

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

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

v.

BITSIGHT TECHNOLOGIES, INC.,
Patent Owner

Case IPR2024-01394
Patent 9,438,615

**DECLARATION OF KEVIN C. ALMEROOTH, PH.D. IN SUPPORT OF
PETITION FOR INTER PARTES REVIEW OF
UNITED STATES PATENT NO. 9,438,615**

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
II.	BACKGROUND AND QUALIFICATIONS	2
III.	LEGAL STANDARDS.....	14
	A. Obviousness	14
	B. Claim Construction	19
IV.	PERSON OF ORDINARY SKILL IN THE ART AND THE TIME OF THE ALLEGED INVENTION.....	20
V.	THE '615 PATENT.....	21
	A. Technical Background	21
	1. Overview of the Domain Name System (DNS).....	22
	2. Open Source Intelligence (“OSINT”) Tools.....	24
	B. Description of the '615 Patent’s Specification	40
	C. The '615 patent’s Prosecution History	41
	D. Interpretation of the '615 patent’s Challenged Claims	41
VI.	PUBLIC AVAILABILITY OF CERTAIN EXHIBITS	42
	A. McNab.....	42
	B. Internet Archive Pages.....	43
	C. Public Videos.....	44
VII.	STRUCTURE OF THE CLAIMS OF THE '615 PATENT	46
VIII.	GROUND 1: LONGO AND MCNAB RENDER OBVIOUS CLAIMS 1-4, 9-14, 16, 33-37, 42-46, 48, 50-54, 59-63, AND 65.....	49
	A. Rationale to Combine	49
	B. Claim 1.....	52
	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, ...	52
2.	1[b]: generating of the map comprising online discovery of information about the technical assets; and	54
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, ...	55
4.	1[d]: the technical assets comprising network-related information, ...	56
5.	1[e]: and in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or Regional Internet Registry.	56
C.	Claim 2: The method of claim 1 in which the technical assets comprise network-related information.	57
D.	Claim 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.	57
E.	Claim 4: The method of claim 2 in which generating the map comprises online discovery of information about the technical assets.	57
F.	Claim 9: The method of claim 1 comprising presenting to the user non-technical-asset information through the user interface.	57
G.	Claim 10: The method of claim 9 in which the non-technical-asset information comprises information about the entities.	58
H.	Claim 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.	58
I.	Claim 12: The method of claim 1 in which the tool enables the user to assign a technical asset to an entity.....	58
J.	Claim 13: The method of claim 1 comprising invoking external crowd-sourced services to aid in generating the map.	59
K.	Claim 14: The method of claim 1 comprising enabling a separate review and approval of entities proposed to be included in the map.....	59

L.	Claim 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.....	60
M.	Claims 33-37, 42-46, 48.....	60
N.	Claims 50-54, 59-63, 65.....	61
IX.	GROUND 2: LONGO, MCNAB, AND MALTEGO RENDER OBVIOUS CLAIM 1-4, 9-14, 16-17, 33-37, 42-46, 48-49, 50-54, 59-63, 65-66, 84-88, 93-98, AND 100	61
A.	Maltego Documents	61
B.	Rationale to Combine	61
C.	Limitation 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,.....	63
D.	Claim 9: The method of claim 1 comprising presenting to the user non-technical-asset information through the user interface.	69
	Claim 10: The method of claim 9 in which the non-technical-asset information comprises information about the entities.	69
	Claim 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.....	69
E.	Claim 17: The method of claim 1 comprising generating graphs of relationships among entities based on their associations with technical assets.	70
F.	Remaining limitations and claims.....	71
G.	Claims 33-37, 42-46, 48-49	71
H.	Claims 50-54, 59-63, 65-66	72
I.	Claims 84-88, 93-98, 100.....	72
X.	GROUND 3: LONGO, MCNAB, AND ZHANG (WITH OR WITHOUT MALTEGO) RENDER OBVIOUS CLAIMS 5-8, 13, 17, 18-29, 31-32, 33-46, 48-49, 55-58, 63, 66, 84-98, 100.....	72
A.	Rationale to Combine	72

B.	Claim 5: The method of claim 4 in which the information about the technical assets is discovered through passive DNS queries.	77
C.	Claim 6: The method of claim 5 comprising identifying from the passive DNS queries associations between domain names and network addresses.	78
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.	78
E.	Claim 8.....	78
1.	8[pre]: The method of claim 6 in which identifying from the passive DNS queries associations between domain names and network addresses comprises:.....	78
2.	8[a]: receiving a first domain name for an entity from the entities;	79
3.	8[b]: sending a first passive DNS query to identify first name servers for the first domain name;.....	79
	8[c]: receiving a list of the first name servers for the first domain;.....	79
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;.....	84
	8[e]: receiving a list, for each of the first name servers, of the second domain names for which the name server is authoritative;.....	84
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;	92
	8[g]: receiving a list of the host names and the Internet Protocol addresses for the host names; and	92
6.	8[h]: mapping each of the Internet Protocol addresses to an attribute for the entity.....	99
F.	Claim 13: The method of claim 1 comprising invoking external crowd-sourced services to aid in generating the map.	99
G.	Claim 17: The method of claim 1 comprising generating graphs of relationships among entities based on their associations with technical assets.	99
H.	Claims 18-29, 31-32.....	100
I.	Claims 33-46, 48-49.....	100

J.	Claims 55-58, 63, 66.....	101
K.	Claims 84-98, 100.....	101
XI.	GROUND 4: LONGO, MCNAB, AND TIPPETT (WITH OR WITHOUT MALTEGO) RENDER OBVIOUS CLAIMS 15, 47, 64, 67-71, 76-83, AND 99.....	101
A.	Rationale to Combine	101
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.....	104
C.	Claims 47, 64, 99	104
D.	Claims 67-71, 76-83.....	104
XII.	GROUND 5: LONGO, MCNAB, ZHANG, AND TIPPETT (WITH OR WITHOUT MALTEGO) RENDER OBVIOUS CLAIMS 30, 47, 72-75, 80, 83 AND 99.....	105
A.	Rationale to Combine	105
B.	Claim 30.....	106
C.	Claim 47.....	106
D.	Claims 72-75, 80, 83.....	106
E.	Claim 99.....	107
XIII.	GROUND 6: SHULL RENDERS OBVIOUS CLAIMS 1-4, 9-16, 33-37, 42-48, 50-54, 59-65, 67-71, AND 76-82.....	107
A.	Claim 1.....	107
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, ...	107
2.	1[b]: generating of the map comprising online discovery of information about the technical assets; and	108
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, ...	109

4.	1[d]: the technical assets comprising network-related information,.....	110
5.	1[e]: and in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or Regional Internet Registry.	110
B.	Claim 2: The method of claim 1 in which the technical assets comprise network-related information.	111
C.	Claim 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.	111
D.	Claim 4: The method of claim 2 in which generating the map comprises online discovery of information about the technical assets.	111
E.	Claim 9: The method of claim 1 comprising presenting to the user non-technical-asset information through the user interface.	111
	Claim 10: The method of claim 9 in which the non-technical-asset information comprises information about the entities.	111
	Claim 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.	111
F.	Claim 12: The method of claim 1 in which the tool enables the user to assign a technical asset to an entity.	112
G.	Claim 13: The method of claim 1 comprising invoking external crowd-sourced services to aid in generating the map.	112
H.	Claim 14: The method of claim 1 comprising enabling a separate review and approval of entities proposed to be included in the map.	113
I.	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.	113
J.	Claim 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.	113
K.	Claims 33-37, 42-48, 50-54, 59-65, 67-71, 76-82	114

XIV. GROUND 7: SHULL, MCNAB, AND ZHANG RENDERS OBVIOUS CLAIMS 5-8, 13, 17, 18-32, 33-49, 55-58, 63, 66, 72-75, 80, 83, 84-100	114
A. Rationale to Combine	114
B. Claims 5-8, 13, 17	118
C. Claims 18-32, 33-49, 55-58, 63, 66, 72-75, 80, 83, 84-100	118
XV. CONCLUSION	119

I. INTRODUCTION

1. I have been retained by Normshield Inc. d/b/a Black Kite Inc. (“Petitioner”) as an independent expert consultant in this proceeding before the United States Patent and Trademark Office (“PTO”) against BitSight Technologies, Inc. (“Patent Owner”) regarding U.S. Patent No. 9,438,615 (“the ’615 patent”) (Ex. 1001).¹ I have been asked to submit this Declaration on behalf of Petitioner.

2. I have been asked to consider whether certain references disclose or render obvious the features recited in claims 1-100 (collectively, the “Challenged Claims”) of the ’615 patent. My opinions are set forth below. Based on my experience and expertise, it is my opinion that the prior art renders obvious all limitations of the Challenged Claims, as I discuss in detail below.

3. I am being compensated at a rate of \$850 per hour for my work in this proceeding. My compensation is in no way contingent on the nature of my findings, the presentation of my findings in testimony, or the outcome of this or any other proceeding. I have no other interest in this proceeding.

4. All of my opinions stated in this Declaration are based on my own personal knowledge and professional judgment. I am over 18 years of age and, if I am

¹ Where appropriate, I refer to exhibits that I understand are to be attached to the petition for *Inter Partes* Review of the ’615 patent.

called upon to do so, I would be competent to testify as to the matters set forth in this Declaration.

II. BACKGROUND AND QUALIFICATIONS

5. I am currently a Professor Emeritus in the Department of Computer Science at the University of California, Santa Barbara (UCSB). While at UCSB, I held faculty appointments and was a founding member of the Computer Engineering (CE) Program, Media Arts and Technology (MAT) Program, and the Technology Management Program (TMP). I also served as the Associate Director of the Center for Information Technology and Society (CITS) from 1999 to 2012. I have been a faculty member at UCSB since July 1997.

