ly] belong to which organizational entities.”); EX1003, ¶¶113-14.

4. 1[d]: the technical assets comprising network-related information,

Longo discloses mapping *IP addresses* with entities. IP addresses are *network-related information.* EX1001, 2:1-3 (“The network-related information includes ... Internet Protocol addresses ...”); EX1003, ¶115.

5. 1[e]: and in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or Regional Internet Registry.

As discussed above, Longo discloses using ARIN or RIPE to discover information about technical assets. EX1026, ¶17; EX1003, ¶116. ARIN and RIPE

are the ***Regional Internet Registries*** for North America and Europe, respectively.

EX1005, 23-24.

C. Claim 2

See limitation 1[d], which is identical to this claim. EX1003, ¶117.

D. Claim 3: The method of claim 2 in which the network-related information comprises at least one of Internet Protocol addresses, ...

As discussed above for limitation 1[d], Longo’s technical assets include ***IP addresses***. EX1026, Abstract; EX1003, ¶118.

E. Claim 4

See limitation 1[b], which is identical to this claim; EX1003, ¶119.

F. Claim 9: The method of claim 1 comprising presenting to the user non-technical-asset information through the user interface.

Claim 10: The method of claim 9 in which the non-technical-asset information comprises information about the entities.

Claim 11: The method of claim 10 in which the information about the entities comprises at least one of descriptions of entities ...

Longo discloses that the user interface includes “an organization selection control” which “auto-complete[s] organization names entered on the form.” EX1026, ¶19. Such information is “non-technical-asset information” about ***entities*** because it would include organization names with which the asset is not associated. EX1003, ¶120. In some cases, the names are ***descriptions of the entity*** that it names, such as “Legal Department, NPE Suit Mfg. Ltd.” EX1026, ¶43; EX1003, ¶¶121-22.

G. Claim 12: The method of claim 1 in which the tool enables the user to assign a technical asset to an entity.

As discussed above for limitation 1[c], Longo describes providing a webpage with “an organization selection control,” *i.e. interactive tool*, that allows a user to select an organization which when submitted, will associate the selected organization with the user’s IP address, *i.e. enables the user to assign a technical asset to an entity*. EX1026, ¶19; EX1003, ¶123.

H. Claim 13: The method of claim 1 comprising invoking external crowd-sourced services to aid in generating the map.

McNab discloses obtaining information about entities by “scour[ing] ... newsgroups,” which a POSITA would understand to be an *external crowd-sourced service*. EX1005, 18. McNab discloses that newsgroups can uncover information such as “[c]ontact details” and “[p]hysical addresses of offices and other locations.” *Id.* A POSITA would understand that such information would be useful for providing information about entities within Longo’s “organization database 134” which stores information “such as mailing address” and “area code.” EX1026, ¶18; EX1003, ¶124.

I. Claim 14: The method of claim 1 comprising enabling a separate review and approval of entities proposed to be included in the map.

Longo discloses that if a user determines that the organization with which its IP address is associated with is incorrect, the user can provide a different organization name to be included in the map. Specifically, Longo discloses that “[i]f

the user specifies that he or she belongs to a different organization than the one originally suggested, the welcome control sends the IP address and the specified organization name to the interaction identification module 131,” which then adds the association to the association database. EX1026, ¶19. For an entity that has not been entered in the database, this constitutes a *separate review and approval* by the user of an entity that the user proposes to be included in the map. EX1003, ¶125.

J. Claim 16: The method of claim 1 in which the part of the generating of the map that is done automatically comprises ... associating technical assets with entities ...

As discussed above for limitation 1[b], Longo discloses that the association of IP address (*i.e. technical assets*) with organizations (*i.e. entities*) is performed automatically by the “association server” which “receives and *processes* interaction information from which IP addresses ... can be identified as candidates for association with organizational entities.” EX1026, ¶¶14-15; EX1003, ¶126.

K. Claim 33-37, 42-46, 48, 50-54, 59-63, 65

As discussed above in Section V.C.1, claims 33-37, 42-46, 48, 50-54, 59-63, and 65 all recite limitations from claims 1-4, 9-14 and 16. . Thus, they are rendered obvious by Longo and McNab. EX1003, ¶¶127-28.

VIII. GROUND 2: LONGO, MCNAB, MALTEGO

As discussed above in the Technical Background, Maltego is a well-known security tool that provides a user interface for mapping assets like the IP addresses of hosts to organizations. As discussed below, it would have been obvious to

combine Maltego as described in Exhibits 1018, 1021, and 1025 with Longo and McNab. This combination provides an alternative reason why limitation 1[c] and claims 9-11 are rendered obvious. EX1003, ¶129.

