

Filed on behalf of: Orca Security Ltd.

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

WIZ, INC.,
Petitioner,

v.

ORCA SECURITY LTD.,
Patent Owner.

Case IPR2024-01109
Patent 11,726,809

**DECLARATION OF DR. DAVID Kaeli
IN SUPPORT OF
PATENT OWNER'S RESPONSE**

Table of Contents

I.	Introduction	3
II.	Qualifications	4
III.	Summary Of Materials Reviewed And Considered	7
IV.	Legal Standards.....	8
V.	Level Of Skill In The Art And Perspective Applied In This Declaration	12
VI.	Background	15
A.	Overview of the '809 Patent.....	15
B.	Claim Construction.....	20
1.	“Analyzing the Snapshot”	20
VII.	Overview of the Asserted References.....	24
A.	Overview of Veselov	24
B.	Overview of Mohanty.....	32
C.	Overview of Czarny.....	35
D.	Overview of Hutchins.....	35
VIII.	The Claims Of The '809 Patent Would Not Have Been Obvious...	37
A.	Ground 1: Claims 1-10 and 12-23 Would Not Have Been Obvious Over Veselov and Mohanty	38
1.	A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov and Mohanty and Would Not Have Had a Reasonable Expectation of Success in Doing So	38
2.	Even If Combined, Veselov and Mohanty Do Not Disclose or Suggest Multiple Limitations in the Independent Claims	66
B.	Ground 2: Claims 1-10 and 12-23 Would Not Have Been Obvious Over Veselov, Mohanty, and Czarny.....	88

1.	A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov, Mohanty, and Czarny and Would Not Have Had a Reasonable Expectation of Success in Doing So	89
2.	Even If Combined, Veselov, Mohanty, and Czarny Do Not Render Claims 1-10 and 12-23 Obvious.....	91
C.	Ground 3: Claim 11 Would Not Have Been Obvious Over Veselov, Mohanty, and Hutchins	93
1.	A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov, Mohanty, and Hutchins and Would Not Have Had a Reasonable Expectation of Success in Doing So	93
2.	Even If Combined, Veselov, Mohanty, and Hutchins Do Not Render Claim 11 Obvious	95
D.	Ground 4: Claim 11 Would Not Have Been Obvious Over Veselov, Mohanty, Hutchins, and Czarny	97
1.	A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov, Mohanty, Czarny, and Hutchins and Would Not Have Had a Reasonable Expectation of Success in Doing So	97
2.	Even If Combined, Veselov, Mohanty, Czarny, and Hutchins Do Not Render Claim 11 Obvious.....	98
IX.	Objective Indicia of Non-obviousness.....	99
A.	Orca's SideScanning™ Uses The '809 Patent Inventions	99
B.	Long-Felt, Unresolved Need	100
C.	Orca's claimed inventions overcame industry skepticism to become widely accepted--and widely praised—in the industry.	105
D.	Extensive Industry Praise Supports Non-Obviousness of the Claims	109
E.	Evidence Shows Wiz Copied Orca's SideScanning™.....	116
1.	Evidence shows Wiz had access to Orca's invention	117
2.	Evidence shows Wiz's Workload Scanning is substantially similar to Orca's SideScanning™	134
F.	Commercial Success	157
X.	CONCLUSION	159

I. Introduction

1. I, David Kaeli, submit this declaration on behalf of Orca Security Ltd. (“Patent Owner”) to provide my opinions concerning the validity of claims 1-23 of U.S. Patent No. 11,726,809 (the “’809 patent,” Ex.1001), in support of Patent Owner’s Response to the Petition for *Inter Partes* review of the ’809 patent, which I understand was filed by Wiz, Inc. (“Petitioner”) and accompanied by a Declaration of Dr. Angelos Stavrou (the “Stavrou Declaration”), among other exhibits. I further understand that on January 16, 2025, this *Inter Partes* review (IPR) was instituted (Paper 9).

2. I have previously provided my opinions in a Declaration (Ex.2001) in support of the Patent Owner’s Preliminary Response to the Petition. Although I believe that declaration is accurate in all respects, this new Declaration represents my considered opinion and supersedes my previous Declaration.

3. I am more than eighteen years of age, and I am a citizen of the United States, currently residing in the Commonwealth of Massachusetts.

