

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 IPR2024-01393
Patent 9,973,524

**DECLARATION OF KEVIN C. ALMEROOTH, PH.D. IN SUPPORT OF
PETITION FOR INTER PARTES REVIEW OF
UNITED STATES PATENT NO. 9,973,524**

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 '524 PATENT.....	21
	A. Technical Background	21
	B. Description of the '524 Patent's Specification	34
	C. The '524 patent's Prosecution History	34
	D. Interpretation of the '524 Patent's Challenged Claims	35
VI.	PUBLIC AVAILABILITY OF CERTAIN EXHIBITS	36
	A. McNab.....	36
	B. Internet Archive Pages.....	37
	C. Public Videos	38
VII.	GROUND 1: McNab and McGovern RENDER OBVIOUS CLAIMs 1-8, 10- 16, and 18-20.....	40
	A. Rationale to Combine	40
	1. McGovern	40
	2. Nesus	45
	B. Claim 1	46
	1. 1[pre], 1[a]: A computer implemented method comprising: on one or more computers, maintaining in a database an entity map that maps technical assets to respective companies or other entities with which the assets are associated, wherein mapping technical assets to respective companies or other entities comprises determining IP	

	information by annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses, ...	46
2.	1[b]: automatically usinging [<i>sic</i>] sensors on the Internet to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities and associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name, ...	55
3.	1[c]: automatically deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (i) a number of technical assets that have been reported to be malicious and (ii) a duration of detected malicious activity associated with the technical assets, ...	63
4.	1[d]: automatically generating a cyber-security rating for each of the entities using the entity map and the derived observations, and	71
5.	1[e]: through a communication network, exposing to users in a user interface of a portal, information about the cyber-security ratings of the entities that is useful to decisions about entities with which to do business.	73
C.	Claim 2: The method of claim 1 in which the technical assets comprise an IP space.	74
D.	Claim 3: The method of claim 1 in which maintaining an entity map comprises querying a regional Internet registry.	75
E.	Claim 4: The method of claim 3 in which maintaining an entity map comprises submitting a name of one of the entities to the regional Internet registry.	75
F.	Claim 5: The method of claim 4 in which submitting the entity name includes submitting spelling variations of the entity name.	76
G.	Claim 6: The method of claim 1 in which maintaining an entity map comprises using a domain name associated with the entity.	77
H.	Claim 7: The method of claim 6 in which maintaining an entity map comprises using a dig tool on the domain name to determine a published IP address.	78
I.	Claim 8: The method of claim 1 in which maintaining an entity map comprises probing IP addresses around a published IP address.	79

J.	Claim 10: The method of claim 1 in which maintaining an entity map comprises using a name server lookup to verify the derived server name.	81
K.	Claim 11: The method of claim 2 comprising using the IP space for collecting the cyber-security characterizations.	82
L.	Claim 12: The method of claim 1 in which the technical assets that are associated with cyber-security characterizations of the respective entities comprise technical assets that the entity owns, controls, uses, or is affiliated with.	83
M.	Claim 13: The method of claim 1 in which the technical assets comprise ranges of IP addresses.	84
N.	Claim 14: The method of claim 1 in which the technical assets comprise servers.	84
O.	Claim 15: The method of claim 1 in which maintaining an entity map comprises using information received from Internet Service Providers.....	84
P.	Claim 16: The method of claim 1 in which maintaining an entity map comprises using heuristics.	85
Q.	Claim 18: The method of claim 1, wherein annotating comprises prepending to the domain name at least one of “mail,” “smtp,” and “ftp”.....	85
R.	Claim 19.....	86
1.	19[pre], 19[a]: A computer implemented method comprising: on one or more computers operating with a database, maintaining an entity map that maps technical assets that comprise an IP space to respective entities with which the assets are associated, ...	86
2.	19[b]: the maintaining of the entity map comprising (a) automatically using sensors on the Internet to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities, (b) associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name, (c) submitting names of the entities to regional Internet registries using various spellings of the names, (d) using domain names associated with the entities, (e) probing IP addresses around published IP addresses, (f) using a name server lookup to verify IP addresses associated with domain names, (g) using information received from Internet Service Providers, (h) annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses, and (i) applying heuristics,.....	86

3.	19[c]: automatically deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (a) a number of technical assets that have been reported to be malicious and (b) a duration of detected malicious activity associated with the technical assets,.....	88
4.	19[d]: automatically generating a cyber-security rating for each of the entities using the entity map and the derived observations, and.....	89
5.	19[e]: through a communication network, exposing to users in a user interface the cyber-security ratings of the entities.	89
S.	Claim 20.....	89
1.	20[pre], 20[a], 20[b]: A system comprising: one or more computers programmed to facilitate operations comprising: maintaining in a database an entity map that maps technical assets to respective companies or other entities with which the assets are associated, wherein mapping technical assets to respective companies or other entities comprises determining IP information by annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses,.....	89
2.	20[c]: collecting, using sensors on the Internet, externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities and associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name,.....	89
3.	20[d]: deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (i) a number of technical assets that have been reported to be malicious and (ii) a duration of detected malicious activity associated with the technical assets,.....	90
4.	20[e]: generating a cyber-security rating for each of the entities using the entity map and the derived observations, and	90
5.	20[f]: through a communication network, exposing to users in a user interface of a portal, information about the cyber-security ratings of the entities that is useful to decisions about entities with which to do business.	90

VIII. Ground 2: Cole in view of McNab and McGovern Render Obvious Claims 1-8, 10-16, 18-20..... 90

A.	Rationale to Combine	91
----	----------------------------	----

1.	McNab.....	92
2.	McGovern	96
B.	Claim 1	99
1.	1[pre], 1[a] A computer implemented method comprising: on one or more computers, maintaining in a database an entity map that maps technical assets to respective companies or other entities with which the assets are associated, wherein mapping technical assets to respective companies or other entities comprises determining IP information by annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses	99
2.	1[b] automatically using [sic] sensors on the Internet to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities and associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name	101
3.	1[c] automatically deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (i) a number of technical assets that have been reported to be malicious and (ii) a duration of detected malicious activity associated with the technical assets.....	105
4.	1[d] automatically generating a cyber-security rating for each of the entities using the entity map and the derived observations.....	108
5.	1[e] through a communication network, exposing to users in a user interface of a portal, information about the cyber-security ratings of the entities that is useful to decisions about entities with which to do business.....	109
C.	Claim 2: The method of claim 1 in which the technical assets comprise an IP space.	113
D.	Claim 3: The method of claim 1 in which maintaining an entity map comprises querying a regional Internet registry.....	114
E.	Claim 4: The method of claim 3 in which maintaining an entity map comprises submitting a name of one of the entities to the regional Internet registry.	114
F.	Claim 5: The method of claim 4 in which submitting the entity name includes submitting spelling variations of the entity name.....	114

G.	Claim 6: The method of claim 1 in which maintaining an entity map comprises using a domain name associated with the entity.....	114
H.	Claim 7: The method of claim 6 in which maintaining an entity map comprises using a dig tool on the domain name to determine a published IP address.....	114
I.	Claim 8: The method of claim 1 in which maintaining an entity map comprises probing IP addresses around a published IP address.	115
J.	Claim 10: The method of claim 1 in which maintaining an entity map comprises using a name server lookup to verify the derived server name.	115
K.	Claim 11: The method of claim 2 comprising using the IP space for collecting the cyber-security characterizations.	115
L.	Claim 12 The method of claim 1 in which the technical assets that are associated with cyber-security characterizations of the respective entities comprise technical assets that the entity owns, controls, uses, or is affiliated with	115
M.	Claim 13: The method of claim 1 in which the technical assets comprise ranges of IP addresses.....	116
N.	Claim 14: The method of claim 1 in which the technical assets comprise servers	116
O.	Claim 15: The method of claim 1 in which maintaining an entity map comprises using information received from Internet Service Providers.....	116
P.	Claim 16: The method of claim 1 in which maintaining an entity map comprises using heuristics	117
Q.	Claim 18: The method of claim 1, wherein annotating comprises prepending to the domain name at least one of ‘mail,’ ‘smtp,’ and ‘ftp’	117
R.	Claim 19.....	117
1.	19[pre], 19[a]: A computer implemented method comprising: on one or more computers operating with a database, maintaining an entity map that maps technical assets that comprise an IP space to respective entities with which the assets are associated, ...	117
2.	19[b]: the maintaining of the entity map comprising (a) automatically using sensors on the Internet to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities, (b) associating the observable cyber-security characterizations with a respective entity	

	based at least in part on the derived server name, (c) submitting names of the entities to regional Internet registries using various spellings of the names, (d) using domain names associated with the entities, (e) probing IP addresses around published IP addresses, (f) using a name server lookup to verify IP addresses associated with domain names, (g) using information received from Internet Service Providers, (h) annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses, and (i) applying heuristics,.....	118
3.	19[c]: automatically deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (a) a number of technical assets that have been reported to be malicious and (b) a duration of detected malicious activity associated with the technical assets,.....	120
4.	19[d]: automatically generating a cyber-security rating for each of the entities using the entity map and the derived observations, and.....	120
5.	19[e]: through a communication network, exposing to users in a user interface the cyber-security ratings of the entities.	120
S.	Claim 20.....	120
1.	20[pre], 20[a], 20[b]: A system comprising: one or more computers programmed to facilitate operations comprising: maintaining in a database an entity map that maps technical assets to respective companies or other entities with which the assets are associated, wherein mapping technical assets to respective companies or other entities comprises determining IP information by annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses,.....	120
2.	20[c]: collecting, using sensors on the Internet, externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities and associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name,.....	121
3.	20[d]: deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (i) a number of technical assets that have been reported to be malicious and (ii) a duration of detected malicious activity associated with the technical assets,.....	121

4. 20[e]: generating a cyber-security rating for each of the entities using the entity map and the derived observations, and 121

5. 20[f]: through a communication network, exposing to users in a user interface of a portal, information about the cyber-security ratings of the entities that is useful to decisions about entities with which to do business. 122

IX. CONCLUSION 122

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. 9,973,524 (“the ’524 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 1-8, 10-16, and 18-20 (collectively, the “Challenged Claims”) of the ’524 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 ’524 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.

8. 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.

9. 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.

10. 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.

11. 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.

12. 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.

13. 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.

14. 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.

15. 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.

16. 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).

17. 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.

18. 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.

19. 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.

20. 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.

21. 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.

22. 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.

23. 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.”

24. 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.

25. 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.

26. 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).

27. 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.

28. 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.

29. 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.

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

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

32. 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

33. 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.

34. 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

35. 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.

36. 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.

37. 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.

38. 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.

39. 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.

40. 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.

41. 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.

42. 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.

43. 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.

44. 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.

45. 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

46. 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 '524 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.

47. 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

48. I have been asked to assume that the '524 patent is entitled to its earliest alleged priority date of September 24, 2010. *See* Ex. 1001. Beyond this

assumption, I have not undertaken an analysis to determine the earliest priority date to which the '524 patent is entitled.

49. Based on the materials and information I have reviewed and based on my experience in the technical areas relevant to the '524 patent, a person of ordinary skill in the art at the time of the alleged invention of the '524 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 '524 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 '524 patent.

V. THE '524 PATENT

A. Technical Background

50. Below, I provide a brief background summary of the technology at issue.

51. U.S. Pat. No. 9,973,524 (the "'524 patent") is directed to collecting information about network resources, generally from external information sources.

Prior to 2010, 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.

52. For example, there were a number of well-known OSINT tools for collecting IP address information about assets. The McNab reference (Ex. 1005) devotes an entire chapter to such methods for “Internet Host and Network Enumeration,” *i.e.* enumerating the various hosts within a network, including their IP addresses.

53. 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
```

54. McNab also describes a well-known, free tool called “dig” which can perform a number of functions related to DNS queries. 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.
```

55. Once the network hosts have been enumerated, they can then be scanned to assess information about the network. McNab itself explains that “Network

scanners are used to perform bulk automated scanning of IP ranges to identify vulnerable network service components. The two most popular open source network scanners are Nmap and Nessus.” McNab at 13. Nessus, for example, “is a vulnerability assessment package that can perform many automated tests against a target network, including ICMP, TCP, and UDP scanning, testing of specific network services (such as Apache, MySQL, Oracle, Microsoft IIS, and many others), and rich reporting of vulnerabilities identified.” *Id.* By identifying vulnerabilities, the network can be given a risk assessment score.

56. This process of enumerating network assets and performing vulnerability scans to generate risk assessment scores was described by other sources as well. For example, a 2001 paper titled “An Introduction to Internet Attack & Penetration” (Ex. 1015) describes an “Anatomy of a Hack” that lists the following steps:

1. Identification of Internet-based points of presence
2. Network scanning & reconnaissance
3. Accessing the target system or network
4. Carrying out objectives
5. Manipulation of logs and system files

57. The paper then goes on to describe attacks such as the exploitation of vulnerabilities, or various forms of malware. Ex. 1015 at 9-14.