A. Public Availability of Maltego Documents

Maltego, and in particular its user interface, is described in Exhibits 1018, 1021, and 1025. EX1003, ¶130. As discussed above, these were publicly available before September 2013.

B. Motivation to Combine

A POSITA would have been motivated to combine Longo and McNab with teachings from the Maltego documents (Exhibits 1018, 1021, and 1025) that describe a user interface for assigning IP address assets to organizations. EX1003, ¶¶131-34. In particular, Maltego was a third-party software that allowed a user to assign an IP address to an organization using drag-and-drop.

A POSITA would have been motivated to incorporate this teaching into Longo in order to provide a user-friendly interface for mapping IP addresses to organizations, which is the purpose of Longo's system. Although Longo describes a system for automatically associating IP addresses with an organization by using a scoring system to assess accuracy of proposed associations, a POSITA would have found it desirable to provide a mechanism for a user to manually assign IP addresses to an organization if that user definitively knows that the association is accurate.

EX1003, ¶132. For example, although Longo allows individual users to update the organization for which the user's IP address is associated with, it would be beneficial to allow an administrator to also manually update those associations if they are aware of changes. *Id.* In addition, if the administrator knows that, for example, IP addresses from certain domain names belong to the same organization, it would be beneficial to provide a user interface to manually create that association. *Id.*


A POSITA would have reasonably expected the combination to succeed. Maltego was a well-known tool for mapping IP address assets to organizations. The ability to manually assign IP addresses to organizations using a user interface would have been within the skill of an ordinary artisan and not a difficult task to implement. Indeed, the Maltego software itself could have been integrated with Longo and McNab, because Maltego describes exporting its associations, which could have then been used with Longo to store in its association database. EX1003, ¶133.

Finally, Maltego is analogous art to the '651 patent because it is directed to enumerating IP addresses for an entity. *See, generally*, EX1018, EX1021; EX1003, ¶¶134.

C. 1[c]


The Maltego documents describe a *user interface* that presents *data about the technical assets of entities* and an *interactive tool for associating entities with entities*. EX1003, ¶135.

“presenting to the user through a user interface (a) data about the technical assets of entities”: Maltego has a user interface that can depict data about IP addresses (*i.e. technical assets*) of entities. EX1003, ¶136. For example, in the below screenshot, it shows the name of the MX server (which a POSITA would understand has an IP address associated with it and is therefore a ***technical asset*** of the baytsp.com entity). EX1018, 70.




Other hostnames associated with the baytsp.com entity, which have IP addresses (*i.e. technical assets*), can also be displayed, along with data about them

such as their names or their associations with the organization. EX1018, 71; EX1003, ¶137.



This is also shown in the YouTube demonstration of Maltego. EX1021, 0:41; EX1031, 6; EX1003, ¶138.




The Maltego user guide also describes other information about these *network assets* such as their “weight,” an “entity icon,” and the number of incoming and outgoing links. EX1025, 8; EX1003, ¶139.





“presenting to the user through a user ... (b) an interactive tool for associating the technical assets with the identities of the entities”: The Maltego

documents describe that the user interface can allow a user to manually associate a technical asset with the domain name of an entity (*i.e. identity of the entity*). EX1003, ¶140. For example, for each domain name's IP address, the Maltego interface can manually create an IP address block for which additional IP addresses (*i.e. technical assets*) can be associated with the entity. EX1018, 73.



The YouTube demonstration of Maltego also teaches that for each organization, a user can use the user interface to discover additional “web sites for each domain” and their IP addresses. EX1021, 5:16, 5:34; EX1031, 31, 33; EX1003, ¶141.

Petition for *Inter Partes* Review
U.S. Patent No. 9,438,615




D. Claims 9-11

Maltego's user interface can display non-technical-asset information about entities, such as the description of those entities or the employee count. EX1003, ¶142.

For example, Maltego's user interface displays the domain name and weight of an entity, which are *descriptions of the entities*. EX1003, ¶143. The domain name could also be an *industry classification*, since some top-level domain names are indicative of the industry, such as .edu, .mil, and .com. EX1025, 8.



For example, Maltego's user interface can display email addresses associated with the entity, which corresponds with an *employee count*. EX1018, 74; EX1003, ¶¶144.



E. Claim 17: The method of claim 1 comprising generating graphs of relationships among entities based on their associations with technical assets.


Maltego discloses that its user interface can generate *graphs of relationships among entities based on their associations with technical assets*. EX1003, ¶¶145-46.

For example, Maltego can depict a graph in a “Centrality view” where “Nodes that are calculated to be the most central to the graph are given larger nodes.” EX1025. As Exhibit 1018 describes it, this graph shows “relationships” between

entities, and depicts each entity based on associations with technical assets.