4. I am being compensated for the time I spent preparing this Declaration at my customary rate. The opinions I express in this Declaration are my own. My compensation does not depend upon the outcome of this IPR, the contents of this Declaration, or any testimony I may provide in this IPR.

5. This Declaration summarizes the opinions I have formed to date. I reserve the right to modify or update my opinions, if necessary, based on my review and analysis of information or materials that I may receive after this Declaration is submitted. I further reserve the right to respond to any new or updated positions Wiz or its expert(s) may take in this IPR.

II. Qualifications

6. In formulating my opinions, I have relied upon my knowledge, training, and experience in the relevant art. My qualifications, along with a more detailed summary of my background and publications, are stated more fully in my *curriculum vitae*, which has been provided as Exhibit 2002. Here, I provide a brief summary of my qualifications.

7. I am currently the Director of Northeastern University's Computer Architecture Research Laboratory and a Distinguished Professor of Electrical and Computer Engineering. I have been a faculty member at Northeastern since September 1993, in the Electrical and Computer Engineering Department.

8. For over thirty years, I have pursued research across a broad range of computer architectures and technologies, including memory systems, storage systems, virtualization, and secure computing systems. I have published over 400 refereed journal and conference papers in the areas of computer architecture and computer security. I also have supervised over fifty Ph.D. theses in these areas.

9. I am a Fellow of the Institute of Electrical and Electronic Engineers (IEEE) and a Fellow of the Association of Computing Machinery (ACM). I have served two terms as the Chair of the Executive Committee of the IEEE Technical Committee on Computer Architecture. I am serving in my second terms as the Editor-in-Chief of the ACM Transactions on Architecture and Code Optimization. I am also an Associate Editor for Elsevier's Journal of Future Generation Computer Systems.

10. From 2011-2014, I served as Chief Science Advisor of Apposable Software. From 2007-2011, I served as Chief Technology Officer of NUIC Technologies, where I worked on designing secure virtual appliances. From 2005-2011, I consulted for Akorri Networks, where, among other things, I helped develop intellectual property relating to storage virtualization. In 2000, I served as a consultant for InCert Corporation.

11. Before I joined Northeastern's faculty in 1993, I spent twelve years working for IBM. From 1981-1986, I worked on the design and delivery of the IBM System/370 3090 mainframe system. From 1986-1993, I worked at the IBM T.J. Watson Research Center, designing, and analyzing IBM Personal Systems 2 computers and associated software. During that time, I also worked on the design of a new Application System/400 multi-processor system.

12. I am an active educator and researcher and have extensive experience in the broader field of computer architecture. I am a co-inventor on patents targeting virtualization security. *See, e.g.,* Moffie, et al., *VMM-Based Intrusion Detection System*, (US Patent 8,719,936, May 6, 2014). I am a co-inventor on patents relating to virtual machine management. *See, e.g.,* Smirnov, et al., *Systems and Methods for Determining Placement of Virtual Machines*, (US Patent 8,099,487, January 17, 2012). And, I have conducted academic research on, and published multiple papers relating to, hardware and software security, including papers discussing security on virtualized infrastructure, computer system vulnerabilities and protections, and monitoring virtualized computing environments. *See, e.g.,* Azmandian et al., *Securing Virtualized Execution Environments Through Machine Learning-based Intrusion Detection*, IEEE International Workshop on Machine Learning for Signal Processing (2015); Azmandian et al., *Virtual Machine Monitor-based Lightweight Intrusion Detection*, ACM SIGOPS Operating Systems Review (2011); Sabbagh et al., *Overdrive Fault Attacks on GPU*, IEEE/IFIP International Conference on Dependable Systems (2021); Ahn et al., *Trident: A Hybrid Correlation-Collison GPU Cache Timing Attack for AES Key Recovery*, IEEE International Symposium on High Performance Computer Architecture (2024); Arzuaga, et al., *Increasing Power/Performance Resource Efficiency on Virtualized Enterprise Servers*, IEEE

Computing Frontiers (2011); Azmandian, et al., *Workload Characterization at the Virtualization Layer*, IEEE Annual International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, (2011).

