

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SOPHOS LTD., ET AL.,
Petitioners

v.

OPEN TEXT INC.
Patent Owner

Case IPR2023-00491
Patent No. 9,413,721

**DECLARATION OF PROFESSOR ALESSANDRO ORSO IN SUPPORT
OF PATENT OWNER**

Mail Stop PATENT BOARD
Patent Trial and Appeal Board
U.S. Patent & Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

Declaration in Support of Patent Owner’s
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721

TABLE OF CONTENTS

EXHIBIT LIST	iii
I. QUALIFICATIONS	2
II. UNDERSTANDING OF RELEVANT LEGAL PRINCIPLES	6
III. STATE OF THE PRIOR ART	9
IV. THE ’721 Patent	11
A. Summary of the ’721 Patent.....	11
B. Claim Construction	15
C. Level of Ordinary Skill in the Art.....	16
V. Overview of the Prior Art.....	18
A. Overview Of Liu	18
B. Overview Of Kester.....	21
C. Overview Of Tuvell.....	27
VI. Ground 1 Does Not Disclose or Render Obvious the Challenged Claims....	32
A. The Prior Art Does Not Disclose or Render Obvious the “increasing/increase a count” Limitation	33
1. Neither Liu Nor Tuvell Disclose “[increasing/increase] a count”	33
2. No Motivation to Modify Liu in view of Tuvell	35
3. No Reasonable Expectation of Success in Modifying Liu in View of Tuvell	37
B. The Prior Art in Ground 1 Does Not Disclose or Render Obvious The “central server” Limitation.....	38
C. The Petition Fails To Establish That Kester Is Analogous Art.....	43

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721

1.	Kester is not in the same field of endeavor.....	44
2.	Kester is not reasonably pertinent to the '721 patent	46
VII.	DEPENDENT CLAIMS.....	49
VIII.	OTHER LIMITATIONS NOT ADDRESSED	49
IX.	SECONDARY OBJECTIVE INDICIA OF NONOBVIOUSNESS	49
X.	Conclusion	50

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721

EXHIBIT LIST

Exhibit	Description
2001	Declaration of Professor Alessandro Orso, Ph. D.
2002	<i>Curriculum Vitae</i> of Professor Alessandro Orso, Ph. D.
2003	<i>Webroot, Inc., and Open Text Inc., v. CrowdStrike, Inc. and CrowdStrike Holdings, Inc.</i> , Case No. 6:22-cv-00241-ADA-DTG, First Amended Complaint for Patent Infringement, Dkt. 60 (W.D. Tex., Sept. 8, 2022) (excerpted)
2004	<i>Sonrai Memory Ltd. v. Kingston Tech. Co. and Kingston Tech. Corp.</i> , Case No. 6:21-cv-1284-ADA, Order Denying Kingston Tech. Co. and Kingston Tech. Corp.'s Opposed Motion to Stay Pending Resolution of <i>Inter Partes</i> Review, Dkt. 94 (W.D. Tex., Oct. 18, 2022)
2005	<i>mCom IP, LLC v. Cisco Systems, Inc.</i> , Case No. 6:22-cv-00261-ADA, Order Denying Defendant's Motion to Stay Pending <i>Inter Partes</i> Review, Dkt. 42 (W.D. Tex., Oct. 20, 2022)
2006	<i>Intellectual Ventures I LLC et al. v. Hewlett Packard Enterprise Co.</i> , No. 6:21-cv-00596-ADA, Order Denying Hewlett Packard Enterprise Company's Opposed Motion to Stay Pending Resolution of <i>Inter Partes</i> Review, Dkt. 104 (W.D. Tex., Dec. 22, 2022)
2007	<i>Webroot, Inc. and Open Text Inc., v. AO Kaspersky Lab. et al.</i> , Case No. 6:22-cv-243-ADA-DTG, Claim Construction Order, Dkt. 236 (W.D. Tex. Mar. 16, 2023)
2008	<i>Webroot, Inc. and Open Text Inc., v. AO Kaspersky Lab et al.</i> , Case No. 6:22-cv-00243-ADA-DTG, Scheduling Order, Dkt. 50 (W.D. Tex., Sept. 28, 2022)
2009	<i>Open Text Inc. et al., v. AO Kaspersky Lab</i> , Case No. 6:22-cv-00243-ADA-DTG, First Amended Complaint for Patent Infringement, Dkt. 80 (W.D. Tex., Oct. 27, 2022)

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721

Exhibit	Description
2010	<i>Webroot, Inc. and Open Text Inc., v. Sophos, Ltd.</i> , Case No. 6:22-cv-00240-ADA-DTG, Second Amended Complaint for Patent Infringement, Dkt. 76 (W.D. Tex., Oct. 24, 2022)
2011	Law360, “Catching Up on Patent Litigation with Judge Albright”
2012	June 30, 2022 U.S. Courts Statistics & Reports Comparison Report
2013	<i>Webroot, Inc. and Open Text Inc., v. ForcePoint LLC</i> , Case No. 6:22-cv-00342-ADA-DTG, Defendant ForcePoint LLC’s Preliminary Invalidity Contentions (W.D. Tex., Sept. 13, 2022)
2014	Application No. 15/285,875, Appeal No. 2021-003606, Decision on Appeal at 7-8 (PTAB July 26, 2022)

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. Patent No. 9,413,721

I, Professor Alessandro Orso, declare as follows:

1. I have been engaged by King & Spalding LLP on behalf of OpenText Inc. ("Patent Owner") to provide this Declaration concerning the technical subject matter relevant to the petition for Inter Partes Review ("Petition") of U.S. Patent No. 9,413,721 ("the '721 patent").

2. I am over 18 years of age. I have personal knowledge of the facts stated in this Declaration and could testify competently to them if asked to do so.

3. I have reviewed and am familiar with the specification and the claims of the '721 patent. In general, I will cite to the specification of a United States patent using the following formats: (Exhibit No., Col: Line Number(s)) or (Exhibit No., Paragraph Number(s)). For example, the citation (EX-1001, 1:1–10) points to the '721 Patent specification at column 1, lines 1–10. Also, for convenience, I use italics to denote limitations from the challenged claims.

4. All the opinions in this Declaration are based on the documents I reviewed and my knowledge and professional judgment. In forming the opinions expressed in this Declaration, I reviewed the Petition for *Inter Partes* Review of U.S. Patent No. 9,413,721 and supporting exhibits.

5. I have been asked to provide my technical opinions regarding how a person of ordinary skill in the art would have understood the claims of the '721 patent at the time of the alleged invention. I have also been asked to provide my

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721

technical opinions on how concepts in the '721 patent specification relate to claim limitations of the '721 patent. In reaching the opinions provided herein, I have considered the '721 patent, its prosecution history, and the references cited above and have drawn as appropriate on my own education, training, research, knowledge, and personal and professional experience.

I. QUALIFICATIONS

6. In formulating my opinions, I have relied on my knowledge, training, and experience in the relevant field, which I will summarize briefly. A more detailed summary of my background, experience, and publications I set forth in my curriculum vitae (CV), which I submitted as Exhibit 2002.