6. I hold three degrees from the Georgia Institute of Technology: (1) a Bachelor of Science degree in Information and Computer Science (with minors in Economics, Technical Communication, and American Literature) earned in June 1992; (2) a Master of Science degree in Computer Science (with specialization in Networking and Systems) earned in June 1994; and (3) a Doctor of Philosophy (Ph.D.) degree in Computer Science (Dissertation Title: Networking and System Support for the Efficient, Scalable Delivery of Services in Interactive Multimedia System, minor in Telecommunications Public Policy) earned in June 1997. During my education, I took a wide variety of courses as demonstrated by my minor. My undergraduate degree also

included a number of courses more typical of a degree in electrical engineering including digital logic, signal processing, and telecommunications theory.

7. One of the major concentrations of my research over the past 30+ years has been the delivery of multimedia content and data between computing devices, including through various network architectures. In my research, I have studied large-scale content delivery systems, and the use of servers located in a variety of geographic locations to provide scalable delivery to hundreds or thousands of users simultaneously. I have also studied smaller-scale content delivery systems in which content is exchanged between individual computers and portable devices. My work has emphasized the exchange of content more efficiently across computer networks, including the scalable delivery of content to many users, mobile computing, satellite networking, delivering content to mobile devices, and network support for data delivery in wireless networks.

8. In 1992, the initial focus of my research was on the provision of interactive functions (e.g., VCR-style functions like pause, rewind, and fast-forward) for near video-on-demand systems in cable systems; in particular, how to aggregate requests for movies at a cable head-end and then how to satisfy a multitude of requests using one audio/video stream broadcast to multiple receivers simultaneously. This research has continually evolved and resulted in the development of techniques to scalably deliver on-demand content, including audio, video, web documents, and other

types of data, through the Internet and over other types of networks, including over cable systems, broadband telephone lines, and satellite links.

9. An important component of my research has been investigating the challenges of communicating multimedia content, including video, between computers and across networks including the Internet. Although the early Internet was used mostly for text-based, non-real time applications, the interest in sharing multimedia content, such as video, quickly developed. Multimedia-based applications ranged from downloading content to a device to streaming multimedia content to be instantly used. One of the challenges was that multimedia content is typically larger than text-only content, but there are also opportunities to use different delivery techniques since multimedia content is more resilient to errors. I have worked on a variety of research problems and used a number of systems that were developed to deliver multimedia content to users. One content-delivery method I have researched is the one-to-many communication facility called “multicast,” first deployed as the Multicast Backbone, a virtual overlay network supporting one-to-many communication. Multicast is one technique that can be used on the Internet to provide streaming media support for complex applications like video-on-demand, distance learning, distributed collaboration, distributed games, and large-scale wireless communication. The delivery of media through multicast often involves using Internet infrastructure, devices and protocols, including protocols for routing and TCP/IP.

10. Starting in 1997, I worked on a project to integrate the streaming media capabilities of the Internet together with the interactivity of the web. I developed a project called the Interactive Multimedia Jukebox (IMJ). Users would visit a web page and select content to view. The content would then be scheduled on one of a number of channels, including delivery to students in Georgia Tech dorms delivered via the campus cable plant. The content of each channel was delivered using multicast communication.

11. In the IMJ, the number of channels varied depending on the capabilities of the server including the available bandwidth of its connection to the Internet. If one of the channels was idle, the requesting user would be able to watch their selection immediately. If all channels were streaming previously selected content, the user's selection would be queued on the channel with the shortest wait time. In the meantime, the user would see what content was currently playing on other channels, and because of the use of multicast, would be able to join one of the existing channels and watch the content at the point it was currently being transmitted.

12. The IMJ service combined the interactivity of the web with the streaming capabilities of the Internet to create a jukebox-like service. It supported true Video-on-Demand when capacity allowed, but scaled to any number of users based on queuing requested programs. As part of the project, we obtained permission from Turner Broadcasting to transmit cartoons and other short-subject content. We also

connected the IMJ into the Georgia Tech campus cable television network so that students in their dorms could use the web to request content and then view that content on one of the campus's public access channels.

13. More recently, I have also studied issues concerning how users choose content, especially when considering the price of that content. My research has examined how dynamic content pricing can be used to control system load. By raising prices when systems start to become overloaded (i.e., when all available resources are fully utilized) and reducing prices when system capacity is readily available, users' capacity to pay as well as their willingness can be used as factors in stabilizing the response time of a system. This capability is particularly useful in systems where content is downloaded or streamed on-demand to users.

14. As a parallel research theme, starting in 1997, I began researching issues related to wireless devices and sensors. In particular, I was interested in showing how to provide greater communication capability to "lightweight devices," i.e., small form-factor, resource-constrained (e.g., CPU, memory, networking, and power) devices. Starting in 1998, I published several papers on my work to develop a flexible, lightweight, battery-aware network protocol stack. The lightweight protocols we envisioned were similar in nature to protocols like Bluetooth, Universal Plug and Play (UpnP) and Digital Living Network Alliance (DLNA).

15. From this initial work, I have made wireless networking—including ad hoc, mesh networks and wireless devices—one of the major themes of my research. My work in wireless networks spans the protocol stack from applications through to the encoding and exchange of data at the data link and physical layers.

16. At the application layer, even before the large-scale “app stores” were available, my research looked at building, installing, and using apps for a variety of purposes, from network monitoring to support for traditional computer-based applications (e.g., content retrieval) to new applications enabled by ubiquitous, mobile devices. For example, my research has looked at developing applications for virally exchanging and tracking “coupons” through “opportunistic contact” among mobile wireless devices (i.e., communication among devices moving into communication range with each other). In many of the courses I have taught there is a project component. Through these projects I have supervised numerous efforts to develop new “apps” for download and use across a variety of mobile platforms.

17. Toward the middle of the protocol stack, my research has also looked to build wireless infrastructure support to enable communication among a set of mobile devices unaided by any other kind of network infrastructure. These kinds of networks are useful either in challenged network environments (e.g., when a natural disaster has destroyed existing infrastructure) or when suitable support for network communication never existed. The deployment of such networks (or even the use of traditional network

support) are critical to support services like disaster relief, catastrophic event coordination, and emergency services deployment.

18. Yet another theme is monitoring wireless networks, in particular different variants of IEEE 802.11 compliant networks, to (1) understand the operation of the various protocols used in real-world deployments, (2) use these measurements to characterize use of the networks and identify protocol limitations and weaknesses, and (3) propose and evaluate solutions to these problems. I have successfully used monitoring techniques to study wireless data link layer protocol operation and to improve performance by enhancing the operation of such protocols. For wireless protocols, this research includes functions like network acquisition and channel bonding.

19. One theme in my wireless network research has been cross-layer solutions and innovations. As mentioned above, with greater wireless device use and network support, we envisioned new application paradigms and services, for example, when mobile devices come into contact with each other. Instead of relying on existing infrastructure to relay communication, the devices are able to discover each other and communicate directly. Other examples include discovering and using location information to enhance users' experiences. Network support and novel applications span use a variety of network architectures supporting users on foot, in vehicles, and across varying terrains and environments. Finally, we studied how communication

efficiency can be supported through intelligent handoffs as well as location and movement prediction.

20. Protecting networks, including their operation and content, has been an underlying theme of my research almost since the beginning of my research career. Starting in 2000, I have been involved in several projects that specifically address security, network protection, and firewalls. After significant background work, a team on which I was a member successfully submitted a \$4.3M grant proPOSITAl to the Army Research Office (ARO) at the Department of Defense to propose and develop a high-speed intrusion detection system. Key aspects of the system included associating streams of packets and analyzing them for viruses and other malware. Once the grant was awarded, we spent several years developing and meeting the milestones of the project. A number of my students worked on related projects and published papers on topics ranging from intrusion detection to developing advanced techniques to be incorporated into firewalls. I have also used firewalls, including their associated malware detection features, in developing techniques for the classroom to ensure that students are not distracted by online content.

21. My recent work ties some of the various threads of my past research together. I have investigated content delivery in online social networks and proposed reputation management systems in large-scale social networks and marketplaces. On the content delivery side, I have looked at issues of caching and cache placement,

especially when content being shared and the cache has geographical relevance. We were able to show that effective caching strategies can greatly improve performance and reduce deployment costs. Our work on reputation systems showed that reputations have economic value, and as such, creates a motivation to manipulate reputations. In response, we developed a variety of solutions to protect the integrity of reputations in online social networks. The techniques we developed for content delivery and reputation management were particularly relevant in peer-to-peer communication and recommendations for downloadable “apps.”

22. As an important component of my research program, I have been involved in the development of academic research into available technology in the market place. One aspect of this work is my involvement in the Internet Engineering Task Force (IETF). The IETF is a large and open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. I have been involved in various IETF groups including many content delivery-related working groups like the Audio Video Transport (AVT) group, the MBone Deployment (MBONED) group, Source Specific Multicast (SSM) group, the Inter-Domain Multicast Routing (IDMR) group, the Reliable Multicast Transport (RMT) group, the Protocol Independent Multicast (PIM) group, etc. I have also served as a member of the Multicast Directorate (MADDOGS), which oversaw the standardization of all things related to multicast in

the IETF. Finally, I was the Chair of the Internet2 Multicast Working Group for seven years.