13. I earned a Bachelor of Science degree in Electrical Engineering from Rutgers University in 1981. I earned a Master of Science degree in Computer Engineering from Syracuse University in 1985. I earned a Ph.D. in Electrical Engineering from Rutgers University in 1992.

III. Summary Of Materials Reviewed And Considered

14. The opinions contained in this Declaration are based on the documents I reviewed and my knowledge and professional judgment. In forming the opinions expressed in this Declaration, while drawing on my experience, I reviewed all of the documents cited in this declaration, as well as the Petition and the Institution Decision.

15. My opinions are additionally guided by my appreciation of how a person of ordinary skill in the art would have understood the claims of the '809 patent at the time of the invention, which I have been asked to assume is January 28, 2019. I will sometimes refer to this time as the "relevant time."

IV. Legal Standards

16. I am not an attorney. However, I have been informed by counsel of certain legal standards and have applied those standards, described below, in formulating my opinions.

17. I understand that the claims of a patent define the limits of a patentee's exclusive rights. I understand that to determine the scope of the claimed invention, courts may construe claim terms, the meaning of which the parties may dispute. I understand that claim terms should generally be given their ordinary and customary meaning as understood by a person of ordinary skill in the art at the time of filing of the patent application.

18. I understand that claims must be construed in light of and consistent with the intrinsic evidence. In this context, I understand that intrinsic evidence includes the claims themselves, the written disclosure in the specification, and the patent's prosecution history. I understand that the specification is always highly relevant to the claim construction analysis and often is the single best guide to the meaning of a disputed term. I understand that extrinsic evidence may also be considered when construing claims and may include, for example, technical dictionaries, technical publications and books, treatises, and expert testimony, though the intrinsic record will be more persuasive as to the meaning of a particular term.

19. I understand that one or more prior art references can render a patent claim obvious to a person of ordinary skill in the art if the differences between the subject matter set forth in the patent claim and the prior art are such that the subject matter of the claim would have been obvious at the time the claimed invention was made.

20. I understand that for a reference to qualify as a relevant prior art reference in an obvious analysis, the reference must be “analogous”—meaning it is from the same field of endeavor as the patented technology, and if not, then the reference must be reasonable pertinent to the particular problem facing the inventors of the patent at issue. I also understand that a reference is reasonably pertinent if a person of ordinary skill in the art would have reasonably consulted those references and applied their teachings in seeking a solution to the problem facing the inventors.

21. In analyzing obviousness, I understand that it is important to consider the scope of the claims, the level of skill of a person of ordinary skill in the relevant art, the scope and content of the prior art, and the differences between the prior art and the claims. I understand that claimed subject matter can be obvious when a person of ordinary skill in the art would have been motivated to combine pre-existing disclosures in the prior art, and that combination yields no more than what a person of ordinary skill in the art would have expected. However, I further

understand that obviousness cannot be based on the hindsight combination of components selectively culled from the prior art.

22. I also understand from counsel that to guide against hindsight, obviousness must be considered in view of any evidence demonstrating certain objective factors that may show that an invention was not obvious from the prior art, even if the differences between the prior art and the claimed invention seem slight. I have been informed by counsel that “objective indicia” of non-obviousness, can include, for example, evidence of commercial success of the invention, evidence of a long-felt need that was solved by an invention, evidence that others copied an invention, evidence that an invention achieved an unexpected result, evidence of industry praise of an invention, evidence of skepticism by others, evidence of licensing an invention, or evidence that others failed to solve the same problem that an invention solved.

23. I have been informed that in order to have relevance and weight in an obviousness analysis, the patentee must show that there is a “nexus” between the objective indicia and the merits of the claimed inventions. I understand that the focus of the nexus requirement looks to the patent claims as a whole, and that nexus does *not* have to be based only on limitations that are allegedly not shown in a primary prior art reference. For example, even if every limitation of a claim is shown to be *individually* known in the prior art, that does not preclude nexus where

the objective indicia evidence is tied to the combination of features as a whole and the combination was not previously known.

24. I have been informed that commercial success is relevant as an indication of nonobviousness because the law presumes an idea would successfully have been brought to market sooner, in response to market forces, had the idea been obvious to persons skilled in the art. Thus, the law deems evidence of (1) commercial success, and (2) some causal relation or “nexus” between an invention and commercial success of a product embodying that invention, helpful in determining whether an invention was nonobvious.