7. I am currently a Professor in the School of Computer Science and an Associate Dean in the College of Computing, at the Georgia Institute of Technology (Georgia Tech). I received my MS degree in Electrical Engineering (1995) and my PhD in Computer Science (1999) from Politecnico di Milano, Italy. My main areas of expertise are program analysis, software testing and debugging, computer security, and program comprehension.

8. Overall, I have been working in the software field for almost 30 years. While preparing my MS thesis, in 1994, I worked for one year as the system administrator for a network of 16 public libraries, which involved setting up,

connecting, and maintaining a group of Linux servers that were exchanging updates overnight.

9. After getting my MS degree, in 1995, I worked for several months as a free-lance consultant, helping develop IT solutions for companies and performing Linux system installations.

10. Additionally, in my initial years at Georgia Tech, when I was working as a Postdoctoral Fellow, I was in charge of setting up my research group's lab, which involved installing operating systems and software, setting up networks, and configuring routers and firewalls for these networks. In the same period, I became a certified Linux system administrator.

11. During my PhD, I did research in the general area of software testing. My dissertation was about novel techniques for testing systems written using object-oriented languages. After completing my PhD, in 1999, and before moving to Georgia Tech, I worked for a year as a Postdoctoral Fellow at Politecnico di Milano, where I extended my dissertation work. In the same year, I spent the summer at the University of Illinois at Chicago as a visiting researcher.

12. I joined Georgia Tech in 2000 and worked there in different capacities: Postdoctoral Fellow (2000-2001), Research Scientist (2001-2004), Assistant Professor (2004-2008), Associate Professor with tenure (2008-2014), and Full Professor (2014 to date). From January 2015 to December 2017, I was Associate

School Chair, which involved various kinds of administrative responsibilities. I became an Associate Dean in 2020.

13. In 2011, during my sabbatical year, I spent time as a visiting scholar at Politecnico di Milano, Italy (my alma mater), IBM Research, India, University of Lugano, Switzerland, and University of Buenos Aires, Argentina.

14. I have published over 100 refereed conference and journal papers in the general area of computer science. I have received several Distinguished Paper Awards and Best Paper Awards. I also received four Impact Awards, which are given to recognize the impact of research work in retrospect, typically considering the last ten years. I have been invited to give numerous guest lectures, talks, and distinguished lectures about my research. I have received grants to support my research from several major companies and national agencies, including Google, IBM, Microsoft, DARPA, DHS, Facebook, NSF, and ONR.

15. I have developed and taught several classes on software engineering, for both graduate and undergraduate students. I also developed an on-line class on software engineering that is being offered through the Udacity platform, as part of Georgia Tech's Online MSCS, and as part of Georgia Tech's professional education curriculum. My online class has been taken by thousands of students and has been used by several colleagues as teaching material for their software engineering courses.

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. Patent No. 9,413,721

16. In addition to teaching in the academic context, I also taught a variety of professional education courses in various topics, including software quality, project management, software design, and system administration, to IT professionals in various companies and countries, including Argentina, Italy, Paraguay, South Africa, and United States.

17. I am an expert in the area of computer security. In the last 20 years, I have received multiple grants to support my security-related work from various agencies, including DARPA, DHS, NSF, and ONR, and I have published several well-cited papers about computer security. For example, one line of research that I successfully pursued in my career focuses on analysis and testing of applications developed using web technology. In particular, I have developed two new techniques that can detect and prevent injection attacks to web applications. Together the three main papers that describe these techniques have been cited over 2,000 times by the scientific community. One of these papers has also received a Most Influential Paper Award at the IEEE/ACM International Conference on Automated Software Engineering (ASE 2020). Developing these kinds of techniques requires a deep understanding of the concepts, mechanisms, and technologies involved, including web and network security, multi-tier architectures, databases, JavaScript, SQL, HTML, HTTP and other underlying network protocols.

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. Patent No. 9,413,721

18. I am active in the software engineering research community in a variety of ways. I served as Program Chair for all the top conferences in my area and regularly serve on Program Boards and Program Committees for these conferences. I also regularly perform peer reviewing for various journals and for grant proposals submitted for funding to the National Science Foundation and other agencies (e.g., DARPA). I have been a member of the Steering Committee of the flagship conference in software engineering. I have also served on the Steering Committee for the two main conferences in the area of Software Testing and Program Analysis: The ACM International Symposium on Software Testing and Analysis (ISSTA) and the IEEE International Conference on Software Testing, Verification and Validation (ICST). Finally, I served as Associate Editor of the ACM Transactions on Software Engineering and Methodology and am currently serving as Associate Editor of the IEEE Transaction on Software Engineering.

19. I am a Distinguished Member of the ACM and an IEEE Fellow.

II. UNDERSTANDING OF RELEVANT LEGAL PRINCIPLES

20. I am not a lawyer, and I will not provide any legal opinions. Although I am not a lawyer, I have been advised certain legal standards are to be applied by technical experts in forming opinions regarding the meaning and validity of patent claims.

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. Patent No. 9,413,721

21. I understand that a patent claim is invalid if it is anticipated or obvious in view of the prior art, and that a claim can be unpatentable even if all of the requirements of the claim cannot be found in a single prior-art reference. I further understand that invalidity of a claim requires that the claim be anticipated or obvious from the perspective of a person of ordinary skill in the art at the time the invention was made.

22. I have been informed that a patent claim is invalid if it would have been obvious to a person of ordinary skill in the art. In analyzing the obviousness of a claim, I understand the following factors may be taken into account: (1) the scope and content of the prior art; (2) the differences between the prior art and the claims; (3) the level of ordinary skill in the art; and (4) any so called "secondary considerations" of nonobviousness if they are present.

23. I have been informed that a claim can be obvious in light of a single prior art reference or multiple prior art references. I understand that to prove that prior art or a combination of prior art renders a patent obvious, it is necessary to:

1. Identify the particular references that, singly or in combination, make the patent obvious;
2. Specifically identify which elements of the patent claim appear in each of the asserted references; and

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. Patent No. 9,413,721

3. Explain why a person of ordinary skill in the art would have combined

the references, and how they would have done so, to create the inventions claimed in the patent. I further understand that exemplary rationales that may support a conclusion of obviousness include:

- Combining prior art elements according to known methods to yield predictable results;
- Simple substitution of known element for another to obtain predictable results;
- Use of known technique(s) to improve similar devices (methods or products) in the same way;
- Applying a known technique to a known device (method or product) ready for improvement to yield predictable results;
- “Obvious to try” – choosing from a finite number of identified, predictable solutions with a reasonable expectation of success;
- Known work in one field of endeavor may prompt variations of the work for use in either the same field or a different field based on design incentives or other market forces if the variations are predictable to a person of ordinary skill in the art; and
- Some teaching, suggestion, or motivation in the prior art that would have led a person of ordinary skill in the art to modify the prior art

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721
reference or to combine the prior art reference teachings to arrive at
the claimed invention.

24. I have been informed that, in considering obviousness, hindsight reasoning derived from the patent-at-issue may not be used.

III. STATE OF THE PRIOR ART

25. “Malware” generally refers to an executable computer file or an object that is or contains malicious code, such as trojans, worms, and the like. (EX-1001 1:33-35.) Different versions of malware can function in different manners, and what works as a detection for one type of malware (*e.g.*, a virus) may be less effective against another type (*e.g.*, a worm). (*See generally* EX-1001 1:33 – 2:42.)