23. My involvement in the research community extends to leadership positions for several academic journals and conferences. I am the co-chair of the Steering Committee for the ACM Network and System Support for Digital Audio and Video (NOSSDAV) workshop and on the Steering Committees for the International Conference on Network Protocols (ICNP), ACM Sigcomm Workshop on Challenged Networks (CHANTS), and IEEE Global Internet (GI) Symposium. I have served or am serving on the Editorial Boards of IEEE/ACM Transactions on Networking, IEEE Transactions on Mobile Computing, IEEE Network, ACM Computers in Entertainment, AACE Journal of Interactive Learning Research (JILR), and ACM Computer Communications Review. I have co-chaired a number of conferences and workshops including the IEEE International Conference on Network Protocols (ICNP), IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON), International Conference on Communication Systems and Networks (COMSNETS), IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS), the International Workshop On Wireless Network Measurement (WiNMee), ACM Sigcomm Workshop on Challenged Networks (CHANTS), the Network Group Communication (NGC) workshop, and the Global

Internet Symposium, and I have served on the program committees for numerous conferences.

24. Furthermore, in the courses I taught at UCSB, a significant portion of my curriculum covered aspects of the Internet and network communication including the physical and data link layers of the Open System Interconnect (OSI) protocol stack, and standardized protocols for communicating across a variety of physical media such as cable systems, telephone lines, wireless, and high-speed Local Area Networks (LANs). The courses I have taught also cover most major topics in Internet communication, including data communication, multimedia encoding, and mobile application design. My research and courses have covered a range of physical infrastructures for delivering content over networks, including cable, Integrated Services Digital Network (ISDN), Ethernet, Asynchronous Transfer Mode (ATM), fiber, and Digital Subscriber Line (DSL). For a complete list of courses I have taught, see my curriculum vitae (CV).

25. In addition, I co-founded a technology company called Santa Barbara Labs that was working under a sub-contract from the U.S. Air Force to develop very accurate emulation systems for the military's next generation internetwork. Santa Barbara Labs' focus was in developing an emulation platform to test the performance characteristics of the network architecture in the variety of environments in which it was expected to operate, and, in particular, for network services including Ipv6,

multicast, Quality of Service (QoS), satellite-based communication, and security. Applications for this emulation program included communication of a variety of multimedia-based services, including video conferencing and video-on-demand.

26. In addition to having co-founded a technology company myself, I have worked for, consulted with, and collaborated with companies for nearly 30 years. These companies range from well-established companies to start-ups and include IBM, Hitachi Telecom, Turner Broadcasting System (TBS), Bell South, Digital Fountain, RealNetworks, Intel Research, Cisco Systems, and Lockheed Martin.

27. Through my graduate education, leadership with CITS, involvement in TMP, role in the development of the Internet2 infrastructure, and consulting with ISPs, I have gained a strong understanding in the role of the Internet in our society and the challenges of deploying large-scale production networking infrastructure. CITS, since its inception, has looked at the role of the Internet in society, including how the evolution of technology have created communication opportunities and challenges, including, for example through disruptive technologies like P2P. TMP looks to focus on non-purely technical issues, including, for example, state-of-the-art business methods, strategies for successful technology commercialization, new venture creation, and best practices for fostering innovation. Through my industry collaborations and Internet2 work, I have developed significant experience in the challenges of deploying, monitoring, managing, and scaling communication

infrastructure to support evolving Internet services like streaming media, conferencing, content exchange, social networking, and e-commerce.

28. I am a Member of the Association of Computing Machinery (ACM) and a Fellow of the Institute of Electrical and Electronics Engineers (IEEE).

29. Additional details about my employment history, fields of expertise, courses taught, and publications are further included in my CV attached as Ex. 1004.

30. Based on my professional experience, I believe I am qualified to testify as an expert on matters related to the patent at issue.

III. LEGAL STANDARDS

31. Petitioner' attorneys have explained to me the legal standards that apply in this case. My understanding of those standards is described below. I am not an attorney, and I do not have formal training in the law regarding patents. I have used my understanding of the following legal principles set forth in this section in reaching my opinions.

32. I understand that, in this proceeding, Petitioner have the burden of proving that the challenged claims are invalid by a preponderance of the evidence.

A. Obviousness

33. I have been informed that a claim is invalid as obvious under 35 U.S.C. § 103 (pre-AIA) if the differences between the claimed subject matter and the prior art are such that the subject matter as a whole would have been obvious at the time of the

invention to a person of ordinary skill in the art. I have been informed that the following matters are relevant to determining whether the claimed invention would have been obvious: (1) the scope and content of the prior art, (2) the difference or differences between the patent claim and the prior art, (3) the level of ordinary skill in the art at the time the invention of the patent, and (4) any secondary considerations or objective indicia of non-obviousness.

34. I have been informed that the combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results. When a claim simply arranges prior art elements with each performing the same function it had been known to perform and yields no more than one would expect from such an arrangement, then such a combination is obvious. When a patent claims a structure already known in the prior art altered by the mere substitution of one element for another known in the field, the combination is likely to be obvious unless the combination yields an unpredictable result.

35. I have been informed that when a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill in the art can implement a predictable variation, such a variation is likely unpatentable. For the same reason, if a technique has been used to improve one device, and one of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the

technique is obvious unless its actual application is beyond his or her skill. One question to consider is whether the improvement is more than predictably using prior art elements according to their established functions.

36. I have been informed that it may often be necessary, in a validity analysis, to consider whether there was an apparent reason to combine the known elements in the fashion claimed by the patent at issue. This can be accomplished by looking to interrelated teachings of multiple patents or other publications or pieces of prior art; the effects of demands known to the design community or present in the marketplace; and the background knowledge possessed by one of ordinary skill in the art.

37. I have been informed that a validity analysis need not seek out precise teachings directed to the specific subject matter of the challenged claim; it is appropriate to take account of the inferences and creative steps that a person of ordinary skill in the art would employ. I have been informed that a person of ordinary skill in the art is a person of ordinary creativity, not an automaton.

38. I have been informed that a claim composed of several elements is not proved obvious merely by demonstrating that each element was, independently, known in the prior art. I have been informed that it can be important to identify a reason that would have prompted a person of ordinary skill in the art in the relevant field to combine the elements in the way the claimed new invention does. I am told that one

way that subject matter can be proved obvious is by noting there existed at the time of the invention a known problem for which there was an obvious solution encompassed by the patent's claims. I have been informed that any need or problem known in the field of endeavor at the time of the claimed invention and addressed by the patent can provide a reason for combining the elements in the manner claimed.

39. I have been informed that one should not assume that a person of ordinary skill in the art attempting to solve a problem will be led only to those elements of prior art designed to solve the same problem. Instead, I have been informed that since familiar items may have obvious uses beyond their primary purposes, in many cases a person of ordinary skill in the art will be able to fit the teachings of multiple prior art references together like pieces of a puzzle.

40. I have been informed that, when there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, persons of ordinary skill in the art have good reason to pursue the known options within their technical grasp. If this leads to the anticipated success, the product was likely not accomplished by innovation but by using ordinary skill and common sense. I have been informed that, in such an instance, the fact that the combination was obvious to try may show that the combination was obvious.

41. I have been informed that, when determining whether a claimed combination would have been obvious, the correct analysis is not whether a person of

ordinary skill in the art, writing on a blank slate, would have chosen the particular combination of elements described in the claim. Instead, I have been informed that the correct analysis considers whether one of ordinary skill, facing the wide range of needs created by developments in the field of endeavor, would have seen a benefit to selecting the combination claimed.

42. I have been informed that, when determining whether a claimed invention is obvious, any “secondary considerations” of non-obviousness identified by the patentee should also be considered. These secondary considerations can include:

- commercial success of the invention, causally related to the invention itself rather than to companion factors, such as advertising or attractive packaging;
- the invention taught away from the technical direction followed by those skilled in the art;
- a long-felt but unsatisfied need for the invention while the needed implementing arts and elements had long been available;
- the invention achieves results unexpected to those skilled in the art;

- copying of the invention by competitors as distinguished from their independent development
- unsuccessful attempts by those skilled in the art to make the invention;
- acquiescence by the industry to the patent's validity by honoring the patent through taking licenses or not infringing the patent, or both; and
- skepticism, disbelief in or incredulity by those skilled in the art that the patentee's approach worked.

43. I have been informed that, for the above information to impact the obviousness of a patent claim, there must be a nexus between the alleged secondary considerations and the claims. In addition, I have been informed that the burden of introducing evidence of secondary considerations generally is on the Patent Owner. If the Patent Owner or its expert should assert secondary considerations of non-obviousness, I reserve the right to provide a Declaration addressing assertions of non-obviousness due to secondary considerations.

B. Claim Construction

44. I have been informed that claim terms are typically given their plain and ordinary meanings, as would have been understood by a person of ordinary skill

in the art at the time of the earliest alleged priority date. I have further been informed that when considering the meaning of any terms in the Challenged Claims of the '615 patent, I should apply the plain and ordinary meaning of those terms. I have further been informed that in considering the meaning of the claims, one must consider the language of the claims, the specification, and the prosecution history of record.

45. I have been informed that in general, a preamble limits the invention if it recites essential structure or steps, or if it is necessary to give life, meaning, and vitality to the claim. I have further been informed that a preamble is not limiting where a patentee defines a structurally complete invention in the claim body and uses the preamble only to state a purpose or intended use for the invention. I have further been informed that dependence on a particular disputed preamble phrase for antecedent basis may limit claim scope because it indicates a reliance on both the preamble and claim body to define the claimed invention. I have further been informed that clear reliance on the preamble during prosecution to distinguish the claimed invention from the prior art transforms the preamble into a claim limitation because such reliance indicates use of the preamble to define, in part, the claimed invention.