25. Similarly, I have been informed that recognition by others of skill in the art that the inventor’s contribution was an invention may support a finding of nonobviousness. This is because the law recognizes that industry participants, especially competitors, are not likely to praise an obvious advance over the known art. Accordingly, the fact that a product which embodies the patent claims garnered awards or otherwise received praise will weigh against an assertion that the same claim would have been obvious.

26. I understand that copying of the claimed inventions can also support an inference of nonobviousness, because it can both show a type of “praising” of the invention and suggesting that the prior art was insufficient. I have been informed that, like all other objective indicia, copying requires a showing of a

nexus between the alleged copying and the claimed invention. I also understand that not every competing product that arguably falls within the scope of a patent is evidence of copying, because otherwise every infringement suit would automatically confirm the nonobviousness of the patent. Rather, a nexus can be shown for copying based on, for example, that a competitor had access to the patentee's product or a disclosed embodiment in the patent and then created a product having substantial similarity to the patented product.

V. Level Of Skill In The Art And Perspective Applied In This Declaration

27. I understand that certain issues relating to validity must be judged from the perspective of a person of ordinary skill in the art. I further understand that factors that may be considered in determining the level of skill of a person of ordinary skill in the art include, but are not limited to, (1) the type of problems encountered in the art; (2) prior art solutions to those problems; (3) rapidity with which innovations are made; (4) sophistication of the technology; and (5) the educational level of active workers in the field.

28. I have been informed and understand that a person of ordinary skill in the art is a hypothetical person who is presumed to be aware of all pertinent prior art, thinks along conventional wisdom in the art, and is a person of ordinary creativity. Importantly, this hypothetical person is of ordinary skill in the art to which the claimed subject matter pertains.

29. With respect to the '809 patent, a person of ordinary skill in the art at the time of the earliest priority date, January 28, 2019, would have been familiar with techniques for securing virtual cloud assets, and underlying virtualization technologies used with such techniques.

30. A person of ordinary skill in the art would have had at least a Bachelor's degree in computer science, computer engineering, or a related field, and two years of industry experience or academic research experience in cyber security and virtualization, including in cloud computing cybersecurity. In my view, more education can supplement relevant experience and vice versa.

31. At the time of the earliest priority date for the '809 patent (January 28, 2019), I met the qualifications to be a person of ordinary skill in the art, and I believe that I also have a sufficient level of knowledge, experience, and education to provide expert opinions in the field of the '809 patent.

32. While my own level of skill exceeded that of the ordinary level of skill as of January 2019, I am well-acquainted with the actual performance and capabilities of a person of ordinary skill in the art as defined above. For example, during the relevant timeframe, I was supervising and evaluating the performance of such individuals, such as my graduate students who were performing research in this area. Furthermore, my participation in professional organizations and in

conferences and other events exposes me to the state of the industry and the level of experience of individuals who, like I did, had at least the ordinary level of skill.

33. I have applied the above-defined perspective of a person of ordinary skill in the art, as of January 2019, in forming and expressing my opinions. Unless otherwise stated, my testimony below refers to the knowledge of one of ordinary skill in the art at that relevant time.

34. In the Stavrou Declaration (Ex.1002), Dr. Stavrou sets out the following articulation of the level of ordinary skill:

In my opinion, a person of ordinary skill in the art (“POSA”) at the relevant time would have had (1) at least a bachelor’s degree in computer science, computer engineering, electrical engineering, or a related degree, and (2) would also have 2-3 years of professional experience with cyber security analysis and virtualization. Such experience includes, for example, malware analysis, security analysis of cloud computing systems, and security analysis of virtual machines. ... Additional experience can compensate for less education and vice versa.

Ex.1002 (Stavrou Decl.), ¶22. At the time of the earliest priority date for the ’809 patent (January 28, 2019), I met Dr. Stavrou’s proposed qualifications to be a person of ordinary skill in the art, and my opinions in this Declaration would

not change if the Board were to adopt and apply that proposed level of ordinary skill.

VI. Background

A. Overview of the '809 Patent

35. The '809 patent is titled “Techniques for Securing Virtual Machines by Application Existence Analysis” and generally relates to innovative systems and methods for securing virtual cloud assets, such as VMs. Ex.1001 ('809 patent), Abstract, cls. 1, 16, 19.