26. The '721 patent explains that prior anti-malware products used signature matching to detect malware, either locally or at a central server. (EX-1001 1:37-2:14.) The local anti-malware products suffered from delays in identifying new malware threats and obtaining signatures for them so they could be blocked. (*Id.* 1:37-55.) Central servers stored signatures in the cloud but only signatures or very basic information were sent to the central server for matching. (*Id.* 1:56-2:2.) If the object was unknown, a copy had to be sent to the central server for investigation by a human, a time-consuming and laborious task. (*Id.* 2:5-7.) In a network environment, it was unrealistic for a human to investigate each new object due to the high number of such objects that can be seen across a network on a daily basis.

(*Id.* at 2:7-10) Thus, under these approaches, “malevolent objects may escape investigation and detection for considerable periods of time.” (*Id.* at 2:10-13.)

27. Also, prior methods of classifying malware had technical drawbacks when used on a distributed network. For example, a distributed network that required each server to maintain rules for determining what is malware required each server to deal with huge amounts of largely common data. (EX-1001 12:20-24.) It was also generally impractical to store the required data on each server because, for example, there were problems determining whether or not the data—which is both massive and constantly changing—is common and up to date in real-time. (*Id.* at 12:24-27.) Accordingly, there was a long-felt need for an improved technique to efficiently determine whether files were safe or malware. The '721 patent recognized the need for efficiency given that the “large amounts of commonality in the data, which can be used for determining rules on what objects are malware.” (*Id.* at 12:20–23.) To solve this problem, the '721 explicitly rejected prior art approaches involving simply duplicating the data across all web servers. (*Id.* at 12:24 – 25 (“[I]t is impractical to store all data on all web servers 61 on all regions.”).)

28. The novel three-tiered architectures described and claimed in embodiments of the '721 patent provide an effective solution for distributed computer networks by, *inter alia*, reducing the workload across the network. (*Id.* at 12:28-59.)

29. A peer-to-peer network is a decentralized network in which there is no central server, and the computers in the network can connect directly to each other and act as both clients and servers. The computer in the network can share files and other resources directly, without the need for a central server that acts as shared storage. By contrast, client-server networks are centralized networks with a central server in which clients and servers are distinguished. In client-server networks, clients initiate connections to the servers and request data or services, and servers send the clients the results of the processing of such requests.

IV. THE '721 Patent

A. Summary of the '721 Patent

30. The '721 patent is entitled "Methods and Apparatus for Dealing with Malware," was filed on February 13, 2012, and was duly and legally issued by the United States Patent Office on February 5, 2013. The systems and methods described and claimed in the '721 patent are directed to improved techniques for detecting and classifying malware.

31. To address the shortcomings of the prior art, the '721 patent describes and claims unconventional, novel distributed system architectures, such as remote computers that may be allocated to "threat" servers, with "central" servers sitting behind the "threat servers". (EX-1001 9:16-57.) These enhanced computer architectures provide a technical solution to the technical problem of detecting and

classifying malware in a computer network environment, thus improving network security while identifying and classifying malware threats in real-time without delays engendered by use of human analysts. (EX-1001 1:60-2:7.)

32. In particular, the '721 patent describes and claims embodiments including three-tiered architectures of remote computers, threat servers, and a central server that provide a technical enhancement to the computer network itself (improving upon the two-tiered architectures of traditional systems having only remote computers and a central server) by enabling the central server to keep, for example, a master list “of all data objects, their metadata and behaviour seen on all of the remote computers,” and propagate it back to the threat servers. (EX-1001, 12:28-54.) This novel network architecture improves the operation of the computer network over traditional networks because, for example and as described in the '721 patent, “[t]his scheme has been found to reduce workload and traffic in the network by a factor of about 50 compared with a conventional scheme.” (Id. 12:55-57.) An example of this novel network architecture is illustrated in the figure below, which includes a base computer 3 (having Fx Layer 60, and ENZO layer 70) in communication with agent program 10 via the internet through DNS Distribution 50. (Id. 9:13–59.)




Fig. 3

(EX-1001, Fig. 3, 9:13-59.)

33. “[B]y being able to query and analyze the collective view of an object, i.e., its metadata and behaviours, across all agents [] that have seen it, a more informed view can be derived, whether by human or computer, of the object. In addition, it is possible to cross-group objects based on any of their criteria, i.e., the metadata and behaviour.” (Id. 18:17-22.) Thus, systems and methods disclosed in the ’721 patent enable better malware identification than conventional systems (e.g., using human analysis) in addition to providing an efficiency benefit. The

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721

patent explicitly notes that “the work in processing the raw data [] is too large of a task to be practical for a human operator to complete.” (Id. 18:50-52.)

34. Further, the '721 patent implements its server and database architecture using cloud-computing, which provides location-independent computing and is not tethered to a particular physical infrastructure. (EX-1001 10:31-39) (“In an exemplary embodiment, some or all of the servers 61, 74 and databases 62, 71, 72, 73 of the FX and ENZO layers 60, 70 are implemented using cloud computing. Cloud computing is a means of providing location independent computing, whereby shared servers provide resources, software, and data to computers and other devices on demand.”) This feature “provides a large degree of flexibility for the user.” (EX-1001 10:42-43.)

35. The systems and methods described and claimed in the '721 patent provide further improvements. For example, the information collected at the central server may include additional information about the object being classified as well as a count associated with the number of times that the first computer object has been seen. (*Id.* at cl. 1.) As explained above, using information about the object (such as behavior information) being classified, embodiments described and claimed in the '721 patent provide an approach that is more effective than traditional code or signature matching techniques for classifying objects as malicious. (*Id.* at 1:52-2:14.)

36. Accordingly, the '721 patent discloses and claims, among other things, an unconventional technological solution to the inherently network-centric technical issue of identifying malware in computer systems. The solutions implemented by the '721 patent provide a specific and substantial improvement over prior malware classification systems, including by introducing novel computer network architecture elements combined in an unconventional manner. These approaches improve the function and working of malware detection services by, for example, utilizing multiple threat servers and central servers and performing the analysis and communication carried out by each type of server in an unconventional and efficient manner. These elements and their combination represent a marked improvement in the functioning of computer systems used to identify and detect malware in computer networks.

B. Claim Construction

37. Petitioner does not offer a construction for any claim term and purports to apply the ordinary and customary meaning of the claim terms as understood by a POSITA with the exception of the “wherein” clause of claims 1 and 8. (Petition 5-7.) Petitioner challenged Patent Owner’s plain and ordinary meaning interpretation of these terms in the district court, but argues that the Board should “apply PO’s interpretation for purposes of this proceeding.” (Petition 7.) Patent Owner maintains that a plain and ordinary meaning interpretation of these terms is

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721

correct. For purposes of responding to the petition, Patent Owner applies the plain and ordinary meaning of these terms.