58. In 2002, the National Institute of Standards and Technology (“NIST”) (an organization part of the U.S. Department of Commerce) published a paper called

“Risk Management Guide for Information Technology Systems” (Ex. 1016), which provides guidelines and methodologies for conducting risk assessments, mitigating risks, and ensuring the security of IT systems within the federal government. Section 3.1.2 of the paper outlines various information-gathering techniques, such as “Automated Scanning Tool[s]” such as “network mapping tool [to] identify the services that run on a large group of hosts and provide a quick way of building individual profiles of the target IT system(s).” Ex. 1016 at 11-12. This results in a “picture of the IT system environment, and delineation of system boundary.” *Id.* The paper then goes on in Chapter 3.3 to discuss scanning these services to identify vulnerabilities, noting that “[t]he analysis of the threat to an IT system must include an analysis of the vulnerabilities associated with the system environment.” *Id.* at 15. Based on this analysis, the paper recommends producing a risk rating, which is derived “by multiplying the ratings assigned for threat likelihood (e.g., probability) and threat impact,” resulting in various “risk levels of High, Medium, and Low” within a risk-level matrix. *Id.* at 24.

Table 3-6. Risk-Level Matrix

Threat Likelihood	Impact		
			<i>High</i> (100)
<i>High</i> (1.0)	Low 10 X 1.0 = 10	Medium 50 X 1.0 = 50	High 100 X 1.0 = 100
<i>Medium</i> (0.5)	Low 10 X 0.5 = 5	Medium 50 X 0.5 = 25	Medium 100 X 0.5 = 50
<i>Low</i> (0.1)	Low 10 X 0.1 = 1	Low 50 X 0.1 = 5	Low 100 X 0.1 = 10

*Risk Scale: High (>50 to 100); Medium (>10 to 50); Low (1 to 10)*⁸

59. 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 on to provide resources for discovering employees of an organization, such as through Google searches, or using a tool called Evolution that “is able to cross-reference information from a large set of public data sources.” *Id.* at 8. The chapter then explains 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.*


60. The paper also describes using “the Metasploit Framework” which “includes a number of application discovery modules ... to detect services that are difficult to find otherwise.” *Id.* at 14. For example, the below output searched on an IP range and returned DNS and SNMP services running on hosts within that IP range:

```
[*] Sending 6 probes to xxx.xxx.xxx.0->xxx.xxx.xxx.255 (256 hosts)
[*] Discovered DNS on xxx.xxx.xxx.19 (TinyDNS)
[*] Discovered DNS on xxx.xxx.xxx.249 (BIND 8.4.6-REL-NOESW)
[*] Discovered DNS on xxx.xxx.xxx.250 (Microsoft)
[*] Discovered SNMP on xxx.xxx.xxx.170 (Ethernet Routing Switch)
[*] Discovered SNMP on xxx.xxx.xxx.171 (ProCurve J8692A)
```

Id. The paper goes on to describe various ways to attack these services. *Id.* at 19-33.

61. In 2008, a well-known security researcher named Chris Gates delivered a presentation called “New School Information Gathering” (Ex. 1018). 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. The presentation then describes several tools for gathering such information. For example, the presentation describes one tool called Fierce DNS which “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 “baytsp.com” shows the various hosts for that domain name. *Id.* at 21-22.





Fierce DNS

```
216.133.204.246 dmz.baytsp.com
216.133.204.247 dmz.baytsp.com
216.133.204.248 dmz.baytsp.com
216.133.204.249 dmz.baytsp.com
216.133.204.250 dmz.baytsp.com
216.133.204.251 dmz.baytsp.com
216.133.204.252 dmz.baytsp.com
216.133.204.253 dmz.baytsp.com
216.133.204.254 dmz.baytsp.com
216.133.204.255 dmz.baytsp.com
216.133.204.240 demo.baytsp.com
216.133.204.229 dl.baytsp.com
216.133.204.227 mail.baytsp.com
216.133.221.24 pop.baytsp.com
216.132.49.228 vpn.baytsp.com
206.135.194.198 w3.baytsp.com
216.133.221.15 webmail.baytsp.com
68.178.254.204 www.baytsp.com
68.178.254.204 .baytsp.com
```

```
Subnets found (may want to probe here using nmap or unicornscan):
206.135.194.0-255 : 1 hostnames found.
216.132.49.0-255 : 1 hostnames found.
216.133.204.0-255 : 29 hostnames found.
216.133.221.0-255 : 2 hostnames found.
68.178.254.0-255 : 2 hostnames found.
```


```
Done with Fierce scan: http://hackers.org/fierce/
Found 35 entries.
```

Have a nice day.

CarnalOwne

Classified by SP 6068
Declassify on: OADR
CA# 94-1720 CRR b7C

62. Another tool called “SEAT (Search Engine Assessment Tool)” “uses information stored in search engine databases, cache repositories, and other public resources to scan a site *for potential vulnerabilities*.” *Id.* at 24 (emphasis added). The tool involves providing a domain name as input (such as “baytsp.com” as used in the slides) and returning a set of results about that domain. *Id.* at 26.



63. The presentation also describes methods for harvesting email addresses from the provided domain name. *Id.* at 28-32.

64. Another tool called Metagoofil searches for PDF, Word, and other documents found on the provided domain name and extracts metadata from those documents to, for example, “[r]eveal the creator of a document,” “[r]eveal the application that created the document,” “[r]eveal create date,” all in order to “deliver a directed client side attack for something installed in the enterprise.” *Id.* at 35-40.


65. Finally, the presentation describes a tool called Maltego, which can perform many of the above functions automatically in one piece of software. Maltego

“is a program that can be used to determine the relationships and real world links between: People ... Companies ... Web sites ... Internet infrastructure such as: Domains, DNS names, Netblocks, IP addresses” and it does this “using open source intelligent (OSINT).” *Id.* at 67. Maltego can enumerate network assets given a domain name, such as by finding “MX and NS servers” for the domain name and determining “what hostnames are shared by the DNS servers,” and “brute forc[ing] our domain names and resolve them to IPs.” *Id.* at 67-72. From there, email addresses and documents can be harvested. *Id.* at 73-77. The presentation explains that Maltego can also search blogs, social networks and incoming links. *Id.* at 82. Maltego also provides graphical representations of the mapped network assets, showing them being associated with various entities. *Id.* at 80.



Maltego

- Sort your views to see relationships



Carnal O w n a g e

11/10/94
Classified by 57668 [redacted] 167C
Declassify on: OADR
CAF 94-1720 CRR

66. Maltego is further described and shown at Exhibit 1021, which is a 2009 video entitled “Social Engineering Using Maltego.” The video shows how Maltego can be used to discover network assets for a particular domain, as well as additional domains that share a DNS name server.

67. Thus, as shown above, ordinary artisans had long had experience seeking information from various sources to enumerate network assets and collect information for providing a risk assessment score.

B. Description of the '524 Patent's Specification

68. The '524 patent to Boyer *et al.* was filed November 18, 2015 and claims priority to a provisional application filed September 24, 2010.

69. The '524 patent is directed to conventional techniques for measuring the security vulnerability of a network from external observances. Specifically, it can “creat[e] a composite security rating” from data that is “derived from externally observable characteristics.” Ex. 1001 at Abstract. The claims of the '524 patent further require displaying such ratings on a map of the assets within the network. *See, e.g., id.*, cl. 1.

C. The '524 patent's Prosecution History

70. I have reviewed the prosecution history of the '524 patent and summarize it below.

71. During prosecution, the applicant made a number of amendments that failed to overcome the examiner's rejections. Finally, following an interview with the examiner, applicant added the following limitations: (1) that the entity map is based on determining IP information specifically for “a mail server, an SMTP server, [or] an FTP server,” (2) and the step of “associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name,” and (3) that the characterizations comprise “a number of technical assets” or “duration of detected malicious activity.” Applicant noted that while one of the

references, Wanser, “appears to describe domain name naming conventions in general,” it did not describe it specifically for “a mail server, an SMTP server [or] an FTP server,” nor disclose attributing events to a particular entity.

72. The examiner initiated an interview and proposed an amendment where the characterizations have to comprise both “a number of technical assets” and “duration of detected malicious activity.” Following these amendments, the examiner allowed the claims.

D. Interpretation of the '524 Patent's Challenged Claims

73. 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.

74. 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 24, 2010 filing date of the '524 patent as the date of the invention.

VI. PUBLIC AVAILABILITY OF CERTAIN EXHIBITS

75. I have been asked to opine on the public availability of certain exhibits and specifically whether they were publicly available by September 24, 2010.

A. McNab

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

77. 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.

78. 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.

79. 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

80. Several of the exhibits I have reviewed are from the Internet Archive, including Exhibits 1008-1012, 1015, 1017, 1020, and 1022. I am familiar with the Internet Archive and have reviewed the declaration of Mr. Frank-White regarding the operation of the Internet Archive.

81. 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

82. Exhibits 1019 and 1020 are printouts of webpages that host a video that was publicly available on Vimeo and YouTube, respectively.

83. 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.

84. 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.

85. 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.

86. 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.

87. 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.

88. This 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. GROUND 1: MCNAB AND MCGOVERN RENDER OBVIOUS CLAIMS 1-8, 10-16, AND 18-20

89. McNab is an O'Reilly published book called "Network Security Assessment" and provides various methods for assessing networks. In particular, it discloses the creation of a map of assets for an organization's networks/entities, and furthermore discloses a number of methods for assessing the security of those network assets, including the performance of vulnerability scans from outside the networks.

90. McGovern discloses determining the duration of malicious activity, and in my opinion, it would have been obvious to combine McNab and McGovern.

A. Rationale to Combine

1. McGovern

91. As I explain below, a person of ordinary skill would have found it obvious to combine McNab's network enumeration and security scanning methods with McGovern's teachings of determining the duration of a malicious activity as well as its teachings of quantifying cyber-security observations into a numerical score to be displayed in a web-based user interface.

92. In my opinion, an ordinary artisan would be motivated to incorporate such teachings because they were known to increase the security of a network as well as increase user-friendliness of the system and provide the user an easier assessment of the network. Simplifying the display of security information would be of paramount

interest to a POSITA as it would aid in identifying and addressing security breaches and other security issues within the user's network.

93. McGovern discloses the capability of determining the “average response or patch time after security breach detection.” McGovern at ¶38. A POSITA would be motivated to incorporate this teaching into McNab, which discloses various methods and measurements for evaluating the security of a network, because of the known benefits. A POSITA would understand that it would be beneficial to include the duration of a security breach in the network security assessments of McNab, because it was known that the duration of a security breach is an important factor in determining overall network security. A security breach is an incredibly dangerous and highly undesirable event that puts all of the organization's data and resources at risk of damage or theft. Security breaches can (and usually do) cause damage to an organization in various ways, such as the theft of confidential information, putting the business at a competitive disadvantage, or the theft of customer or employee information, therefore harming said customers or employees, harming the company's reputation, and increasing the liability risk to the company. A POSITA would therefore understand that it would be highly desirable to minimize the duration of any security breach—thus, it would be beneficial for a network administrator to know whether the network and the organization's IT policies are robust and proactive to such security

breaches (resulting in a minimal duration) or if they are particularly vulnerable (as evidenced by a lengthy duration).

94. McGovern also discloses the capability of assigning “a numerical score” to “each risk level or degrees of compliance (associated with each security parameter)” following a network assessment. McGovern at ¶37. McGovern further teaches that this numerical score is further used by “applying one or more scoring algorithms to the data structure(s) to calculate information security assessment (ISA) scores based on input data concerning an information system.” *Id.* at ¶42. McGovern also teaches that the ISA score is “incorporated into a software program with a user interface,” which can be “designed to run as a Web-based service or application” wherein users can “access the scoring and related functionalities remotely via standard browsers or similar user interfaces.” McGovern at ¶¶42, 45.

95. A POSITA would also be motivated to incorporate these teachings into McNab because they provide an easier method of evaluating a network’s security assessment, improving usability and security. I note that McGovern expressly teaches that the numerical score will “quantify the contribution of each security parameter to ultimately reach an overall risk assessment,” indicating that the score is an important factor in providing an “overall risk assessment.” *Id.* at ¶37. McGovern also teaches that the numerical scores that partly make up the ISA score reflects “degrees of compliance with security standards” and “goodwill and negative publicity” of the

entity to which the network belongs. *Id.* These teachings would motivate a POSITA to use McGovern's ISA scores to provide useful evaluations of the network and the entity. Generating an ISA score and displaying it in a user interface increases the ease for a network administrator to identify risks within the network, since a numerical score takes significantly less effort and time to evaluate compared to, for example, a lengthy and wordy report. As a network's assets increase, a network administrator will increasingly need the capability of evaluating the security of the assets in a quick and efficient manner instead of being forced to read lengthy reports before knowing the risk assessment. A POSITA would also be further motivated by McGovern's teachings to allow the ISA score to be viewed in a remote user interface because doing so would allow the network administrators to access the ISA scores from anywhere, increasing ease-of-use and accessibility. Thus, a POSITA would be motivated to incorporate such teachings in order to increase security and usability.

96. In my opinion, a POSITA would have combined McNab and McGovern with a reasonable expectation of success. The majority of the McNab book is dedicated to explaining various methods for performing network intrusion tests. *See* McNab at Chapters 5-16. In my opinion, adding McGovern's capability of determining the duration of a security breach would simply be another test that an administrator could perform that was well-within the capability of a POSITA at the time of the '524 patent. McNab does not state any limit on the number of tests that

could be performed (and some of the software McNab describes performs dozens of tests, such as the Nessus vulnerability scanning I discuss below), and a POSITA would understand that adding an additional test would not pose any significant hurdle in computing power or evaluation time, especially in view of the benefits of said test described above.