36. The '809 patent explains that, around the time of its priority date (January 2019), public cloud service providers, like Amazon, Microsoft, and Google, had “increasingly adapted their applications to be run from multiple cloud computing platforms,” in which “[v]irtualization,” which “allow[s] multiple applications and users to share the same cloud computing infrastructure,” “plays a key role.” Ex.1001 ('809 patent), 1:20-30. The '809 patent also provides context that “virtualization can be achieved by means of virtual machines,” each of which can “emulate[] a number of ‘computers’ or instances, all within a single physical device.” Ex.1001 ('809 patent), 1:31-33.

37. The '809 patent discusses problems that existed with trends in the cybersecurity field. For example, the '809 patent discusses that “applications and virtual machines running on top of virtualization technologies are also vulnerable

to some cyber threats,” e.g., from “vulnerable software applications or infected operating systems.” Ex.1001 (’809 patent), 1:42-46. And, the ’809 patent also explains ways the industry had tried to deal with such issues. For example, the ’809 patent discusses “traffic inspection,” which is “[t]raditionally ... performed by a network device connected between a client and a server ... hosting virtual machines”; using “network scanner[s] deployed out of path” to query servers to determine if applications possessing security threats are executed; “traffic monitor[ing],” which “listens to traffic flows between clients and [a] server”; and using “vulnerability management and security assessment solutions ... based on agents installed in each server in a cloud computing platform or data center.” Ex.1001 (’809 patent), 1:46-2:14. As the ’809 patent also explains, though, these methods suffered from disadvantages: “Traffic inspection may not provide an accurate indication of the security status of [a] server due to inherent limitations, such as encryption” and lack of access to pertinent data. Ex.1001 (’809 patent), 1:52-55. With network scanners deployed out of path, servers “may not respond” or “expose [] necessary data,” and network configuration- and credential-related issues can be problematic. Ex.1001 (’809 patent), 1:60-67. Traffic monitoring “detect[s] threats only based on the monitored traffic” and thus can miss server “misconfiguration[s]” and “vulnerabilities in software executed by the server.” Ex.1001 (’809 patent), 2:4-9. And “[u]sing agents is a cumbersome

solution for a number of reasons, including IT resource management, governance, and performance.” Ex.1001 ('809 patent), 2:14-16. “For example, installing agents in a large data center may take months.” Ex.1001 ('809 patent), 2:16-17. The '809 patent's systems and method “overcome the[se] deficiencies.” Ex.1001 ('809 patent), 2:18-19.

38. The '809 patent aims to solve such problems. Ex.1001 ('809 patent), 3:35-36, 3:47-63. Figure 1B of the '809 patent, reproduced below, illustrates a diagram of a cloud computing platform to which the '809 patent's techniques for scanning for embedded vulnerabilities in machines apply:




FIG. 1B

39. The cloud computing platform of the '809 patent includes a server (shown as no. 115 in Figure 1B) that hosts one or more VMs (shown as no. 119 in Figure 1B). Ex.1001 ('809 patent), 3:48-50. The cloud computing platform also has storage (shown as no. 117 in Figure 1B) which acts as storage space for the platform's server and "emulates" virtual disks (shown as no. 118-1 in Figure 1B) for the VMs. Ex.1001 ('809 patent), 3:50-51, 3:55-60. The storage may be part of the server or, as shown in Figure 1B, separated from the server. Ex.1001 ('809 patent), 3:55-60. In the example shown in Figure 1B, a "high-speed connection" such as an "optic fiber" connects the storage with the server. Ex.1001 ('809 patent), 3:55-60. Figure 1B also shows a client environment (shown as no. 130 in Figure 1B) which interfaces with a security system (shown as no. 140 in Figure 1B) using "APIs" or through "services provided by the cloud computing platform" (not shown in Figure 1B). Ex.1001 ('809 patent), 4:16-21. The security system is configured to perform the steps of the '809 patent's invention and may be separate from the client environment, as shown in Figure 1B, or part of the client environment. Ex.1001 ('809 patent), 4:1-15. The challenged claims of the '809 patent demonstrate the embodiment where the security system interfaces with the client environment through APIs. Ex.1001 ('809 patent), cls. 1, 16, 19.