IV. PERSON OF ORDINARY SKILL IN THE ART AND THE TIME OF THE ALLEGED INVENTION

46. I have been asked to assume that the '615 patent is entitled to its earliest alleged priority date of September 9, 2013. *See* Ex. 1001. Beyond this

assumption, I have not undertaken an analysis to determine the earliest priority date to which the '615 patent is entitled.

47. Based on the materials and information I have reviewed and based on my experience in the technical areas relevant to the '615 patent, a person of ordinary skill in the art at the time of the alleged invention of the '615 patent would have had at least a Bachelor's degree in Computer Science or similar field and at least two years of industry experience in network security. More education can supplement practical experience and vice versa. Based on my knowledge and experience, including as discussed above in Section II, I was a person of ordinary skill in the art at the time of the alleged invention of the '615 patent and can provide opinions regarding the knowledge of a person of ordinary skill in the art as of that time. My opinions herein are, where appropriate, based on my understandings as to a person of ordinary skill in the art at that time. I myself had these capabilities at the time of the alleged invention of the '615 patent.

V. THE '615 PATENT

A. Technical Background

48. U.S. Pat. No. 9,438,615 (the "'615 patent") is directed to mapping or enumerating network assets to entities, including using well-known domain name system (DNS) technologies to discover IP addresses of assets. Below I provide a general background of DNS technologies and their use for network enumeration that

were known for decades before the earliest filing date of the '615 patent, September 9, 2013.

1. Overview of the Domain Name System (DNS)

49. The below overview describes how the DNS system has functioned for decades, including many years before 2013. These technologies are fundamental to how the Internet has been used since the 1980s.

50. Within a network, including the Internet, endpoints like servers and hosts are identified using IP addresses. The most common form of an IP address is known as “IPv4,” which is a series of four numbers, each number ranging from 0 to 255. Every node that is exposed to the Internet has a unique IP address, although it is typical for organizations to create subnetworks within the larger Internet network wherein the nodes in the subnetworks are not directly exposed to the Internet and have IP addresses that are unique within the subnetwork.

51. Web servers, which host websites that Internet users browse, are also identified using an IP address. However, IP addresses are difficult for end-users to memorize. Thus, early on during the development of the Internet, the “Domain Name System” service was developed. The purpose of the DNS service is to map a domain name, such as “google.com,” to an IP address. When an Internet user enters “google.com” in their browser, the browser will query a DNS server to find out the IP address of “google.com” (which is currently 74.125.195.113). When a DNS server

provides this functionality, it is called a “name server.” The records within a name server that map a domain name to an IP address are called “A” records. As discussed below, an A record is a type of DNS record, and there are other types of DNS records as well.

52. Internet users typically query the local name server that their ISP provides. However, every domain name has one or more “authoritative name servers” for which other name servers will rely on to get the authoritative A record and other DNS records for that domain. For example, if a user’s local name server does not have an A record for “google.com,” or there is an indication that the current A record has expired, the local name server will query the authoritative name server for “google.com” to retrieve the A record. Authoritative name servers are often the authoritative name server for multiple domain names, especially domain names of the entity that controls the authoritative name server. A domain name’s authoritative name servers are listed in an NS (Name Server) record, another type of DNS record.

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

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

“drive.google.com” is a hostname of a host within the google.com domain, and its A record maps the hostname to the IP address 74.125.197.100.

2. Open Source Intelligence (“OSINT”) Tools

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

56. Such OSINT tools frequently were used to query the DNS records described above in order to enumerate IP addresses for an entity, and particularly an entity’s domain name. The McNab reference, which is a book entitled “Network Security Assessment” that was published by O’Reilly in 2008, describes several of these well-known tools. For example, McNab describes a tool called “whois,” which can look up various information about a domain name, including its authoritative name servers. McNab at 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
```

57. McNab describes another well-known tool called “nslookup” which “can identify ... hostnames” and “basic domain details” for a given domain, such as “cia.gov” in the example below. McNab at 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
```

58. 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.” McNab at 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.
```

59. McNab explains that using a DNS zone transfer provides information

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

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

60. The dig tool can also query the PTR records of a name server, which can return other domain names for which that name server is authoritative. In the example below, the dig tool queries the name server “relay2.ucia.gov” for all PTR records that are within the 198.81.129.0 network block. McNab at 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.
```

61. 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.” McNab at 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.* at 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  
-----
```


62. Several of the above tools require making multiple queries of DNS records, such as the above TXDNS example which performed at least 98 DNS queries. To increase the efficiency of frequent DNS queries, a technology known as “passive DNS” was created, wherein a passive DNS server and/or sensors (which are typically on the same machine as a regular DNS server) “collect[s] passively” DNS records “by listening to DNS traffic.” Zhang at ¶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.* at ¶79.

The passive DNS server can directly store “DNS resource records” as well. *Id.* at ¶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.* at ¶97.

63. Network enumeration using DNS technologies was also described in the Maltego tool released in the late 2000s. For example, Exhibit 1018 is a presentation given in 2008 by a well-known security researcher named Chris Gates. *See also* Ex. 1019. The presentation describes a “[b]lack box approach to info gathering” where “the customer wants you to find everything about them you can via ‘open source’ methods.” Ex. 1018 at 5. OSINT information gathering involves “no direct contact with victim’s servers OR no non-standard traffic directed toward victim.” *Id.* at 9.

64. One of the tools the presentation describes is Maltego, which “is a program that can be used to determine the relationships and real world links between: People ... Companies ... Web sites ... Internet infrastructure such as: Domains, DNS names, Netblocks, IP addresses” and it does this “using open source intelligent (OSINT).” *Id.* at 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. Ex. 1018 at 70.




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



Maltego

- See what hostnames are shared by the DNS servers




66. 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.* at 72.




Maltego

- Brute force our domain names and resolve them to IPs




67. 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.


68. For a given domain name, its NS record indicating its authoritative NS server can be retrieved. Ex. 1021 at 0:14.



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




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




71.

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



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



74. 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.

75. For example, in the same presentation given by Chris Gates, another tool called Fierce DNS is described. Ex. 1018 at 18. Fierce DNS “locate[s] likely targets both inside and outside a corporate network.” *Id.* at 18. Fierce DNS uses a number of DNS queries to find hosts on the same domain and IP blocks used by the domain. *Id.* at 20. The below output for “bytsp.com” shows the various hosts for that domain name. *Id.* at 21-22.



Fierce DNS

```
cg@WPAD:~/evil/enumeration/dns/fierceDNS$ perl fierce.pl -dns baytsp.com
```

```
DNS Servers for baytsp.com:
```

```
ns1.baytsp.net  
ns13.zoneedit.com  
ns2.baytsp.net
```

```
Trying zone transfer first...
```

```
Testing 216.133.204.226  
Request timed out or transfer not allowed.  
Testing 66.223.40.121  
Request timed out or transfer not allowed.  
Testing 216.132.49.229  
Request timed out or transfer not allowed.
```

```
Unsuccessful in zone transfer (it was worth a shot)  
Okay, trying the good old fashioned way... brute force
```