97. A POSITA would have also expected success in combining McNab with McGovern's numerical score. McGovern provides an algorithm for calculating the numerical score that could be applied to any of the tests described by McNab, including the Nessus vulnerability scanning software. Indeed, McNab discloses that Nessus already "categorizes findings into three severity levels: security hole, security warning, and security info." McNab at 390. Quantifying such levels (such as assigning each level a particular number like 1, 2, and 3) would have been a trivial task well-within the capabilities of a POSITA.

98. In addition, a POSITA would understand that implementing McGovern's measurement of "response or patch time after security breach detection" in McNab's system would have been well-within the capabilities of a POSITA. In particular, as I discuss in fuller detail below, McNab describes a Nessus vulnerability scanner which can detect whether assets were infected by malicious viruses. McNab also discloses that Nessus supports "the ability ... to compare two reports for differences." McNab at 391. Thus, given these teachings, a POSITA would have

understood that McGovern’s “response or patch time” could be measured by evaluating reports by the Nessus vulnerability scanner until the asset is no longer infected and determining the duration it was infected.

99. Finally, in my opinion, McNab and McGovern are directed to the same field of endeavor as the ’524 patent, particularly the field of network security. McNab at xvi (“This book tackles one single area of information security in detail: that of undertaking IP-based network security assessment in a structured and logical way.”); McGovern at Abstract (“Methods for information security assessment and data risk scoring are disclosed.”); Ex. 1001 at Abstract (“A method and system for creating a composite security rating from security characterization data of a third party computer system.”).

2. Nessus

100. As I discuss below, McNab discloses in detail a third-party vulnerability scanner called Nessus. McNab discloses that documentation about Nessus can further be found on the website “<http://www.nessus.org/documentation>”. McNab at 379. McNab further describes that Nessus supports plug-ins which are provided on the website. *Id.* at 377. Thus, in my opinion, documentation and plug-in information about Nessus as described in Exhibits 1008-1012 are incorporated into McNab by reference.

101. I also believe it would have been obvious to combine Exhibits 1008-1012 with McNab. Exhibits 1008-1012 are archived webpages (and all but Exhibit 1011 are from the Nessus website) that provide additional detail about the Nessus product described in McNab. Thus, a POSITA would be motivated to include such additional functionality, such as the plug-ins which detect malicious software such as MSBlaster and So.Big (Exs. 1008-1009) or documentation that describes user interface functionality for Nessus (Exs. 1010-1011), because they describe improved security capabilities and displays. A POSITA would have a reasonable expectation of success because these documents describe the Nessus product and therefore do not require any modifications to the product itself.

B. Claim 1

1. **1[pre], 1[a]: A computer implemented method comprising: on one or more computers, maintaining in a database an entity map that maps technical assets to respective companies or other entities with which the assets are associated, wherein mapping technical assets to respective companies or other entities comprises determining IP information by annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses,**

102. In my opinion, McNab discloses this limitation.

103. McNab discloses several methods for performing external security tests on networks of entities, beginning by maintaining a database of an entity map that maps technical assets to the entities of a particular organization (or multiple

organizations). One such test is called “Forward DNS Grinding” which determines IP information by using mail server and SMTP server naming conventions.

104. With respect to the limitation “*entity map that maps technical assets to respective companies or other entities*,” The ’524 patent describes an “entity map” as “map[ping] data back to an entity,” such as mapping an IP address to an entity. Ex. 1001, 6:15-17.

105. McNab discloses generating an *entity map* by discovering and mapping technical assets for a respective entity. For example, McNab generally “tackles one single area of information security in detail: that of undertaking IP-based network security assessment in a structured and logical way.” McNab at Preface xvi. One method of security assessment is described in Chapter 3, called “Internet Host and Network Enumeration,” which explains how to “*map* your network, from open web searches to DNS sweeping and querying of authoritative name servers.” McNab at Preface xx. Indeed, McNab explains that such network mapping is one of “the first steps you should take when” assessing your network by “assuming the role of an Internet-based attacker.” *Id.* at 17.

106. McNab explains that an *entity map* can be generated by “querying open sources for information relating to the target organization and its networks.” *Id.* These include: “Web and newsgroup search engines,” “Domain and IP WHOIS registrars,” “Border Gateway Protocol (BGP) looking glass sites and route servers,”

“Public DNS servers,” “DNS querying and grinding against specific name servers,” “Web server crawling,” and “SMTP probing.” *Id.* By querying these sources, an assessor can “build a useful *map of your networks* and understand where potential weaknesses may lie.” *Id.*

107. McNab explains that these queries results in obtaining lists of systems of interest, *i.e.* *assets*. *Id.* For example, McNab discloses that “Internet Host and Network Enumeration” techniques discussed in Chapter 3 involve “obtain[ing] useful data about *the structure of the target* network” by gathering “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*), and details of relationships between physical locations.” *Id.* at 5; *see also id.* at 9 (“[T]he network enumeration process ... identif[ies] further network components.”). For example, by using “DNS querying,” one can “enumerate hostnames and sub-domains,” as well as “operating platforms of devices” and “internal network information” within a network. *Id.* at 40. One method McNab describes is “Forward DNS Grinding,” which is a “forward brute-force grinding attack ... to enumerate valid DNS address records and aliases relating to the domain and its hosts,” such as SMTP mail servers of a domain. *Id.* at 35-36. A POSITA would understand that network components, hostnames, and sub-domains such as SMTP mail servers are *assets associated with the entity*. As the ’524 patent explains, “IT assets” include “IP

addresses controlled by the entity and the associated activity of those IP addresses.”

Ex. 1001, 9:65-10:1. A POSITA would understand by performing DNS querying, it obtains the IP addresses of various systems controlled by the entity, and the types of those systems such as SMTP and/or mail servers.

108. With respect to the limitation “*companies or other entities with which the assets are associated*,” McNab discloses that its techniques for building an *entity map* that maps assets to an organization is done with respect to the assets of multiple organizations.


109. McNab discloses that its techniques involve gathering “information relating to the target organization and its networks.” McNab at 17. However, McNab does not limit its techniques to a single organization and a POSITA would understand they could be used on any number of organizations. For example, McNab provides examples of mapping technical assets for different organizations: Nintendo of America (McNab at 24-25), the CIA (*id.* at 34-35), and the Bank of England (*id.* at 35-36). McNab also discloses that its techniques can be used on “companies and organizations.” *Id.* at 5. A POSITA would understand that any of the techniques disclosed by McNab could be used with any domain name associated with a company or organization to generate an entity map that maps assets for that company or organization. For example, Forward DNS Grinding using the “TXDNS” tool will “enumerate valid DNS address records and aliases relating to the domain and its hosts,”

including *any* domain. *Id.* at 35-36. McNab provides an example of using TXDNS on the “bankofengland.co.uk” domain to identify its mail servers by executing the command “txdns -f mail-dict.txt bankofengland.co.uk.” *Id.* A POSITA would understand that such a command could be used with multiple organizations and their domain names to map assets to each organization.

110. McNab also describes that additional assets are enumerated (and thus additional mapping are generated) as “New domain names and IP addresses” are discovered. McNab at 8 (showing that “Network Enumeration” is performed on “New domain names and IP addresses”). Consistent with the specification of the ’524 patent, a POSITA would understand that the plain meaning of “entity” includes a computer system associated with a “domain name” (such as bankofengland.co.uk). For example, the ’524 patent describes that “[m]apping may also include *a domain name* associated with the third party computer system,” the third party computer system being an “entity.” Ex. 1001, 2:35-36, 5:57-58 (“an entity (e.g., a third party computer system)”); *see also id.* 9:4-9 (“It is also possible to query the [RIR] database by an *entity name* and get a list of IP addresses allocated to that *entity*.”). Thus, McNab discloses that as new domain names are discovered (*i.e.* new *entities*), it will discover and map assets to those new entities.

111. In addition, it was well-known that network asset enumeration techniques could be used to map assets to multiple entities. For example, as discussed

above in the Technical Background, numerous tools were available that generated mappings for multiple entities. Maltego is one such exemplary tool that would discover and map assets to new entities as they were discovered. Ex. 1018 at 67-80.



112. With respect to the limitation *“maintaining in a database an entity map,”* McNab discloses that the results of its mapping techniques (*i.e.* “Internet Host and Network Enumeration” techniques) are “gathered” and stored, because such network enumeration information is “used to perform structured bulk network scanning and probing exercises to further assess the target network space and investigate potential vulnerabilities. McNab at 5. McNab also describes a “Cyclic

Assessment Approach” where the mapping is updated as “New domain names and IP addresses” are discovered. *Id.* at 8 (“[A]n analyst may identify previously unknown IP address blocks, which can then be fed back into the network enumeration process to identify further network components.”). A POSITA would therefore understand that the entity map must be stored.

113. A POSITA would have considered it obvious to store the entity map in a database. A POSITA at the time of the patents in 2010 would have been aware of the benefits of storing data in a database. It was well-known that databases were designed to efficiently store and query data, and by 2010, it was common for enterprises to store all manner of data in databases—even the smallest organizations had databases to store information about IT assets such as network users. A POSITA would have known that database storage was fast and space-efficient as a storage mechanism, and been motivated by that. By 2010, there were numerous commercial and non-commercial databases available for use within an organization, including Oracle, Microsoft SQL Server, MySQL, and postgresdb. A POSITA would have also been familiar with the operations for storing and retrieving data in a database, and would have had a reasonable expectation of success in storing McNab’s entity map into a database. For example, it would have been well-within the capabilities of a POSITA to store list of entity assets, including their hostnames and IP addresses, in a

database table, as well as a list of entities to be mapped to such assets. A POSITA would have been familiar with the necessary database operations to store such data.


114. With respect to the limitation *“wherein mapping technical assets to respective companies or other entities comprises determining IP information by annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses,”* McNab discloses that generating an entity map includes a method called “Forward DNS Grinding” where a domain name is annotated with a “dictionary of common mail server names” to determine the IP address of a mail server or SMTP server. McNab at 35.

115. Specifically, McNab discloses that the “map of your networks” includes information from “DNS querying and grinding against specific name servers.” *Id.* at 17. “Forward DNS Grinding” refers to a “forward brute-force grinding attack ... to enumerate valid DNS address records and aliases relating to the domain and its hosts.” *Id.* “Forward DNS Grinding” involves appending conventional prefixes to a chosen domain name, and attempting DNS resolution for that name. If the asset exists, the DNS resolution will resolve its IP address. This is precisely what has been claimed in this limitation.

116. McNab discloses the use of a “very effective forward DNS grinding tool” known as TXDNS. McNab provides an example of a grinding attack against the

“bankofengland.co.uk domain, using a small dictionary of common mail server names.” *Id.* This results in the discovery of three “SMTP mail servers”: “mail.bankofengland.co.uk,” “mail2.bankofengland.co.uk,” and “mailhost.bankofengland” and their IP addresses. *Id.* at 36. 95 other attempted domain names failed.

117. SMTP probing involves “sending an email message to a nonexistent address at a *target domain* [i.e. *SMTP server*] often reveals useful internal network information.” *Id.* at 38-39; *see also* 40 (“SMTP probing: Sending email messages to nonexistent accounts at target domains to map *internal network space* by analyzing the responses from the SMTP system.”). The SMTP server has an SMTP naming convention, as shown below. McNab further discloses that SMTP probing uses a server name annotated with an IP address to derive a likely server name having assigned IP addresses.



118. Later in McNab, it provides another example of using TXDNS against the “apple.com” domain using a different dictionary file “smalllist.txt” that reveals the IP address of an FTP server “ftp.apple.com,” in addition to other servers. *Id.* at 85-86.

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

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

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

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

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

119. Thus, McNab discloses discovering *assets* by *annotating a domain name* (such as bankofengland.co.uk or apple.com) *with naming conventions for at least one of a mail server, an SMTP server, and an FTP server* (such as the mail or ftp prefix) to *derive a likely server name* which can be used to *determine its IP address*.

2. 1[b]: automatically using [sic] sensors on the Internet to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities and associating the observable

**cyber-security characterizations with a respective entity
based at least in part on the derived server name,**

120. McNab discloses a vulnerability scanner called Nessus that runs on the Internet and uses multiple sensors to collect externally observable cyber-security characterizations of each entity's technical assets.

121. With respect to the limitation *“automatically using,”* McNab discloses that external observations are automatically performed after entity mapping. Specifically, McNab discloses a “Network Security Assessment Methodology” where after the entity mapping described above (described as “Network reconnaissance to identify IP networks and hosts of interests” or “Internet Host and Network Enumeration”), the methodology performs “[b]ulk network scanning and probing to identify potentially vulnerable hosts” and “[i]nvestigation of vulnerabilities.” McNab at 4-5. A POSITA would understand that such steps are performed automatically.