EX1018, 80.

The same graph used in the previous section is shown in different views. First the centrality view:



F. Remaining Limitations From Claim 1 and Claims 2-4, 12-14, and 16

The remaining limitations for Claim 1, as well as claims 2-4, 12-14, and 16, are rendered obvious for the same reasons discussed above for Ground 1. EX1003, ¶147.

G. Claim 33-37, 42-46, 48-49, 50-54, 59-63, 65-66, 84-88, 93-98, 100

As discussed above in Section V.C.1, claims 33-37, 42-46, 48-49, 50-54, 59-63, 65-66, 84-88, 93-98, and 100 all recite limitations from claims 1-4, 9-14 and 16-

17. Thus, they are rendered obvious by Longo, McNab, and Maltego. EX1003, ¶¶148-50.

IX. GROUND 3: LONGO, MCNAB, AND ZHANG (WITH OR WITHOUT MALTEGO)

Zhang discloses enumerating IP addresses of an entity by querying passive DNS information. 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 Zhang with Longo and McNab (with or without Maltego).

A. Motivation to Combine

A POSITA would have found it obvious to combine Zhang's teaching of mapping IP addresses to an entity by querying a passive DNS database as well as McNab's teaching of IP address enumeration using DNS information with Longo's system of associating IP addresses with entities. EX1003, ¶¶151-60.

As discussed above, McNab teaches various methods for enumerating IP addresses of an entity's network, and it would have been obvious to include such techniques with Longo. Further examples relevant to this ground that are disclosed in McNab include various methods for querying DNS records to obtain IP addresses of computer within a domain. EX1005, 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.*; EX1003, ¶152.

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 find IP address-entity associations to be stored in Longo’s system. EX1003, ¶153. 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 and therefore likely belonging to the same organization. *Id.*

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, ¶154. 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 Longo 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.*

In addition, Zhang teaches collecting “passive DNS (pDNS) information ..., in which DNS responses are collected passively (e.g., by listening to DNS traffic) from distributed sources (e.g., using passive DNS sensors).” EX1006, ¶47. Zhang discloses that passive DNS 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; ¶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”).

In view of the fact that passive DNS information includes all DNS records that have been queried, a POSITA would be motivated to combine Longo and McNab with Zhang's disclosure of passive DNS information collection and storage within a pDNS database. EX1003, ¶156. A POSITA would be motivated because of the well-known benefits of passive DNS, namely that DNS records can be structured and stored in a manner that improves search capabilities and reliability. As Zhang itself teaches, "passive DNS information can be collected, correlated, and analyzed to determine relationships and associations among various entities in DNS (e.g., IP addresses, domains, BGP prefixes, ASs, DNS name servers, and/or other entities." EX1006, ¶78. For example, passive DNS databases can keep track of changes of DNS records, such as changes to which IP address is mapped to which domain name. *Id.*, ¶81.

Another example is that passive DNS databases can more easily query similarities between records, such as querying which domain names use the same name server or which hostnames are part of the same domain name. For example, passive DNS database can easily record all hostnames belonging to a certain domain like "apple.com," including mail (e.g. mail.apple.com) and FTP servers (e.g. ftp.apple.com). Such information may not be easily obtainable via conventional DNS queries and often require performing brute force searching. EX1003, ¶157. Finally, a passive DNS database provides the benefit of a central repository for DNS

information, whereas normally such information may require querying multiple DNS servers and name servers. *Id.* Thus, in view of these known benefits, a POSITA would be motivated to include Zhang’s teachings of passive DNS used for enumerating IP addresses with Longo and McNab’s methods for enumerating IP addresses.

Zhang also discloses the ability to generate various graphs of relationships among entities and their domains. EX1006, ¶84. A POSITA would have found it obvious to incorporate these teachings into Longo because, as taught in Zhang, such graphs allow for automatically determining “a reputation of new domains” based on its placement in the graph. *Id.* Zhang teaches that this is beneficial from a security perspective, because a new domain can be determined to be malicious or insecure based on its relationship with known malicious or insecure domains. *Id.*; EX1003, ¶158.

A POSITA would have had a reasonable expectation of success in combining Longo and McNab with Zhang. EX1003, ¶159. 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 by querying

DNS records using Zhang's passive DNS database to obtain the desired information for enumerating IP addresses. *Id.* Such information would then be used to associate IP addresses with an entity to be stored in Longo's system. 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.* A POSITA would also reasonably expect combining Longo with Zhang's graph generation teachings to succeed, since generating graphs using the passive DNS information was well-within the capabilities of a POSITA. The above motivations and reasonable expectation of success would not change if Zhang is combined with Longo, McNab and Maltego, as Maltego's user interface would have been entirely compatible with Zhang's passive DNS system.