```
Checking for wildcard DNS...
```


```
Nope. Good.
```

```
Now performing 3149 test(s)...
```

```
216.133.204.235 dmz.baytsp.com  
216.133.204.231 bayreports.baytsp.com  
216.133.204.229 ftp.baytsp.com  
216.133.204.224 dmz.baytsp.com  
216.133.204.225 baytsp-gw.baytsp.com  
216.133.204.232 cat-3000.baytsp.com  
216.133.204.233 cims4.baytsp.com  
216.133.204.234 dmz.baytsp.com  
216.133.204.236 dmz.baytsp.com  
216.133.204.239 dmz.baytsp.com  
216.133.204.240 tc2.baytsp.com  
216.133.204.241 dmz.baytsp.com
```

CarnalOwne

11/10/94
Declassified by [redacted] b7c
Declassify on: OADR
CA# 94-1720 CRR



76. In a 2007 presentation given at the Blackhat conference provides an overview of “gathering and discovery techniques, with a concentration on third-party services and new tools” and then using such tools to “gain[] access to a target network.” Ex. 1017 at 4. Thus, in Chapter 3, the paper states “[t]he first step to any security engagement is the initial discovery process,” and describes “discovery techniques, starting from the outside.” *Id.* at 7. The chapter goes to explain how to use tools to perform “network discovery” using “DNS zone transfers, Whois requests, and reverse DNS lookups.” *Id.* at 9. The chapter then provides descriptions of more advanced tools such as a “‘Domain Dossier’ service which combines the various DNS and Whois

requests into [a] single report, with the option to perform a quick port scan as well.”

Id. There is also the “DomainTools.com web site” which “accepts an IP address or host name as an input and provides a list of all domains that reverse back to that IP.”

Id.

77. All of the above techniques were well-known before the earliest filing date of the ’615 patent.

B. Description of the ’615 Patent’s Specification

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

79. Although the “Background” of the ’615 patent describes the patent as “relate[d] to security risk management” where an “entity may evaluate the magnitude of the risks associated with the other entities to make decisions about its relationships with those other entities,” the 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. Several of the dependent claims further recite three passive DNS queries to obtain IP addresses for an entity: (1) a first query that uses the entity’s domain name to receive a list of the domain name’s name servers, (2) for each name server, a second query that receives a list of other domain names for which that name server is authoritative, and (3) for each of the new

domain names, a third query that receives a list of hostnames and IP addresses for the respective domain name. These IP addresses are then mapped to the entity.

80. The specification of the '615 patent does not explain the purported novelty of its claims, nor does it claim to invent passive DNS. Nor does the '615 patent claim that mapping network assets or entities is novel.

C. The '615 patent's Prosecution History

81. I have reviewed the prosecution history of the '615 patent and summarize it below.

82. 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.

D. Interpretation of the '615 patent's Challenged Claims

83. For the purposes of my analysis in this IPR proceeding, I understand that the words of a claim are given their plain meaning that those words would have had to a POSITA at the time of the alleged invention. I also understand that the structure of the claims, the specification, and the prosecution history may also be used to better construe a claim insofar as the plain meaning of the claims cannot be understood. Moreover, I understand that even treaties and dictionaries may be used, albeit under limited circumstances, to determine the meaning attributed by a POSITA

to a claim term at the time of filing. Furthermore, I understand that a Patent Owner's own apparent interpretation of certain terms in related proceedings can be considered to determine the meaning of patent claims in an IPR proceeding.

84. I have followed this approach in my analysis, and I have applied the plain and ordinary meaning of those terms as they would have been interpreted by a POSITA at the time the invention was made (not today). For purposes of my analysis here, I have used the September 9, 2013 filing date of the '615 patent as the date of the invention.

VI. PUBLIC AVAILABILITY OF CERTAIN EXHIBITS

85. I have been asked to opine on the public availability of certain exhibits and specifically whether they were publicly available by September 9, 2013.

A. McNab

86. In my opinion, McNab was publicly available by 2008.

87. McNab is a book published by O'Reilly Media. I am familiar with O'Reilly Media and know that they are a very well-known and reputable publisher of books related to computer technology. I have reviewed dozens of O'Reilly-published books and have cited many such books in my own publications.

88. In my opinion, based on both my expertise in the field and from my personal experience, a POSITA would have considered O'Reilly Media to be a reputable publisher such that the copyright dates on their publications are a reliable and

trustworthy indicator of when the publication was published. Ordinary artisans rely on the copyright date printed on the inside of O'Reilly-published books. I am not aware of nor have I heard of any O'Reilly-published book having a copyright date that is not commensurate with the date of its public availability.

89. In my opinion, Exhibit 1005 is also an authentic copy of McNab as it was available in 2008. Based on my significant experience as a reader of O'Reilly-published books, Exhibit 1005 is complete and has no visible alterations. Exhibit 1005 also has all the hallmarks of an authentic copy of an O'Reilly-published book, including a cover page, a copyright page, a table of contents, an index, and a back cover page. There are also no missing pages from Exhibit 1005. Thus, in my opinion, Exhibit 1005 is a true and authentic copy of McNab as it was published in 2008.

B. Internet Archive Pages

90. Several of the exhibits I have reviewed are from the Internet Archive, including Exhibits 1017, 1020, and 1025. I am familiar with the Internet Archive and have reviewed the declaration of Mr. Frank-White regarding the operation of the Internet Archive. Ex. 1013.

91. The Internet Archive exhibits are printouts of web pages or files that were archived by the Wayback Machine. Each link to the Wayback machine has a date and timestamp which indicates the date and time that the link was archived. In my experience this date and time an accurate reflection of when the Wayback Machine

archives that particular page. Being familiar with the operation of the World Wide Web and crawlers like those used by the Wayback Machine, a page could not be archived by the Wayback Machine unless it was available to the public. Moreover, crawlers generally operate by accessing pages that were linked by other pages, and therefore a page that was archived by the Wayback Machine is typically archived because it was linked from at least one other webpage. This further indicates that the archived webpage was publicly available at the time it was archived.

C. Public Videos

92. Exhibits 1019 and 1021 are printouts of webpages that host a video that was publicly available on Vimeo and YouTube, respectively.

93. Exhibit 1019 (available at <http://vimeo.com/2745624>) is a video on the Vimeo website entitled “Toorcon X Gates: New School Information.” This video is of Chris Gates presenting the presentation “New School Information Gathering,” a copy of which is at Exhibit 1018. The webpage states that it was uploaded on January 6, 2009. In my opinion, the upload date on a Vimeo webpage is generally a reliable indicator of when a video was publicly available. Ordinary artisans would have found the date on a Vimeo webpage to be a reliable indicator of public availability. I also note that Exhibit 1019 is a confirmation from an employee at Vimeo that the video was uploaded on January 6, 2009.

94. In addition, I note that Exhibit 1020 is the version of Exhibit 1019 as it was archived by the Wayback Machine on March 2, 2009. This indicates that the video was at least publicly available by March 2, 2009.

95. Thus, in my opinion, Exhibits 1018, 1019 and 1020 were all publicly available by early 2009. It is also my opinion that they are true and correct copies of the video on the date they were published, since the video or webpage does not have any indication it was edited or altered since its upload date.

96. Exhibit 1021 (available at <https://www.youtube.com/watch?v=qiv4-wy3mxo>), is a video on the YouTube website entitled “Social Engineering Using Maltego.” Exhibit 1031 are screenshots of relevant portions from this video. The webpage states it was uploaded on September 14, 2009 and that it has thousands of views. In my opinion, the upload date on a YouTube webpage and its view count is generally a reliable indicator of when a video was publicly available. Ordinary artisans would have found the date on a YouTube webpage to be a reliable indicator of public availability. Moreover, the metadata for this 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. Ex. 1030.

97. I also note that the video itself, which is of a person using the Maltego software, often displays various dates. All of these dates are consistent with a public

availability date of 2009. For example, at 0:20, the video shows an NS record that is dated March 23, 2009.

98. Thus in my opinion, Exhibit 1021 was publicly available by 2009. It is also my opinion that they are true and correct copies of the video on the date they were published, since the video or webpage does not have any indication it was edited or altered since its upload date.

VII. STRUCTURE OF THE CLAIMS OF THE '615 PATENT

99. The '615 patent has 100 claims. However, there is a significant amount of overlap between the claims. There are 6 independent claims: claims 1, 18, 33, 50, 67, and 84. Each set of independent claims and corresponding dependent claims recite the same limitations, but in differing order. Specifically, every independent claim recites the same two limitations: “generating a map between ...” and “enabling a user to assist in the generating of the map ...” limitations. However, each independent claim also recites one or more additional limitations that are otherwise found in the dependent claims for the other independent claims. For example, Claim 1 recites “in which the information about the technical assets is discovered from an Internet Assigned Numbers Authority or Regional Internet Registry,” which is not recited in the other independent claims but is recited in one of their dependent claims (*i.e.* claims 20, 37, 54, 71 and 88). Thus, the scope of claims 18-100 are identical to claims 1-17.

100. The below chart shows how each limitation is mapped to each claim within each claim set.

	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 of information about the technical assets”	1[b], 4	18[c]	36	53	70	87
“enabling a user to assist in the generating of the map ...”	1[c]	18[e]	33[c]	50[c]	67[c]	84[c]
“technical assets comprising 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
“presenting ... non-technical-asset information through the user interface”	9	24	42	59	76	93

	Independent Claim Set					
Claim Language	1	18	33	50	67	84
“non-technical-asset information comprises information about the entities”	10	25	43	60	77	94
“descriptions of entities, industry classifications, and employee counts”	11	26	44	61	78	95
“tool enables the user to assign a technical asset to an entity”	12	27	45	62	79	96
“invoking external crowd-sourced services to aid in generating the map”	13	28	33[b]	63	80	97
“separate review and approval of entities proposed to be included in the map”	14	29	46	50[b]	81	98
“providing the map to an application for joining to event data or scoring the security state of the entities”	15	30	47	64	67[b]	99
“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”	16	31	48	65	82	100
“generating graphs of relationships among entities	17	32	49	66	83	84[b]

	Independent Claim Set					
Claim Language	1	18	33	50	67	84
based on their associations with technical assets”						

VIII. GROUND 1: LONGO AND MCNAB RENDER OBVIOUS CLAIMS 1-4, 9-14, 16, 33-37, 42-46, 48, 50-54, 59-63, AND 65

A. Rationale to Combine

101. A POSITA would have found it obvious to combine McNab’s teachings of enumerating assets of an entity by discovering their IP addresses with Longo’s teachings of assigning IP addresses to entities.

102. McNab is an O’Reilly-published book entitled “Network Security Assessment” and describes various tools and methodologies for assessing the security of a network. McNab at 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. McNab at 4-7. For example, McNab discloses various methods such as “DNS Querying” and a “zone transfer” to obtain IP addresses of computer within a domain. McNab, 30-35. McNab also describes querying Regional Internet Registries such as ARIN and RIPE to obtain WHOIS information about network assets. McNab, 23-24.

103. A POSITA would be motivated to incorporate such teachings because they were well-known, effective and conventional methods of determining IP addresses of entities. Indeed, Longo describes that knowing which IP addresses belong

to which entities has “valuable uses.” Longo, ¶5. And the very goal of Longo is to store “actual associations between the candidate address and the organizational entity.” *Id.*, Abstract.

104. Given these teachings in Longo, a POSITA would have been motivated to give such teachings effect by using the techniques disclosed in McNab to identify IP addresses for a particular entity. 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). Network enumeration and security risk assessment were also well-known and ubiquitous methods for helping secure an organization’s networks. 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.

105. 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. 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.” McNab, 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.* at 2.

106. 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.*

107. A POSITA would have reasonably expected success in combining Longo with McNab’s teachings. 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. 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. Longo, ¶15.

108. Finally, Longo, McNab, and the '615 patent are all directed to the same field of determining network assets of an entity. Longo at Abstract (“Network addresses (e.g., IP addresses) are considered candidates for association with organizational entities”); McNab at 5 (“Internet Host and Network Enumeration”); Ex. 1001 at Abstract (“[A] map is generated between (a) technical assets that contribute to security characteristics of respective entities.”).

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,

109. With respect to the limitation “*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. Longo, Abstract. 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*.

110. With respect to the limitation “*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*. 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) ...” McNab, 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. In the alternative, as discussed above, it would have been obvious to combine Longo with these teachings in McNab.

111. With respect to the limitation ***“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. 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.” Longo, ¶¶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

112. 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.” Longo, ¶17. ARIN and RIPE are known as “Regional Internet Registries (RIRs)” “under which all public Internet-based address spaces [for a region] fall.” McNab, 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*.

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,

113. 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. Longo, ¶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.; see also id.*, ¶15 (“[C]orrelation engine 135 ... processes the interaction information and identifies which addresses genuine[ly] belong to which organizational entities.”).

114. A POSITA would understand that by allowing a user to specify an entity to which they belong to, this enables a user to assist in generating a map by using a tool (such as an entity entry form or selector) that associates the user’s IP address (a technical asset) with the selected entity.

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

115. Longo discloses mapping *IP addresses* with entities. A POSITA would understand that *network-related information* includes IP addresses, which the ’615 patent also confirms. *See* Ex. 1001, 2:1-3 (“The network-related information includes ... Internet Protocol addresses ...”).

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

116. As I discussed above, Longo discloses using ARIN or RIPE to discover information about technical assets. Longo, ¶17. A POSITA would

understand that ARIN and RIPE are the *Regional Internet Registries* for North America and Europe, respectively. McNab, 23-24.

C. Claim 2: The method of claim 1 in which the technical assets comprise network-related information.

117. This limitation is identical to limitation 1[d], so is disclosed by Longo for the reasons I discussed above.

D. Claim 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.

118. As I discussed above for limitation 1[d], Longo's technical assets include *IP addresses*. Longo, Abstract.

E. Claim 4: The method of claim 2 in which generating the map comprises online discovery of information about the technical assets.

119. This limitation is identical to limitation 1[b], so is rendered obvious by Longo for the reasons I discussed above.

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

120. Longo discloses that the user interface includes "an organization selection control" which "auto-complete[s] organization names entered on the form." Longo, ¶19. Such information is "non-technical-asset information" because it would include names of organizations and entities with which the asset is not associated. For

example, a POSITA would understand that the selection control could be a drop-down box (which were known to auto-completes as information is entered) containing known entities to select from. Or if the form is a text box, the auto-completion feature would auto-complete different entity names as the entity is being typed. For example, when a user types “m”, it may auto-complete as “Meta, Inc.” but when the user continues to type “mi,” it auto-completes as “Microsoft, Inc.”

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

121. As I discussed above for claim 9, Longo discloses “non-technical asset information” in the form of “organization names,” *i.e. information about the entities*. Longo, ¶19.

H. Claim 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.

122. As I discussed above for claim 9, Longo discloses “non-technical asset information” in the form of “organization names,” and in some cases, the names are *descriptions of the entity* that it names, such as “Legal Department, NPE Suit Mfg. Ltd.” Longo, ¶43.

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

123. As I 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*. Longo, ¶19.

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

124. McNab discloses obtaining information about entities by “scour[ing] ... newsgroups,” which a POSITA would understand to be an *external crowd-sourced service* because information in newsgroups are sourced from the numerous users of that newsgroup, *i.e.* a crowd. McNab, 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.” Longo, ¶18.

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

125. 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, *i.e. a separate review and approval of entities proposed 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. Longo, ¶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.

L. Claim 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.

126. As I 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.” Longo, ¶¶14-15.

M. Claims 33-37, 42-46, 48

127. As I discussed above in Section VII, independent Claim 33 recites the same limitations as limitations 1[a], 1[c] and claim 13. Thus, it is rendered obvious by Longo and McNab for the same reasons discussed above. Dependent claims 34-37, 42-46 and 48 also recite the same limitations as limitations found in dependent claims 1-4, 9-12, 14, and 16 and therefore are also rendered obvious by Longo and McNab for the same reasons discussed above.

N. Claims 50-54, 59-63, 65

128. As I discussed above in Section VII, independent Claim 50 recites the same limitations as limitations 1[a], 1[c] and claim 14. Thus, it is rendered obvious by Longo and McNab for the same reasons discussed above. Dependent claims 51-54, 59-63, and 65 also recite the same limitations as limitations found in dependent claims 1-4, 9-13, and 16 and therefore are also rendered obvious by Longo and McNab for the same reasons discussed above.

IX. GROUND 2: LONGO, MCNAB, AND MALTEGO RENDER OBVIOUS CLAIM 1-4, 9-14, 16-17, 33-37, 42-46, 48-49, 50-54, 59-63, 65-66, 84-88, 93-98, AND 100

129. Maltego provides an alternative basis for rendering obvious limitation 1[c] and claims 9-11, and further discloses claim 17.

A. Maltego Documents

130. Maltego, and in particular its user interface, is described in Exhibits 1018, 1021, and 1025. As I discussed above, each of these exhibits were publicly available before September 9, 2013.

B. Rationale to Combine

131. In my opinion, 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. In particular, Maltego was a third-party software that allowed a user to assign an IP address to an organization using drag-and-drop.

132. 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. 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. 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.

133. 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.

134. Finally, Maltego is analogous art to the '651 patent because it is directed to enumerating IP addresses for an entity. *See, generally*, Ex. 1018, Ex. 1021.

C. Limitation 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,


135. 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*.

136. With respect to the limitation “*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. 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). Ex. 1018, 70.



Maltego

- Find our MX and NS servers



CarnalOwne


11/10/94
Classified by 5668
Declassify on: OADR
CA# 94-1720 CRR b7c

137. 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. Ex. 1018, 71.




Maltego

- See what hostnames are shared by the DNS servers



138. This is also shown in the YouTube demonstration of Maltego. Ex. 1021, 0:41; Ex. 1031, 6.




139. 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. Ex. 1025, 8.



140. With respect to the limitation “*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*). 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. Ex. 1018, 73.



141. 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. Ex. 1021, 5:16, 5:34; Ex. 1031, 31, 33.

The screenshot shows the Maltego v2.0.2 interface. The main window displays a network graph with nodes representing email addresses and domains. A context menu is open over the graph, listing various transforms for domain analysis. The 'Output - Transform execution' pane at the bottom shows the results of the transforms.

Transforms listed in the context menu:

- To DNS Name [Brute] Paterva CTAS
- To DNS Name [ZoneTransfer] Paterva CTAS
- To Documents [SE] Paterva CTAS
- To Domain [more TLDs] Paterva CTAS
- To Email address [whois] Paterva CTAS
- To Email addresses [domain only] [SE] Paterva CTAS
- To Email addresses [PGP] Paterva CTAS
- To Email addresses [SE] Paterva CTAS
- To Files [SE] Paterva CTAS
- To MX record [DNS] Paterva CTAS
- To NS record [DNS] Paterva CTAS
- To Person [PGP] Paterva CTAS

Output - Transform execution:

```

Transform "To Social networks [Rapleaf]" completed with 4 results
- One or more social memberships are still in the process of being determined on email address pablo@...
Transform "To Social networks [Rapleaf]" completed with 5 results
Transform "To Social networks [Rapleaf]" completed with 5 results
Transform "To Social networks [Rapleaf]" completed with 1 results
- One or more social memberships are still in the process of being determined on email address prana@slather.com
  
```

Entity properties:

Entity type	Value	Weight
Website	www.inalaska.net	100

GENERATION INFORMATION:

Result	Source	Gen. by
www.inalaska.net (Website)	inalaska.net (Domain)	To Website [quick lookup] 2009-3-23 15:48

...get web sites for each domain...

The screenshot shows the Maltego v2.0.2 interface. The main window displays a network graph with nodes representing websites and domains. A context menu is open over the graph, listing various transforms for website analysis. The 'Output - Transform execution' pane at the bottom shows the results of the transforms.

Transforms listed in the context menu:

- To Domains [DNS] Paterva CTAS
- To Email Address [Mirror] Paterva CTAS
- To IP Address [DNS] Paterva CTAS
- To Phrase [Blog Tag Technorat] Paterva CTAS
- To URLs [Dump] Paterva CTAS
- To Website Title Paterva CTAS
- To Website [coming links SE] Paterva CTAS
- To Website [Mirror] Paterva CTAS
- To Website [Thumbnail] Paterva CTAS
- _To Webdir (directories found using SE) Paterva PT
- _To Website Title Paterva PT

Output - Transform execution:

```

Transform "To Website [quick lookup]" completed with 1 results
Transform "To Website [quick lookup]" completed with 1 results
Transform "To Website [quick lookup]" completed with 1 results
Transform "To Website [quick lookup]" completed with 1 results
Transform "To Website [quick lookup]" completed with 1 results
Transform "To Website [quick lookup]" completed with 1 results
  
```

Entity properties:

Entity type	Value	Weight
Website	www.inalaska.net	100

GENERATION INFORMATION:

Result	Source	Gen. by
www.inalaska.net (Website)	inalaska.net (Domain)	To Website [quick lookup] 2009-3-23 15:48

...and resolve their IP addresses.

D. 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, industry classifications, and employee counts.

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

143. For example, Maltego's user interface displays the domain name and weight of an entity, which are *descriptions of the entities*. 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. Ex. 1025, 8.




144. Another example is that Maltego's user interface can display email addresses associated with the entity, which corresponds with an *employee count*. Ex. 1018, 74.



Maltego

- Harvest emails for the domain and verify them



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


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

146. 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.”

Ex. 1025. As Exhibit 1018 describes it, this graph shows “relationships” between entities, and depicts each entity based on associations with technical assets. Ex. 1018,

80.

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



F. Remaining limitations and claims

147. 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.

G. Claims 33-37, 42-46, 48-49

148. As discussed above in Section VII, independent Claim 33 recites the same limitations as limitations 1[a], 1[c] and claim 13. Thus, it is rendered obvious by Longo, McNab and Maltego for the same reasons discussed above. Dependent claims 34-37, 42-46 and 48-49 also recite the same limitations as limitations found in dependent claims 1-4, 9-12, 14, and 16-17 and therefore are also rendered obvious by Longo, McNab, and Maltego for the same reasons discussed above.

H. Claims 50-54, 59-63, 65-66

149. As discussed above in Section VII, independent Claim 50 recites the same limitations as limitations 1[a], 1[c] and claim 14. Thus, it is rendered obvious by Longo, McNab and Maltego for the same reasons discussed above. Dependent claims 51-54, 59-63 and 65-66 also recite the same limitations as limitations found in dependent claims 1-4, 9-13, and 16-17 and therefore are also rendered obvious by Longo, McNab, and Maltego for the same reasons discussed above.

I. Claims 84-88, 93-98, 100

150. As discussed above in Section V.C.1, independent Claim 84 recites the same limitations as limitations 1[a], 1[c] and claim 17. Thus, it is rendered obvious by Longo, McNab and Maltego for the same reasons discussed above. Dependent claims 85-88, 93-98, and 100 also recite the same limitations as limitations found in dependent claims 1-4, 9-14, and 16 and therefore are also rendered obvious by Longo, McNab, and Maltego for the same reasons discussed above.

X. GROUND 3: LONGO, MCNAB, AND ZHANG (WITH OR WITHOUT MALTEGO) RENDER OBVIOUS CLAIMS 5-8, 13, 17, 18-29, 31-32, 33-46, 48-49, 55-58, 63, 66, 84-98, 100

A. Rationale to Combine

151. 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.

152. As I 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. McNab, 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.*

153. 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. 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.

154. 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. 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.

155. 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).” Zhang, ¶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 (“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)").

156. 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. 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." Zhang, ¶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.

157. 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. 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. 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.

158. Zhang also discloses the ability to generate various graphs of relationships among entities and their domains. Zhang, ¶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.*

159. A POSITA would have had a reasonable expectation of success in combining Longo and McNab with Zhang. 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. 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. 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. 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.

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

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

161. As I discussed above in the Rationale to Combine section, 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.

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

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

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.

163. McNab discloses that its “Network Enumeration” techniques, which include deriving IP addresses for hostnames, are repeated on a cyclic basis. McNab, 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. 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:

164. As I discuss in the following limitations, 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 I discussed above, it would have been obvious to perform such DNS queries on Zhang’s passive DNS database.

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

165. Longo discloses that the organization databases “lists domain[s]” for the stored organizations, *i.e. domain name for an entity from the entities*. Longo, ¶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.

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;

166. Zhang discloses using “passive DNS information” to collect information about a “first domain,” *i.e. first domain name*. Zhang, ¶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.

167. 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. Zhang, ¶¶79-80. It was well-known that DNS records and WHOIS information includes the authoritative name server of a particular domain. 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.” McNab, 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:

first domain name

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

list of first name
servers

Status: clientTransferProhibited
Updated Date: 04-oct-2006
Creation Date: 20-mar-1995
Expiration Date: 21-mar-2009

Id., 22.


168. 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.” McNab, 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.

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
```




169. 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.

170. 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. 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. Ex. 1018, 70.



171. Similarly, in the YouTube video demonstrating Maltego, the “name servers (NS) for this domain” are retrieved. Ex. 1021, 0:14; Ex. 1031, 3. This retrieves two NS servers for the domain name “shmoo.com”: “krusty.shmoo.com” and “archimedes.shmoo.com.” *Id.*, 0:19; Ex. 1031, 4.



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;

172. 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*. Zhang, ¶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.

173. 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. Zhang, ¶¶79-80. It was well-known that DNS records and Whois information records Name Server information for domain names. 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. McNab, 34.

174. 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.” McNab, 34-35. McNab further teaches that this query is executed for “all accessible name servers.” *Id.*, 35.

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.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.
```


list of second domain names

175. A POSITA would understand that the query will return domain names *for which the designed name server is authoritative*. This is because the dig command queries DNS records (specifically PTR records) for the designated name server, for example, “relay2.ucia.gov.” The DNS records for the designated name server *must* include records for which it is authoritative—since the “authoritative name server” is required to be the name server of last resort for retrieving any DNS records for which it is authoritative. *See, e.g.,* Ex. 1007, 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. 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].

176. 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.

177. In the alternative, this limitation is disclosed under Patent Owner’s infringement theory for an identical limitation for the related patent U.S. Pat. No. 11,652,834. The ’834 patent recites the same three passive DNS queries, including the same second query to name servers to receive domain names for which the name server is authoritative. For said second query, Patent Owner accuses the use of WHOIS queries. Ex. 1027, 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.*



178. 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*. 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.


179. 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.

180. 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. 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.” Ex. 1018, 71.




Maltego


- See what hostnames are shared by the DNS servers



181. 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.” Ex. 1021, 0:19; Ex. 1031, 4.



182. This results in domain names such as “kusia.net” and “legacyhigh.com.” *Id.*, 0:41; Ex. 1031, 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

183. 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.” Zhang, ¶33.

184. 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. Zhang, ¶¶79-80.

185. 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*. McNab, 35-36. 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
-----
```

186. In this example, the intent was to enumerate mail server hosts. 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. McNab, 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

187. 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”). McNab, 33-34. As shown below, the list includes IP addresses for some of the hosts.

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

188. 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.

189. 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. 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. Ex. 1018, 72.



Maltego

- Brute force our domain names and resolve them to IPs



190. Similarly, in the YouTube video demonstrating Maltego, for each domain name discovered, Maltego can further find the hostnames for each domain. Ex. 1021, 5:16; Ex. 1031, 30. Maltego also resolves the IP address for each hostname. *Id.*, 5:34; Ex. 1031, 32.

Maltego v2.0.2

File Edit View Navigate Tools Window Help

Speed/Accuracy # Results

New Graph (1) * x

Mining View Centrality View Edge Weighted View

Properties

Entity properties

All

All Transforms

About the Domain entity

To DNS Name [Brute] Paterva CTAS

To DNS Name [ZoneTransfer] Paterva CTAS

To Documents [SE] Paterva CTAS

To Domain [more TLDs] Paterva CTAS

To Email address [whois] Paterva CTAS

To Email addresses [domain only] [SE] Paterva CTAS

To Email addresses [PGP] Paterva CTAS

To Email addresses [SE] Paterva CTAS

To Files [SE] Paterva CTAS

To MX record [DNS] Paterva CTAS

To NS record [DNS] Paterva CTAS

To Person [PGP] Paterva CTAS

Output - Transform execution

Transform "To Social networks [Rapleaf]" completed with 4 results

- One or more social memberships are still in the process of being determined on email address pablo@slather.com

Transform "To Social networks [Rapleaf]" completed with 5 results

Transform "To Social networks [Rapleaf]" completed with 5 results

Transform "To Social networks [Rapleaf]" completed with 1 results

- One or more social memberships are still in the process of being determined on email address prana@slather.com

5:16 / 6:50

...get web sites for each domain...

Maltego v2.0.2

File Edit View Navigate Tools Window Help

Speed/Accuracy # Results

New Graph (1) * x

Mining View Centrality View Edge Weighted View

Properties

Entity properties

Entity type Website

Value www.inalaska.net

Weight 100

Detail View

GENERATION INFORMATION

Result	Source	Gen. by
www.inalaska.net (Website)	inalaska.net (Domain)	To Website [quick lookup]

Output - Transform execution

Transform "To Website [quick lookup]" completed with 1 results

Transform "To Website [quick lookup]" completed with 1 results

Transform "To Website [quick lookup]" completed with 1 results

Transform "To Website [quick lookup]" completed with 1 results

Transform "To Website [quick lookup]" completed with 1 results

Transform "To Website [quick lookup]" completed with 1 results

Found 17 entries

5:34 / 6:50

...and resolve their IP addresses.

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

191. As I 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 I discussed in the Rationale 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.

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

192. 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)." Zhang, ¶47. Thus, passive DNS information is *crowd-sourced* because it is built from all the DNS traffic from many users that can be detected.

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

193. 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.” Zhang, ¶¶84, 86. These relationships are *based on the entities associations with technical assets*, such as having the same IP address.

H. Claims 18-29, 31-32

194. As discussed above in Section VII, independent Claim 18 recites the same limitations as limitations 1[a] through 1[d] and claim 5. Thus, it is rendered obvious by Longo, McNab and Zhang (with or without Maltego) for the same reasons discussed above. Dependent claims 19-29 and 31-32 also recite the same limitations as limitations found in dependent claims 1-14 and 16-17 and therefore are also rendered obvious by Longo, McNab, and Zhang (with or without Maltego) for the same reasons discussed above in Grounds 1 through 3.