40. The security system of the '809 patent interfaces with the client environment so that it can perform scanning operations and detect any "vulnerable

software, non-secure configurations, exploitation attempts, compromised assets, data leaks, data mining” and other risks or vulnerabilities. Ex.1001 (’809 patent), 4:50-56. For example, the security system can detect vulnerable code (Ex.1001 (’809 patent), 5:31-33), cyber threats that do not represent vulnerabilities such as unencrypted sensitive data (Ex.1001 (’809 patent), 6:21-32), and changes in sensitive machine areas such as “added or changed application files without installation” (Ex.1001 (’809 patent), 6:33-39). The security system is also configured to provide “incident response, anti-ransomware, and cyber-insurance” through access to the system’s “security posture” (Ex.1001 (’809 patent), 4:50-56), generate reports on detected vulnerabilities and cyber threats, including, “for example, as alerts” (Ex.1001 (’809 patent), 7:53-55), and mitigate detected vulnerabilities and cyber threats. Ex.1001 (’809 patent), 7:59-8:6. Mitigation actions may include, for example, blocking traffic from untrusted networks, quarantining an infected VM, or halting the operation of the VM. Ex.1001 (’809 patent), 8:1-6.

41. In order to accomplish the scanning and detection operations I discuss above, the ’809 patent provides a series of steps. The ’809 patent provides multiple techniques for analyzing snapshots to detect different types of vulnerabilities and cyber threats. Ex.1001 (’809 patent), 7:59-8:6. For example, the snapshot may be parsed and analyzed by the security system, the system may match installed

application lists to a known list of vulnerable applications, and the system may check configuration files of applications. Ex.1001 ('809 patent), 5:20-25, 5:34-36, 5:49-54. The security system then generates reports on detected vulnerabilities and cyber threats, including, “for example, as alerts” that may be filtered and prioritized. Ex.1001 ('809 patent), 7:53-55. The security systems and methods also “may cause an execution of one or more mitigation actions,” including disabling a VM from execution, updating a VM with patches, blocking traffic from untrusted networks, quarantining an infected VM, and preventing a VM from launching. Ex.1001 ('809 patent), 6:61-67, 8:1-6.

B. Claim Construction

42. In addition to the explicit constructions discussed below, I have applied the plain and ordinary meanings of the claim terms in the '809 patent, as a person of ordinary skill in the art would have understood them in the context of that patent in approximately January 2019.

1. “Analyzing the Snapshot”

43. Claims 1 and 16 recite “analyzing the snapshot of the at least one virtual disk by matching installed applications with applications on a known list of vulnerable applications.” Similarly, claim 19 recites “analyze the snapshot of the at least one virtual disk by matching installed applications with applications on a

known list of vulnerable applications.” Dr. Stavrou labels these as elements 1.3, 16.3, and 19.3. Ex.1002 (Stavrou Decl.), ¶148.

44. Dr. Stavrou states that he was “asked to assume that ‘[analyze/analyzing] the snapshot’ also encompasses analysis of a VM instantiated from the snapshot.” Ex.1002 (Stavrou Decl.), ¶82; *see also* ¶83 (“Accordingly, my discussion of Ground 1 below applies an understanding of ‘analyze the snapshot’ in claims 1 and 19, as well as ‘analyzing the snapshot’ in claim 16,¹ as encompassing both direct analysis of the snapshot data and analysis of a VM instantiated from snapshot.”). Thus, Dr. Stavrou treats limitations in the challenged claims that recite “[analyze/analyzing] the snapshot” as though they “encompass[] both [1] direct analysis of the snapshot data and [2] analysis of a VM instantiated from [the] snapshot.” Ex.1002 (Stavrou Decl.), ¶¶80-83.

45. I disagree that this claim element encompasses “[2] analysis of a VM instantiated from [the] snapshot” because that is not how a person of ordinary skill in the art would understand the claim element. Dr. Stavrou does not try to justify this claim element encompassing “analysis of a VM instantiated from a snapshot,”

¹ Contrary to this passage of Dr. Stavrou’s declaration, claims 1 and 16 recite “analyzing” and claim 19 recites “analyze.” However, for the purpose of claim construction the distinction does not matter.

but simply states that he was “asked to assume” this claim construction. Ex.1002 (Stavrou Decl.), ¶82; *see* Ex.2077 (1st Stavrou Tr.) 265:17-266:4, 266:13-267:15 (“I didn’t ask why. I was asked to assume.”), 267:16-268:3, 268:4-269:20. Dr. Stavrou also explained that he did not have an opinion on interpreting analyzing the snapshot to include analysis of a VM instantiated from a snapshot.