Finally, Zhang is also analogous art to the '615 patent because it is directed to enumerating IP addresses for an entity. EX1006, Abstract; EX1003, ¶160.

B. Claim 5: The method of claim 4 in which the information about the technical assets is discovered through passive DNS queries.

As discussed above in the Motivation to Combine, it would have been obvious to collect information about technical assets by querying DNS records stored in a passive DNS database as taught by Zhang. EX1003, ¶161.

C. Claim 6: The method of claim 5 comprising identifying from the passive DNS queries associations between domain names and network addresses.

Zhang discloses that passive DNS queries can be used to resolve a domain name to an IP address. EX1006, ¶88 (“‘foo.com’ can be resolved to one IP and has one name server (NS).”); EX1003, ¶1162.

D. Claim 7: The method of claim 6 comprising tracking changes over time in the network addresses that are associated with a given domain name.

McNab discloses that its “Network Enumeration” techniques, which 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, ¶163. As discussed above, it would have been obvious to perform McNab’s DNS queries using Zhang’s passive DNS database.

E. Claim 8

1. 8[pre]: The method of claim 6 in which identifying from the passive DNS queries associations between domain names and network addresses comprises:

As discussed herein, McNab discloses a method for *associating domain names with IP addresses* by, using DNS queries, (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 each of the found domain names. And as discussed above, it would have been obvious to perform such DNS queries on Zhang’s passive DNS database. EX1003, ¶164.

2. 8[a]: receiving a first domain name for an entity from the entities;

Longo discloses that the organization databases “lists domain[s]” for the stored organizations, *i.e. domain name for an entity from the entities*. EX1026, ¶39. It would have been obvious to use this domain name to enumerate IP addresses that can be associated with the entity that owns that domain name by using McNab’s methods and Zhang’s passive DNS database, as discussed in the motivation to combine section. EX1003, ¶165.


3. 8[b]: sending a first passive DNS query to identify first name servers for the first domain name;

8[c]: receiving a list of the first name servers for the first domain;

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, ¶166.

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. It was well-known that DNS records and WHOIS information includes the authoritative name server of a particular domain. EX1003, ¶167. 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, ¶168.

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
```




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, ¶169.

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, ¶¶170-71. 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; EX1031, 3. This retrieves two NS servers for the domain name “shmoo.com”: “krusty.shmoo.com” and “archimedes.shmoo.com.” EX1021., 0:19; EX1031, 4; EX1003, ¶171.



4. 8[d]: sending a second passive DNS query, for each of the first name servers, to identify second domain names for which the name server is authoritative;

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

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, ¶172.

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, ¶173. 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, ¶174.

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

first name
server

```
; <>> 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.
```

list of second
domain names


A POSITA would understand that the query will return domain names *for which the designed name server is authoritative*. EX1003, ¶175. 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. 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, ¶175. 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, ¶176.

In the alternative, this limitation is disclosed under Patent Owner’s infringement theory for related patent U.S. Pat. No. 11,652,834, which recites a limitation that is identical to this one. *See* EX1027. Patent Owner accuses the use of WHOIS queries as the claimed “second passive DNS query.” EX1027, 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, ¶177.




As discussed above for limitation 8[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, ¶178. 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 8[b], Zhang discloses querying and storing WHOIS records, and McNab discloses performing a WHOIS query to obtain authoritative name server information. EX1003, ¶179.


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, ¶¶180-82. 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; EX1031, 4.



This results in domain names such as “kusia.net” and “legacyhigh.com.” *Id.*, 0:41; EX1031, 6.



5. **8[f]: sending a third passive DNS query, for each of the second domain names, to identify host names for the hosts of the second domain name and Internet Protocol addresses for the host names;**

8[g]: receiving a list of the host names and the Internet Protocol addresses for the host names; and

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, ¶¶183-84.

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, ¶185. 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
```

second domain
name

```
-----  
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  
-----
```

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, ¶188.

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 8[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, ¶188.

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, ¶¶189-90. 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; EX1031, 30. Maltego also resolves the IP address for each hostname. *Id.*, 5:34; EX1031, 32.

Petition for *Inter Partes* Review
U.S. Patent No. 9,438,615



6. 8[h]: mapping each of the Internet Protocol addresses to an attribute for the entity.

As discussed above for limitation 1[b], Longo discloses mapping IP addresses to a name of an organization, *i.e. an attribute for the entity*. As discussed in the Motivation to Combine section, it would have been obvious to use the results of the above techniques for enumerating IP addresses in Longo by associating the discovered IP addresses with the name of the entity that owns the domain name. EX1003, ¶191.