I. Claims 33-46, 48-49

195. As discussed above, Claim 33 is rendered obvious because of the reasons discussed in Grounds 1 or 2. This ground provides an alternative read for claim 13, which is also recited in limitation 33[b]. Thus, Claim 33 is also rendered obvious by Longo, McNab, and Zhang (with or without Maltego) for the same reasons discussed above. Dependent claims 34-46 and 48-49 also recite the same limitations as limitations found in dependent claims 1-12, 14, and 16-17 and therefore are also

rendered obvious by Longo, McNab, and Zhang (with or without Maltego) for the same reasons discussed above in Grounds 1 through 3.

J. Claims 55-58, 63, 66

196. As discussed above, Claim 50 is rendered obvious because of the reasons discussed in Grounds 1 or 2. Dependent claims 55-58, 63, and 66 also recite the same limitations as limitations found in dependent claims 5-8, 13 and 17 and therefore are also rendered obvious by Longo, McNab, and Zhang (with or without Maltego) for the same reasons discussed above.

K. Claims 84-98, 100

197. As discussed above in Section VII, independent Claim 84 recites the same limitations as limitations 1[a], 1[c], and claim 17. Thus, it is rendered obvious by Longo, McNab and Zhang (with or without Maltego) for the same reasons discussed above. Dependent claims 85-98 and 100 also recite the same limitations as limitations found in dependent claims 1-14 and 16-17 and therefore are also rendered obvious by Longo, McNab, and Zhang (with or without Maltego) for the same reasons discussed above in Grounds 1 through 3.

XI. GROUND 4: LONGO, MCNAB, AND TIPPETT (WITH OR WITHOUT MALTEGO) RENDER OBVIOUS CLAIMS 15, 47, 64, 67-71, 76-83, AND 99

A. Rationale to Combine

198. A POSITA would have been motivated to combine Tippet's teachings with Longo and McNab (and optionally, Maltego). Tippet 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.” Tippet, ¶¶29-30. Tippet 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.

199. 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. 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.” Tippet, ¶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.” Zhang, ¶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.

200. 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. 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. A POSITA would find this beneficial as it increases the security of the organization's network.

201. 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. 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. This would have been well within the capabilities of an ordinary artisan. 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.

202. Finally, Tippet is analogous art to the '615 patent because it is directed to associating network resources with an entity. Tippet, ¶32 ("The entity 224A may be associated with various information resources").

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.

203. Tippetts discloses the calculation of a “local security index” for one or more organizations, *i.e. scoring the security state of the entities*. Specifically Tippetts 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.” Zhang, ¶¶29-30. Tippetts 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.

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

C. Claims 47, 64, 99

205. Claims 47, 64, and 99 recite the same limitation as claim 15. They are therefore also rendered obvious by Longo, McNab, and Tippetts (with or without Maltego) for the same reasons discussed above.

D. Claims 67-71, 76-83

206. As I discussed above in Section VII, independent Claim 67 recites the same limitations as limitations 1[a], 1[c] and claim 15. Thus, it is rendered obvious by

Longo, McNab and Tippet (with or without Maltego) for the same reasons discussed above for Grounds 1 and 2 and this Ground. Dependent claims 68-71 and 76-83 also recite the same limitations as limitations found in dependent claims 2-4, 9-14, and 16-17 and therefore are also rendered obvious by Longo, McNab, and Tippet (with or without Maltego) for the same reasons discussed above for Grounds 1, 2, and 4.

XII. GROUND 5: LONGO, MCNAB, ZHANG, AND TIPPETT (WITH OR WITHOUT MALTEGO) RENDER OBVIOUS CLAIMS 30, 47, 72-75, 80, 83 AND 99

A. Rationale to Combine

207. As I discussed above in Ground 3, it would have been obvious to combine Zhang with Longo and McNab (with or without Maltego). And as I discussed above in Ground 4, it also would have been obvious to combine Tippet 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. 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.

B. Claim 30

208. As I discussed above, Claim 18 is rendered obvious because of the reasons discussed in Ground 3 (Longo, McNab and Zhang (with or without Maltego)). Dependent claim 30 recites the same limitation as found in dependent claim 15 (which is disclosed by Tippet as discussed in Ground 4) and therefore is also rendered obvious by Longo, McNab, Zhang, and Tippet (with or without Maltego) for the same reasons discussed above in Grounds 3 and 4.

C. Claim 47

209. As I discussed above, Claim 33 is rendered obvious because of the reasons discussed in Ground 3 (Longo, McNab and Zhang (with or without Maltego)). Dependent claim 47 recites the same limitation as found in dependent claim 15 (which is disclosed by Tippet as discussed in Ground 4) and therefore is also rendered obvious by Longo, McNab, Zhang, and Tippet (with or without Maltego) for the same reasons discussed above in Grounds 3 and 4.

D. Claims 72-75, 80, 83

210. As I discussed above, Claim 67 is rendered obvious because of the reasons discussed in Ground 4 (Longo, McNab and Tippet (with or without Maltego)). Dependent claims 72-75, 80, and 83 recite the same limitations as found in dependent claims 5-8, 13, and 17 (which are disclosed by Zhang and McNab as discussed in Ground 3) and therefore are rendered obvious by Longo, McNab, Zhang, and Tippet (with or without Maltego) for the same reasons discussed above in Grounds 3 and 4.

E. Claim 99

211. As discussed above, Claim 84 is rendered obvious because of the reasons discussed in Grounds 2 and 3 (Longo, McNab and Zhang (with or without Maltego)). Dependent claim 99 recites the same limitation as found in dependent claim 15 (which is disclosed by Tippet as discussed in Ground 4) and therefore is also rendered obvious by Longo, McNab, Zhang, and Tippet (with or without Maltego) for the same reasons discussed above in Grounds 2 through 4.

XIII. GROUND 6: SHULL RENDERS OBVIOUS CLAIMS 1-4, 9-16, 33-37, 42-48, 50-54, 59-65, 67-71, AND 76-82

A. 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,

212. With respect to the limitation “*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. 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. Shull, ¶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. Ex. 1001, 3:53-56.

213. With respect to the limitation “*technical assets that contribute to security characteristics of respective entities*,” as I 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*. 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.” Shull, ¶52.

214. With respect to the limitation “*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. Shull, ¶39.

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

215. Shull discloses that it generates the map using “data gathered through a variety of harvesting/investigation information.” Shull, ¶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.

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,

216. Shull describes that data about IP addresses can be “correlated” using other means. Shull, ¶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 as an entity).

217. 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.” Shull, ¶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. 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. This allows administrators to increase the accuracy of the mapping by providing human input where needed. A POSITA would have a reasonable expectation of success since such a user interface is well-within the capabilities of a POSITA.

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

218. Shull discloses mapping *IP addresses* with entities. IP addresses are *network-related information*. Ex. 1001, 2:1-3.

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

219. 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. Shull, ¶40; McNab, 23-24.

B. Claim 2: The method of claim 1 in which the technical assets comprise network-related information.

220. This limitation is identical to limitation 1[d], so is disclosed by Shull for the reasons I discussed above.

C. Claim 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.

221. As I discussed above for limitation 1[d], Shull's technical assets include *IP addresses*. Shull, ¶48.

D. Claim 4: The method of claim 2 in which generating the map comprises online discovery of information about the technical assets.

222. This limitation is identical to limitation 1[b], so is rendered obvious by Shull for the reasons I discussed above.

E. 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, industry classifications, and employee counts.

223. Shull discloses that the user interface may include "surrounding information" which "related to the matching information," *i.e. non-technical-asset*

information. Shull, ¶81. A POSITA would understand it includes information about the entities, such as a “corporate records” or “property records,” which includes *industry classifications*. Shull, ¶40. 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.” Shull, ¶73.

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

224. As I 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*. Shull, ¶81.

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

225. Shull discloses that one information source is a “UDRP database” which “contain[s] data related to UDRP complaints filed against cybersquatters and others.” Shull, ¶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. 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*.

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

226. As I discussed in limitation 1[c], Shull provides a mechanism for an administrator to *review and approve* “matching information” being “associated with the data element.” Shull, ¶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*.

I. 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.

227. Shull discloses that the associations between IP addresses and entities can be used to provide a “reputation score” of the entity. Shull, ¶¶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.*

J. Claim 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.

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

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

229. As I discussed above in Section VII, 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.

XIV. GROUND 7: SHULL, MCNAB, AND ZHANG RENDERS OBVIOUS CLAIMS 5-8, 13, 17, 18-32, 33-49, 55-58, 63, 66, 72-75, 80, 83, 84-100

A. Rationale to Combine

230. 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.

231. As discussed above in Ground 3 and the Background, McNab teaches enumerating IP addresses of an entity's network by obtaining a list of authoritative name servers, a list of domain names for which that name server is authoritative, and a list of hostnames for those received domain names, using well-known DNS tools. McNab, 30-36.

232. A POSITA would be motivated to perform these operations with Shull. 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 Shull's system. 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.

233. 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. 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 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.

234. 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).” Zhang, ¶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

(“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”).

235. 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. 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.” Zhang, ¶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.

236. 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. 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. 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.

237. Zhang also discloses the ability to generate various graphs of relationships among entities and their domains. Zhang, ¶84. 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.*

238. A POSITA would have had a reasonable expectation of success in combining Shull with McNab and Zhang. 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. 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.*

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

B. Claims 5-8, 13, 17

240. As I discussed above 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.

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

241. As I discussed above in Section VII, 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.

XV. CONCLUSION

242. I hereby declare that all statements made of my own knowledge are true and that all statements made on information and belief are believed to be true. I further declare that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of the Title 18 of the United States Code and that such willful false statements may jeopardize the validity of this proceeding.

Executed on this September 6, 2024 by:



Kevin C. Almeroth, Ph. D.