122. One of the tools McNab discloses for vulnerability scanning is called Nessus. *See* McNab at 377-392. McNab describes Nessus as an *“automated system[]”* for “vulnerability scanning” that requires “minimal hands-on qualification and assessment of vulnerabilities.” *Id.* at 3. Nessus was “a free vulnerability scanner that can be used to perform a number of network-wide bulk security check.” *Id.* at 377. Although Nessus is described in Chapter 15 of McNab, a POSITA would have understood that its disclosures are part of the same “Network Security Assessment Methodology” described in Chapter 1, which also includes the entity mapping

described above. *See id.* at 4-5. In addition, McNab discloses that Nessus is part of the “Network Security Assessment Platform” and is one of the “components and tools that make up a professional security consultant’s toolkit for performing tasks including reconnaissance, network scanning, and exploitation of vulnerable software components.” *Id.* at 10, 12-13. In the alternative, a POSITA would have considered it obvious to combine the teachings in these chapters, not least because the McNab reference teaches that Nessus is one of the network assessment tools to be performed on assets discovered from entity mapping. *Id.* at 4-5, 10, 12-13. A POSITA would have also been motivated to use Nessus after entity mapping, since McNab instructs that “network scanning and probing to identify potentially vulnerable hosts” and “[i]nvestigation of vulnerabilities” should be performed after entity mapping, and McNab teaches that Nessus is an effective tool for performing such scans. McNab at 4-5, 377. A POSITA would have had a reasonable expectation of success in combining these chapters in McNab because they would simply involve using off-the-shelf tools in their known, conventional manner.

123. McNab discloses that Nessus performs a vulnerability scan *automatically* because it “significantly [reduces] the amount of time spent during a penetration test performing manual checks.” *Id.* McNab also discloses that Nessus can be “start[ed]” automatically “by using an operating system-appropriate startup file.” *Id.* at 382. Finally, McNab also discloses a “command-line client for scanning,”

and POSITA would have understood that a command line can be automated by using any number of well-known techniques, such as the “cron” command in UNIX or “Task Scheduler” in Windows. A POSITA would have considered it obvious to use the “cron” or “Task Scheduler” utility to automate performing scans using Nessus after network enumeration, as that would provide the well-known benefit of allowing for regularly scheduled vulnerability scans to ensure network assets are not vulnerable. Indeed, it was well-known that vulnerability scanning was often performed automatically after network enumeration, such as described in the NIST paper wherein “Automated Scanning Tool[s]” are used for “network mapping” and then performing vulnerability scanning. Ex. 1016, 11-12, 15. A POSITA would have had a reasonable expectation of success because “cron” and “Task Scheduler” were utilities a POSITA would have been familiar with.

124. With respect to the limitation “*sensors on the Internet to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities,*” McNab discloses that Nessus uses multiple sensors on a server on the Internet to collect externally observable cyber-security characterizations of target assets mapped to entities.

125. McNab discloses that Nessus is “a free *vulnerability scanner* that can be used to perform *a number of network-wide [i.e. multiple sensors]* bulk security check.” *Id.* at 377. The Nessus client “configures the various target, scanning, and

plug-in options, and it reports the findings from the scan to the user.” *Id.* These configurations and plug-ins are each a *sensor* that collects the *cyber-security characterizations* , such as those disclosed in the figures below. *Id.* at 388.




Figure 15-2. NessusClient 3 default scan policy options




Figure 15-3. NessusClient 3 plug-in selection window

126. McNab discloses that Nessus collects *externally observable* cybersecurity characterizations because it “can be run in two different modes: with or without authentication credentials” and “[w]hen run *without credentials*, Nessus will perform *remote network-based security checks*, testing how the target host responds to specific network probes.” *Id.* McNab further discloses that Nessus performs its

scans using a server on *the Internet*. Specifically, the “Nessus server performs all of the scanning and security checks” and is “Internet-connected.” *Id.* at 377, 383.

127. Some of the characterizations that can be performed by Nessus via plug-ins include detecting whether “[t]he remote host [is] infected by the MS Blaster worm” (Ex. 1008) or “check[ing] for the presence of different virii on the remote host” (Ex. 1009). These exhibits are webpages from Nessus’s website archived in 2006 describing certain plug-ins supported by Nessus.

128. Finally, McNab discloses that Nessus collects the *cyber-security characterizations* for the *technical assets that have been mapped to each of the entities*. Specifically, McNab discloses that Nessus performs its scans on “target network addresses or hosts” using “specific hostnames or IP addresses, ranges of network addresses, or through supplying a file with a target host on each line.” *Id.* at 385. As discussed above, McNab discloses (or renders obvious) that Nessus is part of the Network Security Assessment Methodology, where Nessus is performed on hosts identified through entity mapping, *i.e. technical assets that have been mapped to each of the entities*. *Id.* at 4 (“Bulk network scanning and probing to identify potentially vulnerable hosts”), 3 (“Vulnerability scanning uses automated systems (such as Nessus ...)”).

129. With respect to the limitation *“and associating the observable cyber-security characterizations with a respective entity based at least in part on the derived*

server name,” McNab discloses that the Nessus vulnerability scanner associates its cyber-security characterizations with a particular entity based on the derived server name. The term *derived server name* has an antecedent basis to the server name that was derived by “annotating a domain name with naming conventions for ... a mail server, an SMTP server [or] an FTP server,” as recited in limitation 1[a].

130. Specifically, McNab discloses that Nessus requires “target network addresses or hosts to be scanned,” which can be specified using “specific hostnames.” McNab at 385. As discussed above, McNab discloses (or renders obvious) that Nessus is part of the Network Security Assessment Methodology, where Nessus is performed on hosts identified through entity mapping, including the derived hostnames identified with “Forward DNS Grinding.” The Nessus scan is then performed on an entity based on the chosen hostname, resulting in the collection of *observable cyber-security characterizations* associated with that hostname belonging to an entity (e.g. the hostname mail1.bankofengland.co.uk belonging to the entity Bank of England). *Id.* at 389-90. McNab then generates a report, indicating that the characterizations were associated with the entity. *Id.* at 390-91.




Figure 15-6. NessusClient 3 completed scan

3. **1[c]: automatically deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (i) a number of technical assets that have been reported to be malicious and (ii) a duration of detected malicious activity associated with the technical assets,**

131. In my opinion, McNab in view of McGovern render obvious this limitation.

132. With respect to the limitation “*automatically deriving observations about the technical assets from the collected cyber-security characterizations,*” McNab discloses deriving observations from the cyber-security characterizations. These observations include technical assets that have been reported to be malicious. In addition, McGovern discloses deriving observations such as “average response or patch time” following a security breach detection, which indicates duration of detected malicious activity. As discussed above, it would have been obvious to combine McNab with McGovern.

133. McNab itself discloses that the Nessus vulnerability scan results in reports (*i.e. derived observations*) from the *cyber-security characterizations*, such as characterizations about whether a host is “infected by the MS Blaster worm” or some other virus. Exs. 1008-1009. From these characterizations, Nessus generates a report that provides information about those characterizations, such as the one displayed in the below figure. *Id.* at 390-91.




Figure 15-6. NessusClient 3 completed scan


134. While McNab describes that Nessus “categorizes findings into three security levels,” it does not provide an example of how Nessus depicts these security levels. However, webpages that were archived in 2006 and 2007 do provide such examples. For example, Exhibit 1010, which is the Nessus-published document “Nessus 3.0 Client Guide,” archived on March 12, 2007 from Nessus’s website, shows


an example of a report being viewed a web browser for a specific host (*i.e.* 170.20.20.156). Ex. 1010 at 8. As can be seen, the report provides icons for two of the “security levels,” *i.e.* **derived observations**. *Id.* One of them is a red “X” because the “epmap” scan returned a high risk factor, the other is an informational icon from the “general/udp” scan merely returning “information.” *Id.*



135. Similarly, Exhibit 1011 is an article on securityfocus.org that was archived in 2016 and linked from the Nessus website (*see* Ex. 1012), titled “Nessus, Part 3: Analysing Reports.” The article provides several examples of how the “three security levels” are depicted as icons. In Figure 1, each host is associated with an icon


indicating the “Severity” of the security risk, for example, a severity of “Security Warning” has an icon of a yellow triangle with an exclamation mark. In Figure 2, there is a red “X” icon for “High severity.”





136. With respect to the limitation “*wherein the derived observations comprise (i) a number of technical assets that have been reported to be malicious,*” As discussed above for limitation 1[b], Nessus provided plug-ins that could detect the presence of certain viruses on the target host. See Exs. 1008-1009. If detected, these plug-ins will report a “Risk factor” of “High” or “Critical.” *Id.* These viruses cause the infected host to be malicious. See, e.g., Ex. 1022 (“[T]he worm will cause compromised systems to initiate a denial of service attack against windowsupdate.com.”). Thus, these plug-ins provide *derived observations* indicating that a technical asset is *reported to be malicious*.

137. While Nessus does not display the *number of malicious assets*, since the Nessus client user interface provides readily-observable icons indicating high severity for each host that is reported to be malicious, the presence of those icons indicate the *number of malicious assets*. For example, the below screenshot indicates that 10 hosts (*i.e. assets*) have a risk level of “Security Warning.”



138. If those hosts were detected to have viruses (*i.e. reported to be malicious*), then a POSITA would understand they would have red “X” icons indicating that they are malicious. This is especially the case if the administrator decides to only run virus scans on the hosts (and not other scans that also may result in risk factors of high or critical).

139. In the alternative, it would have been obvious to derive the *number of assets reported to be malicious*. Since Nessus already reports which assets are malicious, a POSITA would have found it obvious to collect that information and derive the number of said malicious assets. A POSITA would be motivated to do this because deriving such a number would allow the number to be displayed to the administrator, thereby giving the administrator an easily-accessible number to indicate the health of the network. Deriving the number would also allow the system to alert the administrator if the number exceeds zero. A POSITA would have a reasonable expectation of success in doing this, because since Nessus discloses that its reports can be exported to XML or a database (McNab at 391), it would be well-within the skill of a POSITA to write a script that takes the XML or database report information and derives the number of assets reported to be malicious by checking the status of the virus detecting plug-ins for each host. The script could further email said number to the administrator.

140. With respect to the limitation “*and (ii) a duration of detected malicious activity associated with the technical assets,*” McGovern discloses identifying and collecting “security parameters,” *i.e. cyber-security characterizations*, and evaluating “risk levels or degrees of compliance may be established for each security parameter” (*deriving observations from the collected cyber-security*

characterizations). McGovern at ¶35. McGovern further discloses that the risk levels can be qualitatively, quantitatively, binary, or even have more than two values. *Id.*

141. McGovern discloses that one security parameter is the “average response or patch time after security breach detection,” which a POSITA would understand to be the *duration of detected malicious activity associated with the technical assets* because it indicates how long on average for the technical assets the malicious activity occurred before it was addressed. *Id.* at ¶38. An average response time includes one or more durations of measured response times.

142. As discussed above, a POSITA would have found it obvious to combine McNab with McGovern, and in particular, by using Nessus’s ability to generate reports as to whether an asset has been infected by malicious viruses and “compare two reports for differences.” McNab at 391.

4. 1[d]: automatically generating a cyber-security rating for each of the entities using the entity map and the derived observations, and

143. In my opinion, McNab in view of McGovern render obvious this limitation. McGovern discloses generating a score based on the observations, and it would have been obvious to combine McNab and McGovern.

144. McGovern discloses that “for each risk level or degrees of compliance (associated with each security parameter)” (*the derived observations*), “a numerical score may be assigned” where such numerical score is used to “quantify the

contribution of each security parameter to ultimately reach an overall risk assessment” (*cyber-security ratings*). McGovern at ¶37. Specifically, “a higher score reflects a greater risk exposure” and is “determined by the evaluation of underlying assets,” such as “degrees of compliance with security standards.” *Id.* This numerical score along with “the security parameters” and “risk levels” can further be “organized into one or more data structures” from which “to calculate information security assessment (ISA) scores ... concerning *an information system*,” *i.e.* an *entity*. *Id.* at ¶42.

145. Because McGovern discloses that a security assessment score is generated for the entire entity based on the security assessments of the assets, a POSITA would understand that, in view of McNab’s disclosure of entity maps that map assets to different entities, each entity would be given an numerical score based on the assets mapped to it in the McGovern-McNab combination. A POSITA would be motivated to do this in order to provide to the administrator security assessment scores based on each entity’s network, since each entity would likely have different networks wherein evaluation of the particular entity’s network would provide useful information.

146. As discussed above, a POSITA would have been motivated to combine McNab and McGovern such that the results of Nessus’s security reports (*i.e. derived observations*) would be used in McGovern’s security assessment score. While McGovern does not expressly describe how to generate such a score, since Nessus

describes its security reports as being categorized into “three severity levels,” each severity level would be given different weights within the security assessment score.

5. 1[e]: through a communication network, exposing to users in a user interface of a portal, information about the cyber-security ratings of the entities that is useful to decisions about entities with which to do business.

147. In my opinion, McNab in view of McGovern disclose this limitation.

148. McGovern discloses that the security assessment scores, *i.e. cyber-security ratings of the entities*, “are recorded and organized into one or more [scoring] data structures.” *Id.* at ¶41. Such data structures “may be incorporated in a software program *with a user interface*” and the software can be either a “stand-alone, self-contained package” or “designed to run as a Web-based service or application” wherein users can “*access the scoring* and related functionalities *remotely via standard browsers* or similar *user interfaces*.” McGovern at ¶¶42, 45.

149. A POSITA would understand that the security assessment score is *information useful to decisions about entities with which to do business*. As McGovern explains, the ISA score is “a composite score indicative of an overall security assessment of the information system.” *Id.* at ¶42. These composite scores are generated from “numerical scores” that indicate “degrees of compliance with security standards.” *Id.* at ¶37. The numerical scores also are indicative of “goodwill and negative publicity.” *Id.* All of these are information *useful to decisions about entities with which to do business* because it would be more desirable to do business

with an entity with a better security score and less desirable to do business with a worse security score. This is because the security score is indicative of the security of the *entity's network*, which is important if, for example, doing business with that entity involves transferring confidential or private information.