38. In the parallel district court litigation, the court construed the following term in its claim construction order (EX-2008 at 8), which I've applied in preparing my declaration:

Claim Term	District Court Construction
"object behavior information" (cls. 1, 9, 14)	Plain-and-ordinary meaning wherein the plain-and-ordinary meaning does not include the static attributes of the object

C. Level of Ordinary Skill in the Art

39. I understand that certain issues relating to the validity of the '721 patent must be judged from the perspective of a person of ordinary skill in the relevant art, as I discuss below. I have been asked to define the level of a "person having ordinary skill in the art" or "POSITA" that someone would have had at the time the alleged invention as claimed was made. In deciding the level of a POSITA, I have considered the following factors, which I have been informed are relevant in determining the level of ordinary skill in the art:

- The levels of education and experience of persons working in the field;
- The types of problems encountered in the art;
- The prior art solutions to these problems;

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. Patent No. 9,413,721

- The rapidity with which innovations are made; and
- The sophistication of the technology.

40. Based on these considerations, it is my opinion that a POSITA at the time of the invention of the '721 patent had a bachelor's degree in an accredited program of electrical engineering, computer engineering, or computer science, or a similar discipline, and have 2-3 years of practical work experience in the general fields of electrical engineering, computer science, networking, communications, and/or device and network security. More advanced degrees and/or training in a related discipline can compensate for shorter work experience.

41. I understand that Petitioner has defined the level of ordinary skill in the art as:

[I]n February 2012 would have had at least a Bachelor's degree in Computer Science, Computer Engineering, Electrical Engineering, or an equivalent discipline, and at least two years' worth of experience developing computer networking or network architecture, or a Master's degree in a related field.

EX-1003 ¶17. Alternately, a POSITA could have four or more years of work experience in computer networking and in the development of client-server systems.

(Petition 9.)

42. The minor variations in these two positions would not alter my opinions expressed in this Declaration.

43. Based on either standard, I was at least a person of ordinary skill in the art at the time of the invention of the '721 patent.

V. Overview of the Prior Art

A. Overview Of Liu


44. Liu is directed to improving the detection of and protection against malware by networked computers. (EX-1005 [0002]-[0004].) To identify threats, prior art systems and methods used threat signatures, which are “characteristic[s] of a threat that is unique and, therefore, distinguishes the threat from other potentially benign files or computer programs.” (EX-1005 [0004].) However, “[a] limitation of systems that use threat signatures to detect threats is that these systems do not, typically, possess a threat signature for a previously unknown threat.” (EX-1005 [0004].)

45. To overcome this limitation, Liu does not use a cloud-based distributed architecture but instead implements a “secure peer-to-peer network [that] is used to rapidly distribute information concerning the nature of the potential threat through the threat protection network.” (EX-1005 Abstract.) Liu discloses embodiments including “a client computer connected to a network, a server that stores threat definition data and is connected to the network, [and] an expert system in

communication with the server” to which unidentified threats are sent for analysis

to determine whether the potential threat is an actual threat. (EX-1005 [0006].)

This architecture is demonstrated schematically by Fig. 1:



(EX-1005 Fig. 1 [0066]-[0070]). Figure 1 shows threat protection network 10, including computer 12 connected via network 14, an alert server 16, threat identification server 18, and isolated computers 20, which are isolated from other network components. (EX-1005 [0066].) Liu discloses using alert server 16 and computers 12 to “form a secure peer-to-peer network that facilitates dissemination of the update.” (EX-1005 [0070].)

46. Computers 12 run client software and, if a potential threat is identified, computer 12 will ask alert server 16 for the information it has, if any, about

the potential threat. (EX-1005 [0067].) “[A]lert server 16 can possess information concerning the nature of a potential threat and provide information concerning the nature of the threat directly to the computer 12 that referred the potential threat to the alert server.” (EX-1005 [0068].) However, if alert server 16 does not possess information about the threat, the alert server 16 provides the potential threats to the threat identification server 18, which exposes the isolated computer 20 to the potential threat. (EX-1005 [0068].) Based on the potential threat’s behavior on computer 20, the threat identification server 18 can make a determination about whether the potential threat is an actual threat using a “threat score.” (EX-1005 [0048-50], [0068], [0091].) Depending on whether the potential threat is determined to be an actual threat, alert server 16 can generate an update for distribution to computers 12. (EX-1005 [0068].)

47. Liu benefits from the additional security provided by a peer-to-peer network, which has the effect of isolating the client computers from communications or potential threats outside the peer-to-peer network. (EX-1005 Fig. 7; [0066] (“Each of the isolated computers 20 ... can be isolated (i.e., prevented from sending information) to the broader network.”).) Thus, Liu teaches against the combination of its architecture with a non-peer-to-peer network, lest it lose the benefits from the secured communications architecture disclosed by Liu. (*See* EX-1005 at [0088] (“Once the copy of the threat is received, the threat

identification system is isolated (134) from the network. In several embodiments, isolation is achieved by preventing the threat identification system from sending messages using the network. In one embodiment, the threat identification system is prevented from sending email by implementing a socket layer filter. In other embodiments, isolation is complete and involves the threat identification system neither sending nor receiving messages via the network.”.)

B. Overview Of Kester

48. Kester discloses a software management system that uses a human administrator to determine how to categorize applications and how categorized applications should be distributed to computers on a network. (EX-1007 [0063]-[0064], Fig. 13.) Kester recognized that the rise in Internet connectivity meant an increase in employee communications and with that, an increased risk of employees installing applications or transmitting programs or files their employers may disapprove of. (EX-1007 [0005].) For example, instant messaging, while improving productivity, could pose a risk to sensitive data. (*Id.*) Media streaming, while useful, could drain network bandwidth. (*Id.*) Then-existing traditional filter systems managed how employees access the internet in the workplace, “preserving employee productivity, conserving network bandwidth and storage costs, limiting legal liabilities and improving network security,” but the growth of new threats

meant “employers need[ed] new solutions to manage the broader intersection of employees with their computing environments.” (EX-1007 [0006]-[0007].)

49. Thus, Kester sought to provide a means for “controlling application files,” *i.e.*, controlling when a user (or workstation) was permitted to launch an application. (EX-1007 Abstract, [0003], [0031].) Kester discloses a method and system that is capable of controlling a wide variety of application categories, as well as more than forty different exemplary categories, including those related to productivity (*e.g.*, project managers, word processing), communication (*e.g.*, email, instant messaging), audio/video (*e.g.*, media players), and entertainment (*e.g.*, games). (EX-1007 Fig. 4, [0058]) While one of these forty-four categories is “malicious applets and scripts,” Kester does not discuss why something should be categorized as malicious, how something gets categorized as “malicious” and/or what policy should be implemented for malicious applets or scripts. Kester’s method was not directed to the determination of whether a file was safe or malware, but rather to controlling application files and applying policies defined for these applications. (EX-1007 [0008]–[0010]; [0058] (“For each one of the parent groups and/or categories, the network administrator can select an individual policy or rule to associate therewith. Thus, once the requested application is categorized, the application server module 102 can select the policy or rule that is associated with that category.”).)

50. Because an employer's policies regarding applications and/or their network access are often user- or workstation-specific, Kester discloses a workstation management module ("WMM") that can "select the access privilege for the workstation and/or user" to control the launch of an application. (EX-1007 Figs. 2, 3, [0038], [0048] ("FIG. 3 is a block diagram of the application server module which communicates with the workstation management module 200 (FIG. 2) to upload and download a list of applications comprising properties of applications as well as policies associated with the applications once categorized."), [0050] ("The network administrator, or the like, interfaces with the application server module 102 via the classification user interface 106. The network administrator can classify uncategorized applications from the application inventory database 103 via the classification user interface 106. The network administrator can further interface through the classification user interface 106 to select or create access privileges/policies/rules for users, workstations, and/or groups of user/workstations".))