F. Claim 13

As discussed for Claims 5-8, it would have been obvious to use Zhang’s passive DNS database to *aid in generating the map* between IP addresses and organizations. A POSITA would understand that the passive DNS database is an *external crowd-sourced service*, because it collects passive DNS information “in which DNS responses are collected passively (e.g., by listening to DNS traffic) from distributed sources (e.g., using passive DNS sensors).” EX1006, ¶47. Thus, passive DNS information is *crowd-sourced* because it is built from all the DNS traffic from many users that can be detected. EX1003, ¶192.

G. Claim 17

Zhang discloses the capability of generating graphs of relationships between domains (*i.e. entities*) based on their associations with hosts or IP addresses (*i.e. technical assets*). Specifically, Zhang discloses generating a “DNS relation graph” where each domain is a node and the edges between them represent a relationship

such as “mapped to the same IP” or “hosted by the same name server” or “originated from the same session.” EX1006, ¶¶84, 86. These relationships are *based on the entities associations with technical assets*, such as having the same IP address. EX1003, ¶193.

H. Claims 18-29, 31-32, 33-46, 48-49, 55-58, 63, 66, 84-98, 100

As discussed above in Section V.C.1, claims 18-29, 31-32, 33-46, 48-49, 55-58, 63, 66, 84-98, and 100 all recite limitations from claims 1-14 and 16-17. Thus, they are rendered obvious by Longo, McNab, and Zhang (with or without Maltego). EX1003, ¶¶194-97

X. GROUND 4: LONGO, MCNAB, AND TIPPETT (WITH OR WITHOUT MALTEGO)

Tippett describes collecting security information about the assets of an entity’s network and then generating a composite rating in the form of a security risk index.

A. Motivation to Combine

A POSITA would have been motivated to combine Tippett’s teachings with Longo and McNab (and optionally, Maltego). EX1003, ¶¶198-202. Tippett describes a method for calculating a “local security index” for an organization. Specifically, discloses a “security risk index determining module [which] determin[es] the security risk index” and that it is a “software module[] running in a distributed computing environment.” EX1028, ¶¶29-30. Tippett further discloses

that the “security risk index ... may be a local security risk index representing the security of ... information resources associated with the entity.” *Id.*, ¶45.

It would have been obvious to combine these teachings of a security risk index with Longo in order to provide the additional functionality of evaluating the security of an organization’s network assets. EX1003, ¶199. As Tippet explains it, the security risk index allows an entity to “certify the security of the ... information resources associated with the entity” if the index “reaches a certification level.” EX1028, ¶62. This certification “verif[ies] to outside parties such as clients, partners, insurers, vendors, consultants, associates, competitors, or other parties, that the specific subset of information resources associated with the entity is secure.” *Id.* This is consistent with Longo’s disclosure wherein enumerating the IP addresses of an organization “in a business-to-business (B2B) environment, is useful for a website operator to know the identity of organizations viewing its website.” EX1006, ¶5. Thus, a POSITA would have been motivated to incorporate Tippet’s security risk index in order to provide additional security assurances from one business to another, as disclosed in Longo. EX1003, ¶199.

The security risk index also allows an administrator to determine when associated IP addresses are at risk of a security incident, which a POSITA would find beneficial. EX1003, ¶200. Using Tippet’s teachings, Longo could, for example, notify the organization that an IP address associated with the organization

is at high risk for a security incident and has reduced the overall security of the organization's network. *Id.* A POSITA would find this beneficial as it increases the security of the organization's network. *Id.*

A POSITA would have reasonably expected the combination to succeed because Tippet's provides detailed instructions regarding the calculation of the security risk index based on querying various information resources to collect information about the IP addresses of an organization. EX1003, ¶201. A POSITA would have combined Longo and Tippet by performing the local security risk index algorithm described in Tippet with the entity map described in Longo to generate security risk indices for each entity. *Id.* This would have been well within the capabilities of an ordinary artisan. *Id.* The above motivations and reasonable expectation of success would not change if Tippet is combined with Longo, McNab and Maltego, as Maltego's user interface would have been entirely compatible with Tippet's security risk index. *Id.*

Finally, Tippet is analogous art to the '615 patent because it is directed to associating network resources with an entity. EX1028, ¶32; EX1003, ¶202.

B. Claim 15: The method of claim 1 comprising providing the map to an application for joining to event data or scoring the security state of the entities.