C. Claim 2: The method of claim 1 in which the technical assets comprise an IP space.

150. McNab discloses this claim. As discussed above in limitation 1[a], McNab discloses discovering the IP addresses of assets such as mail, SMTP or FTP servers using “Forward DNS Grinding.” A POSITA would understand that network assets such as mail, SMTP, or FTP servers have IP addresses within an IP space (*i.e.* the space of valid IP addresses for the network) defined by a subnet mask. Indeed, McNab discloses “the subnet mask used by the target host.” McNab at 43. A POSITA would understand that a subnet mask defines the IP space by indicating which bits of the IP address define the subnet IP space. For example, a typical subnet mask is 255.255.255.0, where “255” indicates an 8-bit binary number consisting entirely of 1’s. This indicates that the first three parts of assets within a network must have IP addresses that are the same (*e.g.*, 192.168.1.XXX), whereas the fourth component (which has a corresponding subnet mask of 0) can be used to identify the asset. Thus, a subnet mask of 255.255.255.0 where the gateway has an IP address of 192.168.1.1 indicates that the IP space is 192.168.1.1 through 192.168.1.255.

151. In addition, McNab discloses that Nessus can target “ranges of network addresses” for its vulnerability scans, indicating that the *assets comprise an IP space*. *Id.* at 385.

D. Claim 3: The method of claim 1 in which maintaining an entity map comprises querying a regional Internet registry.

152. As discussed in limitation 1[a], McNab logically walks through the Internet-based options for mapping technical asserts on a network of a target organization. McNab further discloses in Chapter 3 that “*Regional Internet Registries* (RIRs) store useful information (primarily as *network*, *route*, and *person* objects) relating to IP network blocks” and “[e]ach respective *regional registrar’s WHOIS database* contains information relevant to that particular region” where the “IP WHOIS database objects define which areas of Internet space are registered to which organizations.” McNab at 23-24. Thus, “*RIRs can be queried* to glean useful information (including names of technical IT staff, *details of IP network blocks*, and physical office locations).” *Id.*

E. Claim 4: The method of claim 3 in which maintaining an entity map comprises submitting a name of one of the entities to the regional Internet registry.

153. As discussed in Claim 3, McNab discloses querying a regional Internet registry and “regional registrar’s WHOIS database” for information relating to “which areas of Internet space are registered to which organizations.” McNab at 23-24. McNab discloses that “*American Registry for Internet Numbers (ARIN)*” and “*Asia*

Pacific Network Information Centre (APNIC)” are two of many such Regional Internet Registries that can be queried. *Id.* McNab then specifically discloses “[q]uerying WHOIS databases to enumerate objects for a given company” and, in *Example 3-2*, discloses submitting “a query of *nintendo* to enumerate all the objects in the *ARIN database* for Nintendo” as well as in *Example 3-3*, “reissue the query to *APNIC* and *other registrars.*” *Id.* at 25-26.

Example 3-2. Enumerating the Nintendo objects in ARIN

```

$ whois nintendo -h whois.arin.net
Nintendo North America (NNA-21)
Nintendo of America (TEND)
Nintendo of America (NINTEN-1)
Nintendo Of America inc. (NINTEN)
Nintendo of America, Inc. (NINTE-1)
Nintendo of America, Inc. (NINTE-2)
Nintendo Network Administration (NNA12-ARIN) netadmin@noa.nintendo.com +1-425-882-2040
Nintendo Of America inc. (AS11278) NINTENDO 11278
Nintendo North America SAVV-S233299-1 (NET-207-149-2-192-1) 207.149.2.192 - 207.149.2.199
Nintendo North America SAVV-S263732-2 (NET-209-67-111-168-1) 209.67.111.168 - 209.67.111.175
Nintendo North America SAVV-S233299-2 (NET-216-74-145-64-1) 216.74.145.64 - 216.74.145.127
Nintendo of America NET-NOA (NET-206-19-110-0-1) 206.19.110.0 - 206.19.110.255
Nintendo of America NINTENDO-COM (NET-205-166-76-0-1) 205.166.76.0 - 205.166.76.255
Nintendo Of America inc. NOA (NET-192-195-204-0-1) 192.195.204.0 - 192.195.204.255
Nintendo of America, Inc. SAVV-S263732-3 (NET-216-32-20-248-1) 216.32.20.248 - 216.32.20.255
Nintendo of America, Inc. SAVV-S263732-3 (NET-209-67-106-128-1) 209.67.106.128 - 209.67.106.255
NINTENDO ABOV-T461-209-133-66-88-29 (NET-209-133-66-88-1) 209.133.66.88 - 209.133.66.95
NINTENDO ABOV-T461-209-133-66-72-29 (NET-209-133-66-72-1) 209.133.66.72 - 209.133.66.79
Nintendo MFN-N389-64-124-44-48-29 (NET-64-124-44-48-1) 64.124.44.48 - 64.124.44.55

```

F. Claim 5: The method of claim 4 in which submitting the entity name includes submitting spelling variations of the entity name.

154. McNab renders it obvious to submit different spelling variations of an entity to a regional Internet registry. McNab discloses that “[e]ach respective regional registrar’s WHOIS database contains information relevant to that particular region.” *Id.* at 24-25. For example, the Réseaux IP Européens RIPE WHOIS database “doesn’t

contain information about network space and other objects that are found in the Americas” and that American Registry for Internet Numbers ARIN “only contains details of North American objects.” *Id.* A POSITA would have found it obvious that, for example, spelling variants exist between American and British English for even the same English, like *colour* and *color*, or *aluminium* and *aluminum*. Thus, a POSITA would have found it obvious that in some circumstances, for a single entity, different spelling variations of the entity name need to be submitted to different regional Internet registries.

G. Claim 6: The method of claim 1 in which maintaining an entity map comprises using a domain name associated with the entity.

155. As discussed in limitation 1[a], McNab discloses maintaining an entity map. Specifically, McNab discloses that “Internet newsgroup searches can also be queried” and discloses in Figure 3-3 querying Google Groups “using the query string *symantec.com*.” McNab further discloses in *Example 3-1* submitting “a query of *blah.com*, revealing useful information regarding *the domain*, its administrative contacts, and authoritative DNS name servers.” McNab at 21-22. The information gathered in *Example 3-1* shows that the domain “*blah.com*” is associated with the entity “*blah! Sociedade Anonima Serv e Com.*” *Id.* at 22-23.

H. Claim 7: The method of claim 6 in which maintaining an entity map comprises using a dig tool on the domain name to determine a published IP address.

156. McNab discloses using utilities “such as *nslookup*, *host*, and ***dig*** ... to issue DNS requests *relating to domains and IP address blocks* identified” and that the “***dig*** client” is capable of querying DNS servers directly. McNab at 30. Specifically, in *Example 3-8*, McNab discloses using ***dig*** to query domain names “*ucia.gov*” and “*relay2.ucia.gov*” directly to obtain the DNS zone and retrieve “following ***IP address blocks***” and “[i]nternal and external ***IP addresses*** for a number of systems.”

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

```

dig tool
$ 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.

```

Annotations in the image:

- dig tool**: Points to the command line.
- domain names**: Points to the domain names in the output.
- IP addresses**: Points to the IP addresses in the output.

I. Claim 8: The method of claim 1 in which maintaining an entity map comprises probing IP addresses around a published IP address.

157. As discussed in limitation 1[a], McNab discloses maintaining an entity map. Specifically, as discussed above in Claim 7, in *Example 3-8*, McNab discloses gathering “[i]nternal and external **IP addresses** for a number of systems” from domain names “ucia.gov” and “relay2.ucia.gov.” McNab at 32-33. McNab further discloses that an IP address block “198.81.129.0” can be identified from the IP information

gathered. *Id.* For example, in *Example 3-8* from McNab, several hosts have IP addresses that begin with 198.81.129.xxx, such as “mail1.ucia.gov” which has the IP address 198.81.129.68—from these hosts, it is known that they are part of an IP address block “198.81.129.0.” *Id.* Then, in *Example 3-12*, McNab discloses using tools like GHBA to perform reverse DNS sweeping of IP addresses in the IP address block “198.81.129.0” —***around the published IP addresses*** identified in *Example 3-8*. *Id.* at 36-37. For example, one of the published IP addresses identified in *Example 3-8* is 198.81.129.197 and the probed IP addresses in *Example 3-12* include IP addresses around this identified published IP addresses, such as from 198.81.129.196 to 198.81.129.193. *Id.* at 33, 36-37.

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
waits.ucia.gov. 3600 IN CNAME relay2.ucia.gov.
```

published IP addresses
in block 198.81.129.0

Example 3-12. Using GHBA to perform a reverse DNS sweep

```
$ ghba 198.81.129.0
Scanning Class C network 198.81.129...
198.81.129.20 => ddss.ucia.gov
198.81.129.21 => ddssdata.ucia.gov
198.81.129.22 => ddsstest.ucia.gov
198.81.129.23 => ddsstestdata.ucia.gov
198.81.129.68 => mail1.ucia.gov
198.81.129.69 => relay2z.ucia.gov
198.81.129.71 => mail1out.ucia.gov
198.81.129.100 => cia.ucia.gov
198.81.129.101 => www2.ucia.gov
198.81.129.103 => www.intelligence.gov
198.81.129.104 => www.nctc.gov
198.81.129.106 => wits2.nctc.gov
198.81.129.146 => mail2out.ucia.gov
198.81.129.148 => mail2.ucia.gov
198.81.129.186 => relay7.ucia.gov
```

IP addresses around
a published IP

Example 3-12. Using GHBA to perform a reverse DNS sweep (continued)