51. Figure 1 illustrates workstations 1 through N on a local area network 100, each of which includes a WMM:




FIG. 1

(EX-1007 Fig. 1.) Each WMM includes a hash/policy table, in which a hash value corresponding to an application (or user) is optionally associated with a category and also associated with a policy specific to that application/user. (EX-1007 [0031], Fig. 2, [0077]; [0037] (“Once the application that is being requested to run on the workstation is identified, the policy from the hash/policy table 204 which corresponds to that application is also identified.”); [0040] (“If the execution launch detection module 201 does not find the application hash in the hash/policy table 204 (for example, the application is uncategorized), the execution launch detection module 210 then determines how to proceed with the uncategorized application.”).)

52. When the WMM detects a request to launch an application, a “hash” or “application digest” is created from properties unique to the application requesting access. (EX-1007 [0074]-[0075], Fig. 7 (steps 702, 704), [0033].) The WMM then compares the “application digest” with the “hash/policy table.” (EX-1007 [0053], [0075], Fig. 7 (step 706).) “If the hash corresponds to a hash stored in the hash/policy table 204, the process continues to a state 710 where the policy associated with the hash is applied in response to the launch of the requested application.” (EX-1007 [0076], Fig. 7 (step 710).)

53. Kester teaches that a human administrator must review applications and determine the applicable category and policy for each application request and/or workstation. (EX-1007 [0050] (“The network administrator, or the like, interfaces with the application server module 102 via the classification user interface 106. The network administrator can classify uncategorized applications from the application inventory database 103 via the classification user interface 106.”).) Kester further teaches that a human reviewer interfaces with an “application analyst’s classification module” to determine these associations. (EX-1007 [0064] (“The application analyst’s classification module 302 classifies or categorizes applications which are then added to the master application database 300 of categorized applications. A human reviewer interacts with the application analyst’s classification module 302 to perform the categorization or recategorization.”), [0065]

(“For a human reviewer, a set of PC-based software tools can enable the human reviewer to manipulate, scrutinize, and otherwise manage the applications from the master application database 300.”), [0068], Fig. 10, Fig. 13.)

54. Kester first discloses that the human reviewer is provided with a list of applications and determines the appropriate category or categories for each. (*Id.* at [0064] (“A human reviewer interacts with the application analyst’s classification module 302 to perform the categorization or recategorization.”).) A list of applications is extracted from the “uncategorized application database 303 for classification by the human reviewer via the application analyst’s classification module 302.” (EX-1007 [0109].) Kester teaches that the human reviewer conducts an in-depth analysis of the application, related information, internet searches, and/or related documentation to determine the appropriate category(ies) for each application. (EX-1007 [0068]; [0109]-[0110], Fig. 13; *see also* [0008]- [0010].) As discussed above, this could include one or more of the forty-four exemplary categories Kester discloses. (EX-1007 Fig. 4.)

55. These “categories can be further associated with the policies or rules for the workstation 101 and/or user,” including allow, disallow, or an alternative. (EX-1007 [0031], [0037]-[0038].) “[T]he same application may be allowed to run on a workstation but not allowed to run on a different workstation.” (EX-1007 [0091]; [0039] (“For example, if application of the rule by the execution launch

module 210 indicates that the requested application is not allowed to run on the workstation 101 or to be run by the user, a predefined block page can be sent to the user interface explaining that the requested application is not allowed to run and why. Alternatively, the execution launch detection module 210 simply stops the requested application from running on the workstation 101.”) “Thus, once [a] requested application is categorized, the application server module 102 can select the policy or rule that is associated with that category.” (EX-1007 [0058].)

56. Accordingly, Kester enabled a human to control access to applications on a workstation based on information about the application, the workstation, and/or the user. (EX-1007 [0003], [0007]–[0010].)

C. Overview Of Tuvell

57. Tuvell is titled “System and Method of Reporting and Visualizing Malware on Mobile Network” and claims a “network management system [that] monitors malware within a mobile network.” (EX-1006.) It is specifically directed toward mobile networks and mobile devices, given their unique considerations and requirements, including voice data quality requirements for mobile users. (EX-1006 [0004]–[0005].) Tuvell uses “[c]ollection agents [] distributed within a mobile network at various network locations or sites. The collection agents collect executable programs that are being transferred through various protocols, e.g., Bluetooth and WiFi, using both mobile stations and key communications

components in the network, e.g., a GGSN in a GSM network and a PDSN in a CDMA network.” (EX-1006 [0011].) By collecting data from various distributed locations, Tuvell increases the likelihood that a new malware sample is captured once it starts spreading. (EX-1006 [0011], [0073].) Tuvell’s network is illustrated in Figure 11:




FIG. 11

58. Figure 11 illustrates malware sample collection system 1100 for obtaining samples of executable code within a mobile network and sending those samples to a collection center. (EX-1006 [0073].) Collection agents are distributed within the mobile network at various sites to collect executables using key communication components and mobile stations. (EX-1006 [0073].) In operation,

malware-infected devices 1106 and 1004 send connection attempts via protocol

handles 1114a or 1114b. (EX-1006 [0074].) A collection agent 1102 accepts the incoming call attempts from malware infected devices 1104 and 1006 and forwards any transferred executables to sample collection center 1112. Calls from the collection agent 1102 may be switched through the provider's mobile network 302 using a wireless data connection 1108e.

59. A collection agent's process flow is illustrated by Fig. 12:




FIG. 12

(EX-1006 Fig. 12)

60. Since Tuvell is limited to mobile devices and mobile networks, after first determining whether the computer object is an executable, the next decision point 1206 is for Tuvell to determine whether the executable is for a mobile phone. (EX-1006 Fig. 12.) If it is not, Tuvell discards the object and the process flow ends. (EX-1006 Fig. 12, [0086].) Tuvell does not disclose or suggest the deployment of its systems or methods in the context of a non-mobile network. Although the “collection agent” in Tuvell may be implemented using a “general purpose computer” (*see id.* [0092]), Tuvell makes it clear that the “collection agent” is interposed within a mobile network (*see id.* [0077]) (“A collection agent 1102, 202 is a device which is placed at various points in the mobile network 302 in order to collect samples being transmitted over the network executables, wherein a sample is transmitted data containing executable code.”); (*see also id.* [0074]) (“collection agent 1102 accepts the incoming call attempts from the malware infected devices”), [0086] (collection agent “determine[s] if the executable is for a mobile device” and if “the sample is not targeted for a mobile device 308, the collection agent 1102, 202 disards [sic] the executable”), [0089] (“Collection agents 1102, 202 have the following general functionalities: ... checking 1206 if the executable is specifically for mobile devices 308 by looking at the executable file format, if it is not specifically for mobile devices 308, then ignoring the executable;”).