Tippet discloses the calculation of a "local security index" for one or more organizations, *i.e. scoring the security state of the entities*. Specifically Tippet

discloses a “security risk index determining module [which] determin[es] the security risk index” and that it is a “software module[] running in a distributed computing environment.” EX1006, ¶¶29-30. Tippettt further discloses that the “security risk index ... may be a local security risk index representing the security of ... information resources associated with the entity.” *Id.*, ¶45. The “local security risk index” is calculated by using various information related to the security and resilience of the network. EX1003, ¶203.

As discussed above in the Motivation to Combine section, it would have been obvious to provide the mapping to the security risk index algorithm described in Tippettt, which results in a *security state of the entities*. EX1003, ¶204.

C. Claim 47, 64, 67-71, 76-83, 99

As discussed above in Section V.C.1, claims 47, 64, 67, and 99 all recite limitation 15. Claims 68-71 and 76-83 all recite limitations from claims 1-4, 9-14 and 16-17. Thus, they are rendered obvious by Longo, McNab, and Tippettt (with or without Maltego). EX1003, ¶¶205-06.

XI. GROUND 5: LONGO, MCNAB, ZHANG, AND TIPPETT (WITH OR WITHOUT MALTEGO)

As discussed above in Ground 3, it would have been obvious to combine Zhang with Longo and McNab (with or without Maltego). As discussed above in Ground 4, it also would have been obvious to combine Tippettt with Longo and McNab (with or without Maltego). For the reasons discussed in those grounds, it

also would have been obvious to combine Tippet *and* Zhang with Longo and McNab (with or without Maltego). The combination of these references would not pose any incompatibilities because Tippet’s security rating index and Zhang’s use of passive DNS were compatible and separate functions that a POSITA would have found obvious and reasonably expected to succeed when combined. EX1003, ¶207. In particular, the combination would use Zhang’s passive DNS database to enumerate network assets which could then be ranked using Tippet’s security rating index. *Id.*

As discussed above in Section V.C.1, claims 30, 47, 72-75, 80, 83, and 99 all recite limitations from claims 5-8, 13, 15, and 17. Thus, they are rendered obvious by Longo, McNab, Zhang and Tippet (with or without Maltego). EX1003, ¶¶208-11.

XII. GROUND 6: SHULL

A. Claim 1

1. 1[pre], 1[a]

“generating a map between (a) technical assets ... and (b) the identities of the entities that are associated with the respective technical assets”: Shull describes a method for mapping IP addresses (*i.e. technical assets*) to entities such as the owner of the IP address. EX1003, ¶212. Particularly, Shull describes gathering information like the “registered owner[s]” (*i.e. identities of entities*) of IP addresses and then storing such information in a “data store” where the registered owner is

“cross-indexed and/or cross-referenced” with other records for that owner, such as their WHOIS record. EX1032, ¶48. Shull also discloses “data grouping” such as grouping domains that are owned by the same individual (*i.e. identity of an entity*), which results in a “detailed *map* of the individual’s online activities.” *Id.*, ¶50. The ’615 patent discloses that an entity can be an individual. EX1001, 3:53-56.

“technical assets that contribute to security characteristics of respective entities”: As discussed above in Ground 1, a POSITA would understand (and McNab makes explicit) that IP addresses and domains *contribute to security characteristics of respective entities*. EX1003, ¶213. In addition, Shull describes that domains and IP addresses can be used to determine if the entity that owns them “is involved in illicit online activity.” EX1032, ¶52.

“at least part of the generating of the map being done automatically”: Shull discloses that the mapping is performed by a “controller 110” which is “one or more computers operating a controller application,” which a POSITA would understand to perform the above functions automatically. EX1032, ¶39; EX1003, ¶214.

2. 1[b]

Shull discloses that it generates the map using “data gathered through a variety of harvesting/investigation information.” EX1032, ¶47. This includes querying “network databases” to determine “the entity to which a particular IP address ... is allocated” or “DNS tables” or a “WHOIS database.” *Id.*, ¶¶43-44. In addition, as

discussed above, it would have been obvious to combine Shull with McNab's disclosure of gathering information from various sources to enumerate network assets. EX1003, ¶215.

3. 1[c]

Shull describes that data about IP addresses can be “correlated” using other means. EX1032, ¶81. For example, “data for information matching any of the data element(s) and/or groups currently being analyzed” can “be examined ... by a technician.” *Id.* Shull discloses that a “data element” includes an “IP address” as well as “one or more names” of entities. *Id.*, ¶73. Shull describes providing “the matching information (and perhaps a certain amount of surrounding information)” (*i.e. data about technical assets of entities*) “in a pop-up window” (*i.e. user interface*) to allow an “administrator [to] determine whether the matching information and/or the surrounding information ... is associated with the data element being analyzed.” *Id.* In other words, Shull discloses a user interface which displays data that allows an administrator to confirm that one data element (such as an IP address) matches with other data elements (such an entity). EX1003, ¶216.