7 Q. Here you state, "A POSA would have
8 understood the ordinary and customary meaning of
9 '[analyze/analyzing] the snapshot' as at least
10 encompassing direct analysis of the snapshot data";
11 is that right?

12 A. That is correct.

13 Q. And your opinion is that that would be sort
14 of the correct meaning of the term
15 "analyze/analyzing the snapshot" in the context of
16 the '926 patent?

17 MR. ARGENTI: Objection. Form.

18 A. I believe that I -- they were consistent.
19 So I say, "This understanding is consistent with the
20 '926 specification."

Ex.2172 (2nd Stavrou Tr.) 38:7-20

4 Q. And I think in -- in your prior deposition
5 in the context of the '031, '032, and '685 patents,
6 you said that you were just asked to assume that by
7 counsel and you didn't form an opinion about whether
8 it's correct. Is that also true here?

9 A. I don't know if this is exactly what I said
10 in my prior deposition, but I was asked to assume
11 that analyzing and analyze the snapshot also
12 encompasses analysis of a VM instantiated from the
13 snapshot. That is correct.

14 Q. But you don't have an opinion as a technical
15 expert that that is a correct -- a correct way of
16 interpreting analyzing the snapshot consistent with
17 the '926 patent specification; right?

18 A. I was not asked to opine on that, so I don't
19 have an opinion.

Ex.2172 (2nd Stavrou Tr.) 39:4-40:1.

46. A person of ordinary skill in the art would not understand the “[analyze/analyzing] the snapshot” limitations in the '809 patent to encompass “analysis of a VM instantiated from a snapshot.” As an initial matter, there appears to be no disagreement that an instantiated VM is not a snapshot—otherwise Dr. Stavrou would not have enumerated it as separate from “direct analysis of snapshot data” in his assumed construction. Rather, an instantiated VM is a VM, and analysis of a VM instantiated from a snapshot is an analysis of a VM—not a snapshot. Once a VM is instantiated from a snapshot, e.g., when a duplicate VM is instantiated from a snapshot for subsequent monitoring, the instantiated VM is

no longer a “snapshot,” and monitoring, assessing, or evaluating it is not “analyze/analyzing” a “snapshot” as claimed. This would have been readily understood by a person of ordinary skill in the art reading the '809 patent—an instantiated VM is **not** a snapshot including because “[a] VM’s snapshot is a copy of the machine’s virtual disk (or disk file) *at a given point in time.*” Ex.1001 ('809 patent), 4:60-61. The specification also distinguishes analyzing a snapshot by a “security system 140 to detect vulnerabilities” (Ex.1001 ('809 patent), 5:20-21), as claimed, from “another embodiment” where a system “instantiate[s] a copy of the VM” and performs “execution of the instance of the VM... [in] an isolated sandbox” (Ex.1001 ('809 patent), 5:55-60, 7:42-52).

47. In fact, as I discuss further below, it appears that Dr. Stavrou “assume[d] that ‘[analyze/analyzing] the snapshot’ also encompasses analysis of a VM instantiated from the snapshot,” as he was “asked” to (Ex.1002 (Stavrou Decl.), ¶83), simply to try to buttress opinions he offers that depend on disclosures from prior art references, such as Veselov, that discuss evaluating duplicate VMs, not snapshots.

VII. Overview of the Asserted References

A. Overview of Veselov

48. Veselov is titled “Security Assessment of Virtual Computing Environment Using Logical Volume Image” and discloses systems and techniques

that perform a security assessment on “virtualized reproductions of computing resources” such as VMs or instances of VMs. Ex.1007 (Veselov), Title, Abstract, 3:20-23. Figure 1 illustrates a computing environment and components:




FIG. 1

49. As seen in Figure 1, Veselov discloses embodiments of a computing environment involving a computing resource service provider (shown as no. 104), in which user devices (shown as no. 102) request or manage resource allocation for server computers (shown as no. 142). Ex.1007 (Veselov), 4:19-29. These server computers are allocated as virtual computing resources. Ex.1007 (Veselov), 4:19-29. Figure 1 also depicts an API (shown as no. 120) that allows a user to

connect, through, for example, a web-based application, to the computing resource service provider. Ex.1007 (Veselov), 4:29-41.