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. Patent No. 9,413,721

61. The collection agent then proceeds to determine whether this particular executable has been previously seen and sent to the collection center. (EX-1006 [0087].) If the collection agent previously saw the executable, the collection agent notifies the collection center that the malware is being seen and identified again; if not, it sends the sample executable to the collection center. (EX-1006 Fig. 12 [0088].) Tuvell's collection agent's count reflects the number of times that same collection agent has seen a particular executable, not the number of times the particular executable has been seen across an entire network. (EX-1006 [0089].)

62. Critically, "[t]he design, both hardware and software, of a collection agent ... depends on its location in the service provider's mobile network." (EX-1006 [0090].) This requirement is consistent with Tuvell's application-specific goal of improving the security of devices and networks on a mobile network and visualizing the propagation of malware within a mobile network. (EX-1006 [0009] ("The system and method provides a comprehensive means for collecting, reporting, and providing visual depictions of information regarding the propagation and effect of worms, viruses, and other malware on a network.")). This is unlike the '721 patent, which utilizes cloud computing and is not tethered to particular hardware resources. (EX-1001 10:31-36 ("In an exemplary embodiment, some or all of the servers 61, 74 and databases 62, 71, 72, 73 of the FX and ENZO layers 60, 70 are implemented using cloud computing. Cloud computing is a means of

providing location independent computing, whereby shared servers provide re-sources, software, and data to computers and other devices on demand.”); 10:39 – 41 (“Cloud computing users avoid capital expenditure on hardware, software, and services when they pay a provider only for what they use.”); *see also* 23:26–40 (“[T]he embodiments disclosed herein may be implemented as electronic hardware, computer software, or combinations of both.”).)

VI. Ground 1 Does Not Disclose or Render Obvious the Challenged Claims

Ground 1	Section 103(a)	Obvious over Liu in view of Tuvell and Kester	Claims 1, 8, 9, 14, and 15
-----------------	----------------	--	-------------------------------

63. Petitioner’s arguments based on Liu in combination with Tuvell and Kester fail because the petition fails to establish that Ground 1 discloses or renders obvious every limitation of the independent claims, including in particular the “increasing/increase a count” and “central server” limitations. Moreover, the petition attempts to combine incongruous references with starkly different architectures, process flows, and functionality that a POSITA would not be motivated to combine. For these reasons, and as further addressed herein, Petitioner has not shown that it is reasonably likely to prevail on the patentability of any of the five challenged claims.

A. The Prior Art Does Not Disclose or Render Obvious the “increasing/increase a count” Limitation

64. Petitioner relies on Liu and Tuvell as disclosing the limitation “[increasing/increase] a count” in claim elements 1(f) and 9(f). (Petition 46–49, 59.)

However, as addressed below, neither Liu nor Tuvell discloses this limitation.

Limitation	Claim Language
1(e)	<i>“increasing a count associated with a number of times that the first computer object has been seen, and providing the increased count associated with the number of times that the first computer object has been seen to the central server”</i>
9(f)	<i>“increase a count associated with a number of times that the first computer object has been seen; the central server arranged to receive the increased count associated with the number of times that the first computer object has been seen”</i>

1. Neither Liu Nor Tuvell Disclose “[increasing/increase] a count”

65. Instead of relying on any disclosure from Liu, Petitioner relies on speculation about what Liu may do. Petitioner argues that Liu’s disclosure that a server “can provide in a return code ‘whether the potential threat has previously been reported by another client’” suggests that “Liu may do this by keeping a count of a number of times that a computer object has been seen.” (Petition 47)

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721

(emphasis added). Petitioner's use of "may" is instructive—Petitioner identifies no disclosure from Liu suggesting the use of a count and does not argue that this is inherently disclosed. (See Petition 47–48; EX-1005.) Rather, Petitioner speculates on what Liu "may" do to argue that it would have been obvious to use options like "SNMP MIBs," but does not identify any supporting disclosure from the Asserted Art discussing SNMP MIBs. (Petition 47-48.) Instead, Petitioner relies on its expert, but Dr. Newman also does not identify a single citation—from Liu, Kester, Tuvell, or any other reference—to support its SNMP MIB argument. (EX-1003 ¶ 134.) The Board should reject this unsupported argument. Petitioner alternatively cites Tuvell as purportedly disclosing the "increasing a count" limitations because Tuvell's collection agent notifies the sample collection center of the number of times that the collection agent has seen the executable. (Petition 47.) Petitioner argues that "[a] POSITA would have known that notifying the collection center of the number of times the collection agent has seen the executable necessarily requires increasing a count associated with a number of times that the first computer object has been seen." (Petition 47–48.) But this is a static count reflecting "the number of times the collection agent has seen the executable." (*Id.*) The Petition does not argue for or identify any actual disclosure showing increasing any count reflecting the number of times a computer object has been seen.

2. No Motivation to Modify Liu in view of Tuvell

A POSITA would not have been motivated to modify Liu in view of Tuvell to practice the “count” limitation of the ’721 patent because, unlike the ’721 patent, Tuvell has specific hardware requirements for its malware collection agent, *i.e.*, along the communication protocols used by the mobile network and distinct from the mobile device. (See EX-1006 [0011][0076–0080]; *compare* EX-1001, 10:31-36 (“In an exemplary embodiment, some or all of the servers ... and databases ... are implemented using cloud computing. Cloud computing is a means of providing location independent computing, whereby shared servers provide resources, software, and data to computers and other devices on demand.”) *with* (EX-1006 [0090]) (“The design, both hardware and software, of a collection agent 1102, 202 depends on its location in the service provider’s mobile network.”) Meanwhile, in Liu, the user computer is the same as the collection agent that identifies and analyzes malware in the first instance. (EX-1005 [0067]–[0071].) This is illustrated in the side-by-side comparison of figures from Liu and Tuvell, wherein the hardware agent is annotated

3. No Reasonable Expectation of Success in Modifying Liu in View of Tuvell

67. Petitioner suggests a POSITA would have had a reasonable chance of success in modifying Liu in view of Tuvell “because Tuvell teaches different information that can be collected about potential malware threats or malware objects and because including such a count was common in malware detection systems.” (Petition 49.) But this argument ignores the incongruous architectures of Tuvell and Liu. While Liu performs malware analysis by isolating test computers in a peer-to-peer network in order to run and analyze malware, (EX-1005 [0066]), Tuvell detects malware through signature matching of files using collection agents that “collect executable programs that are being transferred through various protocols, e.g., Bluetooth and Wi-Fi, using both mobile stations and key communication components in the network, e.g., a GGSN in a GSM and a PDSN in a CDMA network.”) (EX-1006 [0011].) These starkly different network topologies cannot simply be modified by swapping out components of the communication architecture, as Petitioner suggests. For example, there is no role or place for a collection agent within the peer-to-peer network communication of Liu. As such, and there is no likelihood of success in modifying Liu to meet this limitation by introducing Tuvell’s collection agent, which includes the (static) count reflecting the number of times a collection agent has seen a computer object. (*See* Petition 47.)

68. Petitioner does not reconcile these incompatibilities, thereby failing to show that a POSITA would have had a reasonable chance of success in modifying either Tuvell or Liu to practice the “count” limitation of the ’721 patent.

B. The Prior Art in Ground 1 Does Not Disclose or Render Obvious The “central server” Limitation

69. Petitioner relies on Liu alone as disclosing the limitation “central server” in claim elements 1(g) and 9(g). (Petition 49–51, 59.) However, as addressed below, Liu does not disclose this limitation.