```
198.81.129.193 => relay1.ucia.gov
198.81.129.194 => relayz.ucia.gov
198.81.129.195 => relay2a.ucia.gov
198.81.129.196 => relay2b.ucia.gov
198.81.129.222 => ex-rtr-129.ucia.gov
198.81.129.230 => res.odci.gov
198.81.129.231 => comm.ucia.gov
```

around
198.81.129.197

J. Claim 10: The method of claim 1 in which maintaining an entity map comprises using a name server lookup to verify the derived server name.

158. As discussed in limitation 1[a], McNab discloses maintaining an entity map. Specifically, McNab discloses that utilities “such as *nslookup*, host, and dig are used to issue DNS requests relating to domains and IP address blocks identified.” McNab at 30. A POSITA would have understood that the name “nslookup” comes from “*name server lookup*.”

159. McNab discloses “[u]sing nslookup to enumerate basic domain details” about a domain name. EX1005, 30-31. It would have been obvious that

TXDNS uses nslookups to determine which server names are verified to be real. This is because, as explained above, TXDNS uses known prefixes to generate server name guesses and some methodology to determine whether they exist or not—nslookup would have been an obvious design choice to confirm, since it is the standard method for confirming whether a domain name has an IP address.

K. Claim 11: The method of claim 2 comprising using the IP space for collecting the cyber-security characterizations.

160. As discussed in Claim 2, because a POSITA would understand that the targeted assets disclosed in McNab have IP addresses within an IP space, the collection of cyber-security characterizations for those assets would involve using the IP space. For example, McNab discloses that the Nessus Security Scanner targets assets with IP addresses (*see* below figure showing results of scan for various assets with IP addresses), which means when Nessus collects cyber-security characterizations, it used the IP space (*e.g.*, an IP address within the IP space) to collect the characterizations.




Figure 15-6. NessusClient 3 completed scan

- L. Claim 12: The method of claim 1 in which the technical assets that are associated with cyber-security characterizations of the respective entities comprise technical assets that the entity owns, controls, uses, or is affiliated with.**

161. McNab discloses that the entity assets are technical assets affiliated with the entity or used by the entity. As discussed above for limitation 1[a], McNab discloses generating an entity map for an organization's "networks," *i.e.* entities.

McNab at 17. The assets are therefore affiliated with networks (*i.e. entities*) because they belong to the network.

M. Claim 13: The method of claim 1 in which the technical assets comprise ranges of IP addresses.

162. As discussed above for Claim 2, McNab discloses that the technical assets comprise an IP space because they have subnet masks that define the range of IP addresses for which an asset on the same subnetwork may have. The example given in Claim 2 was a range of IP addresses from 192.168.1.1 to 192.168.1.255. Thus, for the same reasons discussed above for Claim 2, McNab discloses that the technical assets comprise ranges of IP addresses.

N. Claim 14: The method of claim 1 in which the technical assets comprise servers.

163. As discussed above for limitation 1[a], McNab discloses “SMTP gateways” or “mail relay servers” as assets. A POSITA would understand that an SMTP gateway is a server, because it provides the service of routing e-mail messages from the organization for outbound delivery.

O. Claim 15: The method of claim 1 in which maintaining an entity map comprises using information received from Internet Service Providers.

164. McNab discloses that network enumeration involves performing a “BGP query[]” where a domain name’s “Autonomous System (AS)” number is determined from WHOIS information and then queried to determine an IP address block. McNab at 28-29. As McNab explains, the IP address block for a particular AS

number is allocated to an “ISP.” *Id.* Thus, McNab discloses using information received from an ISP to maintain the entity map.

P. Claim 16: The method of claim 1 in which maintaining an entity map comprises using heuristics.

165. The ’524 patent explains that “heuristics” include at least: “Using the ‘dig’ ... tool to determine any IP information published by an entity” or “us[ing] the IP addresses and domain names published to find ranges of IP addresses actually used” or attempting to guess “servers whose names may be guessed” such as adding the “mail,” “smtp” or “ftp” prefix to a domain name. Ex. 1001 at 9:16-38.

166. As I discussed above for limitation 1[a], McNab discloses adding the mail prefix to a domain name to guess server names. And as I discussed for Claim 7, McNab also discloses using “dig.” Thus, McNab discloses *maintaining an entity map by using heuristics*.

Q. Claim 18: The method of claim 1, wherein annotating comprises prepending to the domain name at least one of “mail,” “smtp,” and “ftp”.

167. As discussed in limitation 1[a], McNab discloses cross-referencing common hostnames with known domains to retrieve valid domains. Specifically, in Example 3-11, McNab discloses performing “a hostname grinding attack against the *bankofengland.co.uk* domain, using a small dictionary of *common mail server names*” including “*mail.bankofengland.co.uk*,” “*mail2.bankofengland.co.uk*,” and “*mailhost.bankofengland.co.uk*.” McNab 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



168. Further, in *Example 5-6*, McNab also discloses using a forward DNS grinding tool TXDNS to “reveal valid *apple.com* hostnames” including “*ftp.apple.com*” and “*webmail.apple.com*.” McNab at 85-86.

R. Claim 19

1. **19[pre], 19[a]: A computer implemented method comprising: on one or more computers operating with a database, maintaining an entity map that maps technical assets that comprise an IP space to respective entities with which the assets are associated,**

169. Limitation 1[pre] recites “[a] computer implemented method.”

170. Limitation 1[a] recites “on one or more computers” and “a database.”

171. Limitation 1[a] also recites “maintaining ... an entity map that maps technical assets to respective ... entities with which the assets are associated.”

172. Claim 2 recites that the technical assets “comprise an IP space.”

173. Thus, in my opinion, McNab in view of McGovern renders obvious these limitations for the same reasons I gave for limitations 1[pre], 1[a] and claim 2.

2. **19[b]: the maintaining of the entity map comprising (a) automatically using sensors on the Internet to collect**

externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities, (b) associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name, (c) submitting names of the entities to regional Internet registries using various spellings of the names, (d) using domain names associated with the entities, (e) probing IP addresses around published IP addresses, (f) using a name server lookup to verify IP addresses associated with domain names, (g) using information received from Internet Service Providers, (h) annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses, and (i) applying heuristics,

174. Limitation 1[b] recites “automatically using sensors on the Internet to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities,” which is identical to 19[b]’s step (a).

175. Limitation 1[b] also recites “associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name” which is identical to 19[b]’s step (b).

176. Claim 4 recites “submitting a name of one of the entities to the regional Internet registry” and Claim 5 recites “submitting spelling variations of the entity name,” which is substantively identical to 19[b]’s step (c).

177. Limitation 1[a] describes using a domain name associated with entities, which is substantively identical to 19[b]’s step (d).

178. Claim 8 recites “probing IP addresses around a published IP address” which is substantively identical to 19[b]’s step (e).

179. Claim 10 recites “using a name server lookup to verify the derived server name,” which is substantively identical to 19[b]’s step (f).

180. Claim 15 recites “using information received from Internet Service Providers,” which is substantively identical to 19[b]’s step (g).

181. Limitation 1[a] recites “annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses,” which is substantively identical to 19[b]’s step (h).

182. Claim 18 recites “using heuristics” which is substantively identical to 19[b]’s step (i).

183. Thus, limitation 19[b] is rendered obvious by McNab in view of McGovern for the reasons explained above for limitations 1[a] and 1[b] and claims 4, 8, 10, 15, and 18.

3. **19[c]: automatically deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (a) a number of technical assets that have been reported to be malicious and (b) a duration of detected malicious activity associated with the technical assets,**

184. In my opinion, this limitation is rendered obvious by McNab in view of McGovern for the same reasons I explained above for limitation 1[c].

4. **19[d]: automatically generating a cyber-security rating for each of the entities using the entity map and the derived observations, and**

185. In my opinion, this limitation is rendered obvious by McNab in view of McGovern for the same reasons I explained above for limitation 1[d].

5. **19[e]: through a communication network, exposing to users in a user interface the cyber-security ratings of the entities.**

186. In my opinion, this limitation is rendered obvious by McNab in view of McGovern for the same reasons I explained above for limitation 1[e].

S. Claim 20

1. **20[pre], 20[a], 20[b]: A system comprising: one or more computers programmed to facilitate operations comprising: maintaining in a database an entity map that maps technical assets to respective companies or other entities with which the assets are associated, wherein mapping technical assets to respective companies or other entities comprises determining IP information by annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses,**

187. In my opinion, these limitations is rendered obvious by McNab in view of McGovern for the same reasons I explained above for limitation 1[pre] and 1[a].

2. **20[c]: collecting, using sensors on the Internet, externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities and associating the observable cyber-security characterizations**

with a respective entity based at least in part on the derived server name,

188. In my opinion, this limitation is rendered obvious by McNab in view of McGovern for the same reasons I explained above for limitation 1[b].

- 3. 20[d]: deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (i) a number of technical assets that have been reported to be malicious and (ii) a duration of detected malicious activity associated with the technical assets,**

189. In my opinion, this limitation is rendered obvious by McNab in view of McGovern for the same reasons I explained above for limitation 1[c].

- 4. 20[e]: generating a cyber-security rating for each of the entities using the entity map and the derived observations, and**

190. In my opinion, this limitation is rendered obvious by McNab in view of McGovern for the same reasons I explained above for limitation 1[d].

- 5. 20[f]: through a communication network, exposing to users in a user interface of a portal, information about the cyber-security ratings of the entities that is useful to decisions about entities with which to do business.**

191. In my opinion, this limitation is rendered obvious by McNab in view of McGovern for the same reasons I explained above for limitation 1[e].

VIII. GROUND 2: COLE IN VIEW OF MCNAB AND MCGOVERN RENDER OBVIOUS CLAIMS 1-8, 10-16, 18-20

192. Cole discloses a testing system allowing “highly automated testing of vulnerabilities to intrusion on a target network” including, among other things,

“identification of target network topology and target computers,” “active assessment of vulnerabilities based on information acquired from target computers,” “quantitative assessment of target network security and vulnerability,” and “hierarchical graphical representation of the target network, target computers, and vulnerabilities in a test report.” Cole at Abstract. As discussed above, McNab discloses using “Forward DNS Grinding” to generate likely hostnames using naming conventions and McGovern discloses generating the duration of malicious activity, which would have been obvious to combine with Cole. McNab also discloses various other methods for enumerating network assets, such as querying Regional Internet Registries, performing BGP queries, using the “dig” and “nslookup” tools, and determining an IP space, which also would have been obvious to combine with McNab.

193. As I discussed for Ground 1, McGovern discloses determining the duration of malicious activity, and a POSITA would have likewise found it obvious to combine that teaching with Cole.

A. Rationale to Combine

194. In my opinion, a POSITA would have found it obvious to combine Cole’s network vulnerability assessment methods with McNab’s “Forward DNS Grinding” as well as other methods for generating IP addresses of assets and McGovern’s teachings of determining the duration of a malicious activity as well as its


teachings of displaying network vulnerability assessment scores in a web-based user interface.

1. McNab

195. In my opinion, a POSITA would be motivated to incorporate McNab’s teachings of “Forward DNS Grinding” and other methods for enumerating network assets because such teachings were known to provide another additional reliable methods for mapping additional network assets—which as McNab teaches was useful and necessary for performing network assessments.

196. As discussed in Ground 1, McNab discloses a method of discovering network assets called “Forward DNS Grinding” using a tool called TXDNS where a domain name is annotated with a “dictionary of common mail server names” to determine the IP address of a mail server or SMTP server or another dictionary to determine the IP address of an FTP server. McNab at 35, 85-86. In essence, a domain name is chosen (such as “bankofengland.co.uk” or “apple.com”) and a prefix that is conventionally known to be associated with an SMTP mail server or FTP server (such as “mail” or “ftp”) is prepended to create a possible domain name (such as “mail.bankofengland.co.uk” and “ftp.apple.com”). Dozens of such possible names are created and each one is sent to a DNS server to attempt to resolve into an IP address. *Id.* If the DNS server returns an IP address, then that is a known asset of that domain. For example, in the below, of the 98 possible mail server names that were generated,

three were identified and had IP addresses determined: “mail.bankofengland.co.uk,” “mail2.bankofengland.co.uk” and “mailhost.bankofengland.co.uk.” McNab at 39.



197. A POSITA would have been motivated to incorporate these teachings in McNab to provide a well-known and reliable method of determining the IP addresses of assets. Cole already discloses that as “a first step” it “creates a scan list of IP addresses” for a target network (*i.e.* domain name). Cole at 15:1-3. However, Cole does not describe how to generate the list of IP addresses. A POSITA would have understood that one well-known and reliable method of generating IP addresses of potential network assets is McNab’s “Forward DNS Grinding” method using conventional tools such as TXDNS. A POSITA would have been motivated to give effect to Cole’s instruction to generate a list of IP addresses.

198. In addition, a POSITA would have been motivated to use McNab’s teachings because generating an accurate list of IP addresses is important for properly performing vulnerability assessments on the network. As Cole teaches, the list of IP addresses is used to determine the target computers for vulnerability testing.

Moreover, as discussed in the Technical Background, numerous papers and resources were known to provide network enumeration prior to vulnerability scanning. *See, e.g.*, Ex. 1015 (“Anatomy of a Hack” begins with “Identification of Internet-based points of presence”), Ex. 1016 at 11-12 (“network mapping tool [to] identify the services that run on a large group of hosts and provide a quick way of building individual profiles of the target IT system(s)”), Ex. 1017 at 7 (“The first step to any security engagement is the initial discovery process”). Thus, a POSITA would find it desirable to have a list of IP addresses that includes *all* computers within the network, or at least all computers that are accessible via the Internet. Thus, a POSITA be motivated to use any and all mechanisms that may result in determining assets within the network, as that would increase the reliability of Cole’s vulnerability assessment and increase the overall security of the network. Indeed, if Cole’s system failed to test an asset within an entity’s network, that could prove dangerous if the asset became vulnerable—Cole’s system would be unable to make that determination and the administrators would not be able to timely mitigate the issue.

199. For similar reasons, a POSITA would have been motivated to also incorporate other network enumeration techniques described in McNab. For example, McNab describes querying “Regional Internet Registries” which “store useful information ... relating to IP network blocks.” McNab at 23-24. McNab also describes the use of the “dig” and “nslookup” tools which provide further information for

network enumeration. *Id.* at 30. McNab further describes BGP querying to determine IP address blocks. *Id.* at 28-29. All of these methods would have been obvious to use with Cole for similar reasons described above, namely that they were well-known network enumeration techniques that a POSITA would be motivated to include any technique that would increase the likelihood of enumerating all assets.

200. In my opinion, a POSITA would have reasonably expected success in combining Cole with McNab's "Forward DNS Grinding." McNab teaches that "Forward DNS Grinding" can be performed using a simple command-line tool called TXDNS, and it would have been well-within the capabilities of a POSITA to incorporate that command-line tool in the automated system of Cole to generate a list of IP addresses for performing vulnerability assessments.