At the very least, a POSITA would have considered it obvious that Shull's disclosure of “matching information” being confirmed via a user interface includes information that maps IP addresses or domains with the identities of entities. A POSITA would be motivated because Shull describes that “data elements” is “any

discrete piece or set of data.” EX1032, ¶73. Moreover, Shull discloses that “data elements” can be related because they “share a common characteristic.” *Id.*, ¶76. Thus, a POSITA would be motivated to readily apply such teachings to IP addresses or domains and their associated entities. EX1003, ¶217. This provides the known benefit of allowing an administrator to confirm that IP addresses or domains are being correctly assigned to entities, or to manually assign them to entities if needed. *Id.* This allows administrators to increase the accuracy of the mapping by providing human input where needed. *Id.* A POSITA would have a reasonable expectation of success since such a user interface is well-within the capabilities of a POSITA. *Id.*

4. 1[d]

Shull discloses mapping *IP addresses* with entities. IP addresses are *network-related information*. EX1001, 2:1-3; EX1003, ¶218.

5. 1[e]

Shull discloses that the source of information about IP addresses includes “network registration databases” such as ARIN or RIPE, which a POSITA would understand are Regional Internet Registries. EX1032, ¶40; EX1005, 23-24; EX1003, ¶219.

B. Claim 2

See limitation 1[d]. EX1003, ¶220.

C. Claim 3

Shull’s technical assets include IP addresses. EX1003, ¶221.

D. Claim 4

See limitation 1[b]. EX1003, ¶222.

E. Claims 9-11

Shull discloses that the user interface may include “surrounding information” which “related to the matching information,” *i.e. non-technical-asset information*. EX1032, ¶81. A POSITA would understand it includes information about the entities, such as a “corporate records” or “property records,” which includes *industry classifications*. EX1032, ¶40; EX1003, ¶223. A POSITA would also consider it obvious that such “surrounding information” includes any sort of information disclosed in Shull such as an industry classification or description of the entity, since it may be “any discrete piece or set of data.” EX1032, ¶73; EX1003, ¶223.

F. Claim 12

As discussed for limitation 1[c], Shull describes a “pop-up window” that allows an administrator to confirm that an IP address matches with an entity, *i.e. enable the user to assign a technical asset to an entity*. EX1032, ¶81; EX1003, ¶224.

G. Claim 13

Shull discloses that one information source is a “UDRP database” which “contain[s] data related to UDRP complaints filed against cybersquatters and others.” EX1032, ¶40. UDRP stands for “Uniform Domain-Name Dispute-Resolution Policy” and refers to the ICANN policy for resolving when parties have

disputes about domain names, such as when a well-known corporation wishes to use a domain name owned by a cybersquatter. EX1003, ¶225. The UDRP database described in Shull is an *external crowd-sourced service* because it collects information from a large group of individuals, *i.e. a crowd*. *Id.*

H. Claim 14

As discussed in limitation 1[c], Shull provides a mechanism for an administrator to *review and approve* “matching information” being “associated with the data element.” EX1032, ¶81. This includes information about an entity, and a POSITA would understand that the first time an entity’s information is under review, that constitutes *review and approval of an entity proposed to be included in the map*. EX1003, ¶226.

I. Claim 15

Shull discloses that the associations between IP addresses and entities can be used to provide a “reputation score” of the entity. EX1032, ¶¶58-60. This score is generated by a “‘whitelisting’ service,” *i.e. an application*. *Id.*, ¶58. A POSITA would understand that it receives the mapping so that it can generate a score “based on [the entity’s] associations.” *Id.*; EX1003, ¶227.

J. Claim 16

As discussed above for limitation 1[b], Shull discloses that it automatically collects information about technical assets and entities. EX1003, ¶228.

K. Claims 33-37, 42-48, 50-54, 59-65, 67-71, 76-82

As discussed above in Section V.C.1, claims 33-37, 42-48, 50-54, 59-65, 67-71, and 76-82 all recite limitations from claims 1-4 and 9-16. Thus, they are rendered obvious by Shull for the same reasons discussed above. EX1003, ¶229.

XIII. GROUND 7: SHULL, MCNAB, ZHANG

A. Motivation to Combine

A POSITA would have found it obvious to combine Zhang's teaching of mapping IP addresses to an entity by querying a passive DNS database as well as McNab's teaching of IP address enumeration using DNS information with Shull's system of associating IP addresses with entities. EX1003, ¶¶230-39.