Limitation	Claim Language
1(g)	<i>“receiving, at a second threat server, at least a portion of the contents of the at least one database associated with the central server, wherein the at least a portion of the contents of the at least one database associated with the central server include a subset of the details of the first computer object stored in the second database.”</i>
9(g)	<i>“a second threat server arranged to receive at least a portion of the contents of the at least one database associated with the central server, wherein the at least a portion of the contents of the at least one database associated with the central server include a</i>

	<i>subset of the details of the first computer object stored in the second database"</i>
--	--

70. The Petition relies solely on Liu's disclosure of "mirror servers" for both the "central server limitations" and the "first threat server" and "second threat server" limitations in the challenged '721 patent claims. (Petition 50.) But as explained below, unlike the challenged claims – which expressly distinguish between the central server and the first and second threat servers – Liu's mirror servers are effectively identical to each other. Thus, even if Liu's mirror servers could be interpreted as teaching the claimed threat servers—which Patent Owner disputes—there is no "central server" in Liu's peer-to-peer network of mirrored servers. Moreover, the '721 explicitly teaches away from the mirror server system of Liu. (EX-12:10–12) ("[T]his scheme allows a simple, agent managed way of dynamically load balancing between serves and providing resiliency"); 12:20–27 ("As will be clear from the foregoing, the web servers 61 in a typical implementation will be dealing with huge amounts of data. There will be large amounts of commonality in the data, which can be used for determining rules on what objects are malware. However, it is impractical to store all web data on all web servers 61 in all regions. This presents a problem in determining whether or not the data is common in real time."); 12:55–59 ("This scheme has been found to reduce workload and traffic in the network by a factor of about 50 compared with a conventional scheme where

typically each piece of data on being received by a web server would immediately be propagated to all other servers in the system.”)

71. As the '721 patent specification makes clear, the claimed “central server” collects information from multiple threat servers and serves as the central repository for information about existing and potential malware threats across a network. (EX-1001 9:53–59 (“The ENZO layer 70 has one or more servers 74 (referred to as “central servers” or “ENZO servers” in this disclosure) connected to one or more ‘master’ databases which aggregate all of the information receives from the web servers 61 from all the remote computers 2.”); 15:29-33 (“By collecting hashes, signatures, and statistics of new programs and modules correlated by path and file details, and aggregating this information, it is possible to use a method of analysis that will allow faster identification of new threats.”).) By contrast, threat servers are distributed throughout the network to collect, analyze, and transmit information about computer objects (and their safety) on particular remote computer(s). (*Compare* EX-1001 9:26–30 (“Each web server 61 of the FX layer 60 is connected to a database 62 which contains entries for all of the objects known to the base computer 3 and a classification of whether the object is safe or unsafe.”) with 9:47–52 (“Sitting behind the FX layer 60 is a second layer (called the ‘ENZO layer’ 70 herein). Information about objects received by the FX layer from the remote computers 2 is sent to the ENZO layer 70 such that a master record is

maintained at the ENZO layer 70 of all data received from all remote computers

2.”).)

72. Furthermore, Liu’s disclosure of mirror servers is limited. (*See* EX-1005 [0101].) The use of mirror servers is presented in the context of alternative embodiments in the last paragraph of the specification. (EX-1005 [0101].)

73. Petitioner’s attempt to equate Liu’s mirror servers with the claimed “central server” is wrong at least because it would render the “central server” superfluous and provides a system that suffers from the exact problems that the ’721 patent identified, improved on, and disparaged throughout the specification. For example, the ’721 patent discloses that the central server aggregates relevant data to avoid the problem inherent in Liu’s alleged network of mirror servers—the duplication of data across all servers on a network. (EX-1001 12:20–27 (“[I]t is impractical to store all data on all web servers 61 on all regions.”) 12:55–59 (“This scheme has been found to reduce workload and traffic in the network by a factor of about 50 compared with a conventional scheme where typically each piece of data on being received by a web server would immediately be propagated to all other servers in the system.”); *see* Petition 50.)

74. Petitioner’s mapping of Liu’s “numerous mirrored servers” to the claimed “central server” (Petition 50) results in the central server and every other threat server on the network containing identical information. But this conflicts

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. Patent No. 9,413,721

with the role of the central server in the '721 patent. Because the first, second, etc., threat servers would all contain identical information, there would be no reason for threat servers to pull or pass information to the central server, since all the threat servers already have all the same information—thereby obviating the central server.

75. Mirror servers that blindly duplicate data across all threat servers will drastically increase the volume of data transmissions in order to keep the network up to date with respect to its malware definitions, thus failing to meet the speed and efficiency goals of the '721 Patent. (EX-1001 10:13-17 (“As will be appreciated, in a community having for example 10 million or more remote computers 2, each having on average hundreds to thousands of objects, each having a large number of behaviours and associations with other objects, the amount of data held in the databases 71, 72, and 73 is enormous.”).)

76. Petitioner appears to contend that a POSITA would change the entire architecture of Liu from a peer-to-peer architecture to anyone of a “hub-and-spoke, peer-to-peer, or hybrid architecture.” (Petition 45.) First, Petitioner does not actually explain how it proposes to modify Liu, let alone how Lie would still function for its intended purpose in these hypothetical, undisclosed systems. Its arguments should therefore be rejected.

77. Second, Petitioner provides no motivation for modifying Liu to include an entirely different architecture beyond boiler plate statements that it is a “design choice” or provides efficiency generic benefits (cost, redundancy, efficiency, scalability, etc.) (Petition at 45.) Petitioner does not identify the a “hub-and-spoke” or “hybrid architecture” in the prior art with which it would propose to modify Liu, instead relying on pure hindsight to propose these changes. Petitioner even states, without explanation, that these changes would “not affect the system’s functionality.” (Petition at 45.) But Petitioner does not explain *how* would a central server be used in Liu, and *why* would a POSITA be motivated to add it to the system. Petitioner’s failure is significant here, where it proposes a radical change to a peer-to-peer network, which even its expert Dr. Newman recognizes is different from architectures that use a central server. (EX-1003 ¶129.) Thus, a POSITA would have not understood Liu’s “alert servers”—including a mirrored version of an “alert server” to be the claimed “central server.”

78. Accordingly, Ground 1 neither discloses nor renders obvious the “central server” limitation.

C. The Petition Fails To Establish That Kester Is Analogous Art

79. Petitioner relies on Kester as disclosing limitations in claims 1(b), 1(d), 1(e), 9(b), 9(d), and 9(e). However, Kester is non-analogous art to the ’721

patent, as it is not in the same field of endeavor nor reasonably pertinent to the problems considered by the inventors of the '721 patent.

80. The '721 patent explains that its particular field of endeavor is the rapid determination of whether an object is safe or malware. (EX-1001 1:48-55, 2:29 – 42, cls. 1, 9.) There is no dispute that the '721 patent sought to solve the problem of rapidly identifying and classifying computer objects as potential malware threats, and distributing updated information about threats as the network obtains new information about the potential threats it identifies. (*Supra* Section III.B; Petition 16-17). The '721 patent describes and claims embodiments including novel three-tiered architectures of remote computers, threat servers, and a central server that provide a technical enhancement to the computer network itself (improving upon the two-tiered architectures of traditional systems having only remote computers and a central server) by enabling the central server to keep, for example, a master list “of all data objects, their metadata and behaviour seen on all of the remote computers,” and propagate a subset of that information to other threat servers. (EX-1001 12:28-54.)