201. Finally, in my opinion, Cole and McNab are directed to the same field of endeavor as the '524 patent because all three references are directed to the field of network security. Cole at Abstract ("A system and method provide comprehensive and highly automated testing of vulnerabilities to intrusion on a target network"); McNab at xvi ("This book tackles one single area of information security in detail: that of undertaking IP-based network security assessment in a structured and logical way."); Ex. 1001 at Abstract ("A method and system for creating a composite security rating from security characterization data of a third party computer system.").

2. McGovern

202. In my opinion, a POSITA would be motivated to incorporate McGovern's teachings of calculating a duration of malicious activity as well as a remote web-based user interface because such teachings were known to increase the security of a network as well as increase user-friendliness of the system and provide the user an easier assessment of the network, including by remotely accessing such assessments.

203. McGovern discloses the capability of determining the "average response or patch time after security breach detection." McGovern at ¶38. A POSITA would be motivated to incorporate this teaching into Cole, which discloses various methods and measurements for evaluating the security of a network, because of the known benefits. A POSITA would understand that it would be beneficial to include the duration of a security breach in the network vulnerability assessments of Cole, because the duration of a security breach is an important factor in determining overall network security. A POSITA would understand that a security breach is an incredibly dangerous event that puts all of the organization's data and resources at risk of damage or theft. A POSITA would further understand that a security breach can cause damage to the company via various ways, including the theft of confidential information, putting the business at a competitive disadvantage, or the theft of customer or employee information, therefore harming said customers or employees, harming the company's

reputation, and increasing the liability risk to the company. A POSITA would therefore understand that it would be highly desirable to minimize the duration of any security breach—thus, it would be beneficial for a network administrator to know whether the network and the organization’s IT policies are robust and proactive to such security breaches (resulting in a minimal duration) or if they are particularly vulnerable (because of a lengthy duration).

204. McGovern also discloses the capability of displaying a network assessment score in “a software program with a user interface,” which can be “designed to run as a Web-based service or application” wherein users can “access the scoring and related functionalities remotely via standard browsers or similar user interfaces.” McGovern at ¶¶42, 45. A POSITA would be motivated to incorporate these teachings into Cole because it allows network administrators access to the assessment reports from anywhere an Internet browser is available—at the time of the ’524 patent, that would have included smartphones. A POSITA would therefore be motivated to make this combination in order to increase ease-of-use and accessibility.

205. A POSITA would have combined Cole and McGovern with a reasonable expectation of success. Cole is dedicated to performing vulnerability assessments of target networks—including customizable tests using scripts. Cole at Abstract. Adding McGovern’s teaching of determining the duration of a security breach would simply be another test that an administrator could perform that was well-

within the capability of a POSITA at the time of the '524 patent. Cole does not state any limit on the number of vulnerability tests that could be performed, and a POSITA would understand that adding an additional test would not pose any significant hurdle in computing power or evaluation time, especially in view of the benefits of said test described above. Indeed, because Cole already discloses determining the presence of malicious software, and discloses tracking of fixes for those vulnerabilities, a POSITA would understand that McGovern's teachings could be readily implemented by determining the duration of when the malicious activity was first detected to when it was verified that the vulnerability was fixed. Cole at 70:54-58 ("components for tracking detected vulnerabilities, and for performing targeted vulnerability fix verification tests of specific host computers").

206. A POSITA would have also reasonably expected success in combining Cole with McGovern's remote user interface. Indeed, Cole already describes a testing system that "may be a web-based interface that is accessible to network users via an ordinary web browser" (Cole at 71:50-59) and it would have been well-within the skills of a POSITA to incorporate McGovern's teaching of providing reports through a web-based interface.

207. Finally, Cole and McGovern are directed to the same field of endeavor as the '524 patent, specifically the field of network security. Cole at Abstract ("A system and method provide comprehensive and highly automated testing of

vulnerabilities to intrusion on a target network”); McGovern at Abstract (“Methods for information security assessment and data risk scoring are disclosed.”); Ex. 1001 at Abstract (“A method and system for creating a composite security rating from security characterization data of a third party computer system.”).

B. Claim 1

1. **1[pre], 1[a] A computer implemented method comprising: on one or more computers, maintaining in a database an entity map that maps technical assets to respective companies or other entities with which the assets are associated, wherein mapping technical assets to respective companies or other entities comprises determining IP information by annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses**

208. With respect to the limitation “*A computer implemented method comprising: on one or more computers, maintaining in a database an entity map that maps technical assets to respective companies or other entities with which the assets are associated,*” Cole discloses targeting a network to build lists of assets and their IP addresses and storing them in a database, thus generating an *entity map*. Cole further discloses targeting multiple networks, thereby forming an entity map. Indeed, Cole describes these networks as belonging to “an entity seeking to ‘harden’ its network.” Cole at 1:55-56.


209. Cole discloses a “highly automated” system for “identification of target network topology.” Cole at Abstract. Specifically, Cole discloses that “a first

step” is “defin[ing] the target network and creat[ing] a scan list of IP addresses.” Cole at 15:1-3. From this list of IP addresses, a “ICMP host discovery routine” is executed to “determine that certain target computers [*i.e. technical assets*] are ‘live’ or ‘probably live,’ and the associated IP addresses are added to a respective live *database* 324 or probably live *database* 326 accordingly.” *Id.* at 15:8-17. Thus, this process generates a mapping of *technical assets to an entity which the assets are associated* which is *maintained in a database*.

210. Cole discloses that the discovery of assets can be performed on *multiple entities*. Specifically, Cole discloses that its method “defines the target network” and that also “the method operates on multiple target networks concurrently.” *Id.* at 15:1-3, 18:19-21. Cole describes that every entity has “its network” it seeks to “harden.” *Id.* at 1:55-56. Thus, a POSITA would understand that each target network belongs to a different entity. These mappings are maintained in databases, thereby forming an *entity map*.

211. With respect to the limitation “*wherein mapping technical assets to respective companies or other entities comprises determining IP information by annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses,*” While Cole discloses “creat[ing] a scan list of IP addresses” from which live assets are discovered, Cole does not expressly describe how to create

that initial list of IP addresses. However, as discussed above, McNab discloses a method called “Forward DNS Grinding” which *annotates a domain name with naming conventions for a mail server, SMTP server, and FTP server* (such as appending the prefix mail or ftp to a domain name like “bankofengland.co.uk”) *to derive a likely server name* (such as “mail.bankofengland.co.uk” or “ftp.apple.com”) and then *determining the IP addresses* of those *derived server names* if they are registered in a DNS server.



212. As discussed above, given the known benefits of “Forward DNS Grinding,” it would have been obvious to incorporate this teaching into Cole to generate some of the IP addresses in the scan from which *assets are discovered and mapped to entities*.

2. 1[b] automatically using [sic] sensors on the Internet to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities and associating the observable

**cyber-security characterizations with a respective entity
based at least in part on the derived server name**

213. Cole discloses that its “highly automated” system performs “automatic detection, monitoring and reporting of network vulnerabilities.” Cole at Abstract.

214. With respect to the limitation “*automatically useing* [sic] *sensors on the Internet,*” Cole discloses that its methods are performed by a “network security system 100” which is “outside the target network 102” and “connected to the target network 102 through a network such as *the Internet.*” Cole at 11:34-37. Cole further discloses that this system performs “*automatic detection.*” *Id.* at Abstract. In particular, after the network assets are mapped, Cole discloses performing an “active assessment routine” and a “scoring routine” which “establishes a vulnerability score ... for the network based on the results of the active assessment and based on the vulnerability information in the known vulnerability database 366.” *Id.* at 17:59-18:5. Cole further discloses that it performs this vulnerability testing using “automatic script[s]” for each vulnerability.” *Id.* at 41:11-15. Thus, Cole discloses using “[m]ultiple FASL scripts ... run in parallel,” each of which a POSITA would understand to be one of many *sensors on the Internet.*

215. With respect to the limitation “*to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities,*” Cole discloses that it performs vulnerability testing using “FASL scripts,” which collect *externally observable cyber-security characterizations.* Cole discloses

that such scripts are run on “target system[s],” *i.e. technical assets that have been mapped to an entity*. Cole at 41:11-15; *see also* 53:11-16 (“the m_nIPAddress member of the FASL variable has been set to the IP address of the target machine upon which the FASL script is to run. All network activity that takes place in a FASL script uses this IP address as its destination ...”). Because the security system is external to the network and targets a specific IP address, a POSITA would understand that the FASL scripts collect *externally observable* information. *See also* Cole at 60:53-54 (“Quantitative Scoring of Target Network *External Vulnerability*”); 63:53-59 (“[T]he preceding calculations provide a measure of the security of ports visible *to the Internet* and other *external networks*, the ease with which an unauthorized user could breach the network firewall, and the ease with which an unauthorized user could access data inside the network from a computer *outside of the network*, among other things.”).

216. A FASL script “permits one to perform active assessment of potential vulnerabilities on a known target computer with known open ports.” Cole at 58:55-58. These scripts collect information needed to perform the vulnerability assessment, *i.e. externally observable cyber-security characterizations*. *Id.* at 59:28-31 (“[I]nformation collected from a successful vulnerability test ... is used to advantageously further test the vulnerability of the target network and individual target computers.”). Cole generally describes these assessments as related to “the security of ports visible to the Internet and other external networks, the ease with which an

unauthorized user could breach the network firewall, and the ease with which an unauthorized user could access data inside the network from a computer outside of the network, among other things,” all of which a POSITA would consider to be *cyber-security characterizations*. Cole at 63:53-59.

217. In an example vulnerability test, the script collected information about which open ports were found on a target computer and then performs tests on those ports to determine their vulnerability, *i.e. externally observable cyber-security characterizations*. *Id.* at 59:37-57.

218. With respect to the limitation “*and associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name,*” Cole discloses that “the information collected by the FASL script from the target computer,” *i.e. cyber-security characterizations*, are “stored in a database related to the target computer for later reporting.” Cole at 41:26-31.

219. Because this information is later used to “provide[] objective, quantitative indicia of the overall security of the target *network*, *i.e. the network of the respective entity*, then a POSITA would understand that the characterizations are associated with the *respective entity*, especially since, as discussed for limitation 1[a], the system already associates the target computer with the target network. A POSITA would have also considered it obvious to associate the characterizations with the respective entity (and its network) because they would be motivated to give effect to

the security assessment score that is later generated for the overall network. A POSITA would have also been motivated by the benefits of well-organized data in that subsequent review of the security characterizations would be made much easier if they are associated with the entity for which the assessment was made. A POSITA would have had a reasonable expectation of success because Cole already discloses that the characterizations are stored in a database and it would be well-within the capabilities of a POSITA to reference that data to information about the entity already stored in the database as discussed in limitation 1[a].

220. Although Cole does not describe a *derived server name*, as discussed above, McNab describes using “Forward DNS Grinding” to generate a derived server name (such as ftp.apple.com). A POSITA would have understood that the characterizations are associated with the entity *based on a derived server name* because the characterizations are associated with a target computer that was identified by an IP address that, in accordance with McNab’s “Forward DNS Grinding,” was determined by generating a derived name. In addition, a POSITA would have considered it obvious that the FASL scripts could identify the target computer using the *derived server name* as that would allow the benefit of more easily identifying the target computer and would simply require a single DNS request to resolve the server name to an IP address.

3. 1[c] automatically deriving observations about the technical assets from the collected cyber-security characterizations,

wherein the derived observations comprise (i) a number of technical assets that have been reported to be malicious and (ii) a duration of detected malicious activity associated with the technical assets

221. With respect to the limitation *“automatically deriving observations about the technical assets from the collected cyber-security characterizations,”* Cole discloses that each FASL script will automatically use the information it collects (*i.e. collected cyber-security characterizations*) to “respond with a success or failure result noting whether the target computer was vulnerable or not to the given vulnerability identification code,” *i.e. deriving observations about the technical asset.* Cole at 41:23-26.

222. With respect to the limitation *“wherein the derived observations comprise (i) a number of technical assets that have been reported to be malicious,”* Cole describes that FASL scripts can be programmed to determine the presence of malicious software on an asset. Specifically, Cole notes that one type of vulnerability for which a FASL script will assess is a “trojan.” Cole at 60:20-52 (“TABLE: Vulns”). As Cole explains it (and as a POSITA would have readily understood), a “trojan” or “[t]rojan horse ... can compromise a host security, network security or both” and “allow an attacker to establish a remote connection to the host,” to, for example retrieve “usernames and passwords, enumeration of other network information resources, retrieval of host data, or launching an attack against other systems/networks.” Cole at

67:1-17. Thus, a FASL script that indicates it has detected a trojan in a target computer has *derived an observation reporting the asset to be malicious*.

223. Cole also describes *deriving an observation* where “all vulnerabilities found in the target network have been counted” using a “vulnerability counter.” Cole at 63:3-4, 10. Cole describes that the trojan detected by a FASL script is part of a set of “Vulns” that are scanned for. *Id.* at 60:20-52. In a situation where vulnerability scans are only determining whether assets are reported to be malicious, the “vulnerability counter” would be understood to be a *derived observation of the number of technical assets that have been reported to be malicious*.

224. In the alternative, it would have been obvious to derive the *number of assets reported to be malicious*. Since Cole already reports which assets are malicious, a POSITA would have found it obvious to collect that information and derive the number of said malicious assets. A POSITA would be motivated to do this because deriving such a number would allow the number to be displayed to the administrator, thereby giving the administrator an easily-accessible number to indicate the health of the network. Deriving the number would also allow the system to alert the administrator if the number exceeds zero. A POSITA would have a reasonable expectation of success in doing this, because since Cole describes that its vulnerability scanning is performed using FASL scripts, it would be well-within the skill of a

POSITA to include within the script that records the number of assets to be malicious. The script could further email said number to the administrator.