As discussed in Ground 3 and the Background, McNab teaches enumerating IP addresses using well-known DNS tools. EX1005, 30-36; EX1003, ¶231.

A POSITA would be motivated to perform these operations with Shull. A POSITA would have been motivated to perform these well-known steps in this well-known order to obtain a comprehensive list of IP addresses associated with a particular entity in order to find IP address-entity associations to be stored in Shull's system. EX1003, ¶232. 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 and therefore likely belonging to the same organization. *Id.*

Moreover, as discussed above in the Background, other tools that performed these three steps were known, including the Maltego tool. EX1003, ¶240. Such tools and techniques were presented by ordinary artisans as beneficial for discovering network assets. A POSITA would have been further motivated to combine Shull 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.*

In addition, Zhang teaches collecting “passive DNS (pDNS) information ..., in which DNS responses are collected passively (e.g., by listening to DNS traffic) from distributed sources (e.g., using passive DNS sensors).” EX1006, ¶47. Zhang discloses that passive DNS 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; *see also id.*, ¶80.

In view of the fact that passive DNS information includes all DNS records that have been queried, a POSITA would be motivated to combine Shull and McNab with Zhang’s disclosure of passive DNS information collection and storage within a pDNS database. A POSITA would be motivated because of the well-known benefits

of passive DNS, namely that DNS records can be structured and stored in a manner that improves search capabilities and reliability. EX1003, ¶235. As Zhang itself teaches, “passive DNS information can be collected, correlated, and analyzed to determine relationships and associations among various entities in DNS (e.g., IP addresses, domains, BGP prefixes, ASs, DNS name servers, and/or other entities.” EX1006, ¶78. For example, passive DNS databases can keep track of changes of DNS records, such as changes to which IP address is mapped to which domain name. *Id.*, ¶81.

Another example is that passive DNS databases can more easily query similarities between records, such as querying which domain names use the same name server or which hostnames are part of the same domain name. EX1003, ¶236. For example, passive DNS database can easily record all hostnames belonging to a certain domain like “apple.com,” including mail (e.g. mail.apple.com) and FTP servers (e.g. ftp.apple.com). *Id.* Such information may not be easily obtainable via conventional DNS queries and often require performing brute force searching. *Id.* Finally, a passive DNS database provides the benefit of a central repository for DNS information, whereas normally such information may require querying multiple DNS servers and name servers. *Id.* Thus, in view of these known benefits, a POSITA would be motivated to include Zhang’s teachings of passive DNS used for

enumerating IP addresses with Shull and McNab's methods for enumerating IP addresses. *Id.*

Zhang also discloses the ability to generate various graphs of relationships among entities and their domains. EX1006, ¶84; EX1003, ¶237. A POSITA would have found it obvious to incorporate these teachings into Shull because, as taught in Zhang, such graphs allow for automatically determining "a reputation of new domains" based on its placement in the graph. *Id.* Zhang teaches that this is beneficial from a security perspective, because a new domain can be determined to be malicious or insecure based on its relationship with known malicious or insecure domains. *Id.*

A POSITA would have had a reasonable expectation of success in combining Shull with McNab and Zhang. EX1003, ¶238. 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 by querying DNS records using Zhang's passive DNS database to obtain the desired information for enumerating IP addresses. *Id.* Such information would then be used to associate IP addresses with an entity to be stored in Shull's system. *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.* A POSITA would also reasonably expect combining Longo with Zhang's graph generation teachings to succeed, since generating graphs using the passive DNS information was well-within the capabilities of a POSITA. *Id.*

As discussed above, Shull, McNab, and Zhang are all analogous art to the '615 patent. EX1003, ¶239.

B. Claims 5-8, 13, 17

As discussed in Ground 3, Zhang and/or McNab disclose the limitations in claims 5-8, 13, and 17, and it would have been obvious to combine those teachings with Shull. EX1003, ¶240.

C. Claims 18-32, 33-49, 55-58, 63, 66, 72-75, 80, 83, 84-100

As discussed in Section V.C.1, claims 18-32, 33-49, 55-58, 63, 66, 72-75, 80, 83, and 84-100 all recite limitations from claims 1, 5-8, 13, and 17. Thus, they are rendered obvious by Shull, McNab and Zhang for the same reasons discussed above. EX1003, ¶241.

XIV. *FINTIV*

All *Fintiv* factors weigh against denial. There has been no requested stay, no trial date, no discovery, and no case schedule. Petitioner challenges 100 claims; Patent Owner asserts only 6 claims. And Petitioner presents compelling grounds.

XV. 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

The '615 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

Petition for *Inter Partes* Review

U.S. Patent No. 9,438,615

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 13,938 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. 9,438,615

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