1. Kester is not in the same field of endeavor.

81. Kester is not in the same field of endeavor as the '721 patent. Kester was classified as “data processing” and explicitly defines its “[f]ield of [i]nvention” as “related to computing devices and, more particularly to monitoring and

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721

controlling application files operating thereon.” (EX-1007 [0002]-[0003], Title, cls. 1, 23, 37, 39.) Kester explains that while prior art software was available for “preserving employee productivity, conserving network bandwidth and storage costs, limiting legal liability and improving network security,” Kester sought to provide a “new solution[] to manage the broader intersection of employees with their computing environments” through monitoring and control of application files. (EX-1007 [0006], Title, [0003].) Kester’s field of endeavor is monitoring and controlling application files (EX-1007 [0002]-[0003]), which is distinct from the ’721 patent’s field of rapid determination of whether an object is malware. (*See supra* Section IV.A.)

82. Kester discloses forty-four exemplary application categories in Figure 4, only one of which is “malicious applets and scripts.” (EX-1007 Fig. 4, [0057].) But aside from this single reference, Kester contains no disclosure regarding when, why, or how an application should be categorized as “malicious applets or scripts,” no disclosure regarding what might distinguish an application as safe as opposed to a malicious applet, and no disclosure as to what policy should be applied to applications categorized as a malicious applet or script.

83. Petitioner wrongly asserts that Kester “describes protecting computers from malicious software programs,” “storing information identifying malicious programs in a database associated with a computer;” and “providing the contents

of that database to a central database factory.” (Petition 21-22). Petitioner’s argument that Kester is in the general field of classifying an object as malware is erroneous and based on cherry picked claim language that misrepresents Kester (Petition 22-23.) If one were to accept Petitioner’s assertion that Kester fits in the field of malware classification, it would mean that Kester is also in the field of “instant messaging” classification, “media player” classification, “gambling” classification, and more than forty other fields of endeavor despite Kester’s lack of disclosure as to when or why each of these categories should be applied. (See EX-1007 Fig. 4.) Further, as I explain above, a POSITA would not have looked to Kester for a determination of whether an object is safe or malware because Kester contains no disclosure regarding when, why or how an application should be categorized as “malicious applets or scripts,” no disclosure regarding what might distinguish an application as safe as opposed to a malicious applet, and no disclosure as to what policy should be applied to applications categorized as a malicious applet or script.

84. As Kester is not directed to a determination of whether an object is safe or malware, and certainly not a rapid determination thereof, Kester is in a different field of endeavor than that of the ’721 patent.

2. Kester is not reasonably pertinent to the ’721 patent

85. Kester is also not reasonably pertinent to the problem faced by the ’721 patent inventors. As discussed above, the ’721 patent sought to improve upon

time-consuming prior art methods for determining whether an object was malware or safe, and propagating updated information throughout a network about potentially newly-encountered objects. (EX-1001 2:4-27.)

86. Critically, the '721 specifically identifies a human-analyst based method—as disclosed in Kester—as undesirable, time consuming, and inaccurate. (EX-1001 2:2–14) (“Therefore, in order to analyse whether or not an unknown object is malware, a copy of that object is normally sent to the central server where it is investigated by a human analyst. This is a time consuming and laborious task introducing considerable delay in classifying malware as safe or unsafe. Also, given the considerable volume of new objects that can be seen daily across a community, it is unrealistic to have a skilled human analyst investigate each new object thoroughly. Accordingly, malevolent objects may escape investigation and detection for considerable periods of time during which they can carry out their malevolent activity in the community.”)

87. A POSITA would not have looked to Kester in considering the problems addressed in the '721 patent because Kester embodies the same problems the '721 patent sought to solve. (EX-1001 16:14-34 (“In detecting malware, human analysts and computer automated analysis have different strengths and are often used together to some degree to harness the capabilities of both In the present application, the amount of data received from agent program 10 running on remote

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. 9,413,721

computers 2 is vast and therefore suited to being processed by computer. It would be impractical to have human operators look at all of this data to determine malware.”) Kester’s ability to control application files requires that a human reviewer conducts a detailed analysis of each application (*e.g.*, by performing Internet searches, studying manuals, reading documents), decides how applications should be categorized, and makes policy decisions based on the specific user and/or workstation. (EX-1007 [0064], [0068], [0109]-[0110].) And Kester contains no disclosure regarding why or how an application would be determined to be malicious versus safe, so a POSITA would not have looked to Kester as reasonably pertinent to avoiding detailed analysis and need for decisions of a human analyst that slowed down prior art systems.

88. Accordingly, Kester is not reasonably pertinent to the particular problem that the inventor of the ’721 patent was addressing.

89. Based on the above, it is my opinion that Kester is neither in the same field of endeavor nor reasonably pertinent to the particular problem addressed by the inventor of the ’721 patent. Further, it is my opinion that Kester is non-analogous art to the ’721 patent and cannot serve as a reference for Petitioner’s Section 103(a) grounds.

VII. DEPENDENT CLAIMS

90. Because Petitioner has failed to establish that any of its grounds render obvious all limitations of independent claims 1 and 9, I understand that Petitioner cannot make a showing that such grounds render obvious dependent claims 8, 14, and 15.

VIII. OTHER LIMITATIONS NOT ADDRESSED

91. It is my understanding that this declaration is being provided in support of Patent Owner's Preliminary response. I have provided analysis of key limitations not included in the prior art. The fact that I have not analyzed particular limitations or claims in no way indicates or supports that these limitations or claims are disclosed or rendered obvious by the prior art, and I reserve the right to supplement my analysis if asked to do so.

IX. SECONDARY OBJECTIVE INDICIA OF NONOBVIOUSNESS

92. I understand that in determining obviousness of the claimed invention, objective evidence of secondary considerations pertaining to nonobviousness must be considered, including long-felt but unsolved needs, failure of others, commercial success, etc. Petitioner does not address or consider any such evidence in its petition. (*See generally* Petition.)

93. As discussed above, a long-felt and persistent need existed to improve network forensics methods before the invention of the '721 patent. *See* Section

III.A, *supra*. Prior attempts to develop reliable network security were limited in their ability to effectively and efficiently communicate large volumes of information about potential malware across a network and provide up-to-date, real-time malware protection to devices on the network, making them insufficient to identify and address the threats faced by modern computer networks. *See* Section III.A, *supra*.

94. These secondary indicia further support my opinion that the '721 patent was non-obvious.

X. Conclusion

95. For at least the foregoing reasons, it is my opinion that Petitioner's grounds do not render obvious any of the challenged claims of the '721 patent.

\\

\\

\\

\\

\\

\\

\\

\\

\\

Declaration in Support of Patent Owner's
Preliminary Response to Petition for *Inter Partes* Review
Patent No. Patent No. 9,413,721

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further 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 Title 18 of the United States Code.

Executed this 16th day of May, 2023 in Atlanta, Georgia.

Respectfully submitted,



Alessandro Orso, Ph.D.