225. With respect to the limitation ***“and (ii) a duration of detected malicious activity associated with the technical assets,”*** As discussed above in Ground 1, McGovern discloses deriving an observation that includes ***a duration of detected malicious activity associated with the technical assets.*** McGovern at ¶38 (“average response or patch time after security breach detection”). As discussed above, a POSITA would have found it obvious to combine Cole with McGovern, and in particular, by using Cole’s disclosed “components for tracking detected vulnerabilities, and for performing targeted vulnerability fix verification tests of specific host computers” to determine the duration of when the malicious activity was first detected to when it was verified that the vulnerability was fixed. Cole at 70:54-58.

4. 1[d] automatically generating a cyber-security rating for each of the entities using the entity map and the derived observations

226. With respect to the limitation ***“automatically generating a cyber-security rating for each of the entities using the entity map and the derived observations,”*** Cole discloses that its testing system features “an application of a quantitative score to objectively and comparatively rank the security of the target network.” 2:35-37. For example, Cole’s system discloses steps for derive an objective security score (***cyber-security rating***) for a network including assigning “vulnerability

risk levels ... to each of a plurality of vulnerabilities found on the network” and calculate such an objective security score “that depends at least on the vulnerability risk levels of the vulnerabilities, the vulnerability risk levels of the nodes, and the number of nodes on the network.” Cole at 10:33-41. Cole’s system further discloses a scoring routine that “establishes a vulnerability score for each target computer and for the network based on the results of the active assessment and based on the vulnerability information in the known vulnerability database.” *Id.* at 18:1-5.

227. Cole in particular describes a “FoundScore F [that] provides a quantitative score between 0 and 100 representing an objective indicia of the relative security of the target network” where “the higher the Foundscore ..., the greater the relative security of the target network.” *Id.* at 62:54-58. Cole describes that this Foundscore is determined using the vulnerabilities identified by the FASL scripts. *Id.* at 41:26-31.

5. 1[e] through a communication network, exposing to users in a user interface of a portal, information about the cyber-security ratings of the entities that is useful to decisions about entities with which to do business

228. With respect to the limitation “*through a communication network, exposing to users in a user interface of a portal, information about the cyber-security ratings of the entities,*” Cole’s testing system discloses “compilation of a dynamic, graphical report representing the network topology, network computers, and network vulnerabilities in a hierarchical report including both overview and detail documents”

and demonstrated several examples of such report in Figures 10-34. In particular, Cole discloses that an “exemplary FoundScore Report is set forth below in FIGS. 27-28.”





FIG. 27



FoundScore: Vulnerabilities	Your Results	Deductions	Running Score
High Risk Vulnerabilities	<u>13</u>	50	50
Medium Risk Vulnerabilities	<u>7</u>	0	50
Low Risk Vulnerabilities	<u>3</u>	0	50
Score after Vulnerability Deductions:			50

FoundScore: Exposure	Your Results	Deductions	Running Score
Number of Non-Essential Services	<u>150</u>	20	30
Number of Machines with No Essential Services	<u>9</u>	9	21
UDP Permitted	<u>Yes</u>	10	11
ICMP Permitted	<u>Yes</u>	5	6
Total Score:			6
			Poor

To understand what systems and ports affected your FoundScore, click on the specific item of interest in the "Your Results" column in the above table.

FIG. 28

229. Cole’s testing system further discloses that, for example, Figure 10 “illustrates an embodiment of a hierarchical security report 1000 *presented on a display for viewing* by a user” and a user “may select any of the reports shown in FIG. 10 (or other reports that are not shown) and request additional details.” *Id.* at 69:29-41.

230. Alternatively, a POSITA would have found it obvious that such report can also be accessed through a web-based user interface, because Cole’s testing system also discloses a user interface which “may be a web-based interface that is accessible to network users via an ordinary web browser” for the functionality of the ticket management component. *Id.* at 71:50-59. A POSITA would have been motivated to use that web-based interface to allow viewing of the vulnerability report by any location.

231. In addition, as discussed above in Ground 1, McGovern describes displaying an ISA score through a remote client, and it would have been obvious to combine Cole and McGovern.

232. With respect to the limitation “*information about the cyber-security ratings that is useful to decisions about entities with which to do business,*” Cole describes that the vulnerability reports provide “(1) a dynamic graphical display representing the network topology, the target computers found, and the vulnerabilities found throughout the target network; (2) a comprehensive list of target computers,

vulnerabilities found, and vulnerability explanations; (3) an objective scoring report describing the approximate security rating of the target network; (4) an exposure report describing the ports, services, and (5) detailed information that sets forth testing results on a per-machine, per-port, or per-vulnerability basis.” Cole at 68:63-69:4. As shown in, for example, Figure 28, each report is generated for a particular *entity*.

233. A POSITA would understand that all of this information in the report is *useful to decisions about entities with which to do business* because it would be more desirable to do business with an entity with a better security score and less desirable to do business with a worse security score. This is because the report’s scores are indicative of the security of the *entity’s network*, which is important if, for example, doing business with that entity involves transferring confidential or private information.

C. Claim 2: The method of claim 1 in which the technical assets comprise an IP space.

234. Cole discloses that “target computers” can be “identified by a unique or temporarily unique IP (Internet Protocol) address, typically in the form A.B.C.D, where each of A, B, C and D represent the Class A, Class B, Class C and Class D sub-networks and each has a value between 0 and 255.” 3:41-46. Further, Cole discloses that while “creating a topological representation of a network” the system can identify responsive computers on the network and obtain “*plurality of sequences of IP addresses* by sending to each responsive computer a sequence of packets having

increasing TTL values.” 9:46-50. A POSITA would understand that a sequence of IP addresses is part of an IP space for reasons explained above in Ground 1.

D. Claim 3: The method of claim 1 in which maintaining an entity map comprises querying a regional Internet registry.

235. As discussed in Ground 1, McNab discloses this limitation. It would have been obvious to incorporate such teachings into Cole.

E. Claim 4: The method of claim 3 in which maintaining an entity map comprises submitting a name of one of the entities to the regional Internet registry.

236. As discussed in Ground 1, McNab discloses this limitation. It would have been obvious to incorporate such teachings into Cole.

F. Claim 5: The method of claim 4 in which submitting the entity name includes submitting spelling variations of the entity name.

237. As discussed in Ground 1, McNab renders obvious this limitation. It would have been obvious to incorporate such teachings into Cole.

G. Claim 6: The method of claim 1 in which maintaining an entity map comprises using a domain name associated with the entity.

238. As discussed in Ground 1, McNab discloses this limitation. It would have been obvious to incorporate such teachings into Cole.

H. Claim 7: The method of claim 6 in which maintaining an entity map comprises using a dig tool on the domain name to determine a published IP address.

239. As discussed in Ground 1, McNab discloses this limitation. It would have been obvious to incorporate such teachings into Cole.

I. Claim 8: The method of claim 1 in which maintaining an entity map comprises probing IP addresses around a published IP address.

240. As discussed in Ground 1, McNab discloses this limitation. It would have been obvious to incorporate such teachings into Cole.

J. Claim 10: The method of claim 1 in which maintaining an entity map comprises using a name server lookup to verify the derived server name.

241. As discussed in Ground 1, McNab discloses this limitation. It would have been obvious to incorporate such teachings into Cole.

K. Claim 11: The method of claim 2 comprising using the IP space for collecting the cyber-security characterizations.

242. As discussed in Ground 1, McNab discloses this limitation. It would have been obvious to incorporate such teachings into Cole.

L. Claim 12 The method of claim 1 in which the technical assets that are associated with cyber-security characterizations of the respective entities comprise technical assets that the entity owns, controls, uses, or is affiliated with

243. Cole discloses that the target network itself “is typically defined as an interconnected set of computers, bounded by a specific pre-designated sub-network address, range of IP addresses or sub-addresses” and “defined by one or more *ranges of IP addresses controlled by the target network*, but may contain additional target computers or target sub-networks connected to the target network topographically but not part of the predetermined IP range or ranges.” Cole at 2:54-59, 3:46-51.

M. Claim 13: The method of claim 1 in which the technical assets comprise ranges of IP addresses

244. Cole discloses that “target computers” can be “identified by a unique or temporarily unique IP (Internet Protocol) address, typically in the form A.B.C.D, where each of A, B, C and D represent the Class A, Class B, Class C and Class D sub-networks and each has a value between 0 and 255.” 3:41-46. Further, Cole discloses that while “creating a topological representation of a network” the system can identify responsive computers on the network and obtain “*plurality of sequences of IP addresses* by sending to each responsive computer a sequence of packets having increasing TTL values.” 9:46-50.

N. Claim 14: The method of claim 1 in which the technical assets comprise servers

245. Cole discloses that the target network “typically includes one or more computers, where a computer includes a desktop station running any operating system, a router, a *server*, and/or any other networked device capable of sending and receiving packets through standard internet protocols.” 2:44-50.

O. Claim 15: The method of claim 1 in which maintaining an entity map comprises using information received from Internet Service Providers.

246. As discussed in Ground 1, McNab discloses this limitation. It would have been obvious to incorporate such teachings into Cole.

P. Claim 16: The method of claim 1 in which maintaining an entity map comprises using heuristics

247. The '524 patent explains that “heuristics” include at least: “Using the ‘dig’ ... tool to determine any IP information published by an entity” or “us[ing] the IP addresses and domain names published to find ranges of IP addresses actually used” or attempting to guess “servers whose names may be guessed” such as adding the “mail,” “smtp” or “ftp” prefix to a domain name. EX1001, 9:16-38.

248. As discussed above for limitation 1[a], McNab discloses adding the mail prefix to a domain name to guess server names. Thus, McNab discloses *maintaining an entity map by using heuristics*, and it would be obvious to incorporate such a teaching in Cole.

Q. Claim 18: The method of claim 1, wherein annotating comprises prepending to the domain name at least one of ‘mail,’ ‘smtp,’ and ‘ftp’

249. As discussed in Ground 1, McNab discloses this limitation. It would have been obvious to incorporate such teachings into Cole.

R. Claim 19

- 1. 19[pre], 19[a]: A computer implemented method comprising: on one or more computers operating with a database, maintaining an entity map that maps technical assets that comprise an IP space to respective entities with which the assets are associated,**

250. Limitation 1[pre] recites “[a] computer implemented method.”

251. Limitation 1[a] recites “on one or more computers” and “a database.”

252. Limitation 1[a] also recites “maintaining ... an entity map that maps technical assets to respective ... entities with which the assets are associated.”

253. Claim 2 recites that the technical assets “comprise an IP space.”

254. Thus, in my opinion, Cole in view of McNab and McGovern renders obvious these limitations for the same reasons I gave for limitations 1[pre], 1[a] and claim 2.

2. **19[b]: the maintaining of the entity map comprising (a) automatically using sensors on the Internet to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities, (b) associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name, (c) submitting names of the entities to regional Internet registries using various spellings of the names, (d) using domain names associated with the entities, (e) probing IP addresses around published IP addresses, (f) using a name server lookup to verify IP addresses associated with domain names, (g) using information received from Internet Service Providers, (h) annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses, and (i) applying heuristics,**

255. Limitation 1[b] recites “automatically using sensors on the Internet to collect externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities,” which is identical to 19[b]’s step (a).

256. Limitation 1[b] also recites “associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name” which is identical to 19[b]’s step (b).

257. Claim 4 recites “submitting a name of one of the entities to the regional Internet registry” and Claim 5 recites “submitting spelling variations of the entity name,” which is substantively identical to 19[b]’s step (c).

258. Claim 6 recites “using a domain name associated with the entity,” which is substantively identical to 19[b]’s step (d).

259. Claim 8 recites “probing IP addresses around a published IP address” which is substantively identical to 19[b]’s step (e).

260. Claim 10 recites “using a name server lookup to verify the derived server name,” which is substantively identical to 19[b]’s step (f).

261. Claim 15 recites “using information received from Internet Service Providers,” which is substantively identical to 19[b]’s step (g).

262. Limitation 1[a] recites “annotating a domain name with naming conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses,” which is substantively identical to 19[b]’s step (h).

263. Claim 18 recites “using heuristics” which is substantively identical to 19[b]’s step (i).

264. Thus, limitation 19[b] is rendered obvious by Cole in view of McNab and McGovern for the reasons explained above for limitations 1[a] and 1[b] and claims 4, 6, 8, 10, 15, and 18.

3. **19[c]: automatically deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (a) a number of technical assets that have been reported to be malicious and (b) a duration of detected malicious activity associated with the technical assets,**

265. In my opinion, this limitation is rendered obvious by Cole in view of McNab and McGovern for the same reasons I explained above for limitation 1[c].

4. **19[d]: automatically generating a cyber-security rating for each of the entities using the entity map and the derived observations, and**

266. In my opinion, this limitation is rendered obvious by Cole in view of McNab and McGovern for the same reasons I explained above for limitation 1[d].

5. **19[e]: through a communication network, exposing to users in a user interface the cyber-security ratings of the entities.**

267. In my opinion, this limitation is rendered obvious by Cole in view of McNab and McGovern for the same reasons I explained above for limitation 1[e].

S. Claim 20

1. **20[pre], 20[a], 20[b]: A system comprising: one or more computers programmed to facilitate operations comprising: maintaining in a database an entity map that maps technical assets to respective companies or other entities with which the assets are associated, wherein mapping technical assets to respective companies or other entities comprises determining IP information by annotating a domain name with naming**

conventions for at least one of a mail server, an SMTP server, and an FTP server, thereby deriving a likely server name having assigned IP addresses,

268. In my opinion, these limitations is rendered obvious by Cole in view of McNab and McGovern for the same reasons I explained above for limitation 1[pre] and 1[a].

2. **20[c]: collecting, using sensors on the Internet, externally observable cyber-security characterizations of the technical assets that have been mapped to each of the entities and associating the observable cyber-security characterizations with a respective entity based at least in part on the derived server name,**

269. In my opinion, this limitation is rendered obvious by Cole in view of McNab and McGovern for the same reasons I explained above for limitation 1[b].

3. **20[d]: deriving observations about the technical assets from the collected cyber-security characterizations, wherein the derived observations comprise (i) a number of technical assets that have been reported to be malicious and (ii) a duration of detected malicious activity associated with the technical assets,**

270. In my opinion, this limitation is rendered obvious by Cole in view of McNab and McGovern for the same reasons I explained above for limitation 1[c].

4. **20[e]: generating a cyber-security rating for each of the entities using the entity map and the derived observations, and**

271. In my opinion, this limitation is rendered obvious by Cole in view of McNab and McGovern for the same reasons I explained above for limitation 1[d].

5. **20[f]: through a communication network, exposing to users in a user interface of a portal, information about the cybersecurity ratings of the entities that is useful to decisions about entities with which to do business.**

272. In my opinion, this limitation is rendered obvious by Cole in view of McNab and McGovern for the same reasons I explained above for limitation 1[e].

IX. CONCLUSION

273. 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.