50. Veselov states that the scanning service of Figure 1 implements various disclosed steps for a security assessment. Ex.1007 (Veselov), 8:16-17. The scanning service obtains and processes snapshot data to generate assessment results. Ex.1007 (Veselov), 8:18-21. In implementing its scanning procedures, Veselov attempts to address drawbacks with resource usage during the security assessments. For example, Veselov highlights that there is a problem in the cybersecurity field where security assessments can “reduce the target computing device’s capabilities” such that the target machine “may not be capable of responding to a client as quickly as normal” or “may not be capable of providing content as quickly as normal.” Ex.1007 (Veselov), 2:56-67. Veselov also cautions that security assessments may impact machine capabilities because system components such as processors and memories are used to perform the assessment rather than to execute normal workloads. Ex.1007 (Veselov), 2:56-67. And, Veselov notes that security assessments using a target device’s computing resources “may be costly monetarily” and will require “an extended period of downtime.” Ex.1007 (Veselov), 3:8-11.


51. Veselov’s overarching goal is therefore to avoid downtime of VMs and reduce assessment complexity when performing security assessments.

Throughout its disclosure, Veselov aims to preserve the activity of the original virtual machine instance. For example, Veselov provides systems that “include a scanning service that facilitates duplication of all or a portion of the target computing resource, and then perform[] the security assessment on the duplicate computing resource”—“to avoid consuming processing time, processing power, and storage space of the target computing resource.” Ex.1007 (Veselov), Abstract; *see also* Ex.1007 (Veselov), 10:9-11 (generating a scannable volume from a snapshot includes “launching a duplicate virtual machine instance.”). And, Veselov provides that performing a robust security assessment may reduce the target computing device’s capabilities and may not be as responsive should an assessment be performed on the target machine. Ex.1007 (Veselov), 2:56-67. And Veselov cautions against “downtime.” Ex.1007 (Veselov), 3:8-11. Veselov explains that its solution allows the assessment to “be performed without impacting the customer’s virtual machine,” which can continue to “servic[e] requests from its customers or perform[] other useful work.” Ex.1007 (Veselov), 12:21-43. For the system to perform such “other useful work” Veselov must keep its VMs active during the capture of a snapshot and subsequent vulnerability analysis. Additionally, Veselov describes solutions where the VM is configured to access one or more physical or virtual resources that are isolated in the event data dependencies need to be evaluated during an assessment. Ex.1007 (Veselov),

13:1-22, 13:33-36, 13:43-67. And the scanning service can cause “activity to occur within the assessment virtual machine” to apply tasks in a real-time environment. Ex.1007 (Veselov), 15:41-59, 20:57-21:13, 21:21-32.


52. Veselov proposes two scanning solutions to maintain the active state of the target virtual machine. In both of Veselov’s solutions, Veselov implements snapshot technology to capture the state of the target virtual machine. *See, e.g.*, Ex.1007 (Veselov), Fig. 1, no. 146 (snapshot data). Veselov implements a virtual layer to generate or obtain access to a snapshot or snapshot data to initiate a security assessment. Ex.1007 (Veselov), 3:23-32, 6:25-49. Veselov’s snapshot includes all data necessary to “recreate the state of the computing resource” within a duplicate resource so as to preserve the active state of the machine. Ex.1007 (Veselov), 3:27-29, 3:41-34, 12:21-27.

53. In Veselov’s first solution, Veselov’s scanning service sets up and launches an active virtual machine instance (a duplicate machine for assessments). Ex.1007 (Veselov), 3:54-59. The duplicate machine has the same configuration as the original machine instance at the time of the snapshot capture. Ex.1007 (Veselov), 3:54-59. Figures 3A and 3B are informative.



54. In Figure 3A, the scanning service performs the security assessment on the assessment machine via a security agent program. Ex.1007 (Veselov), 14:13-27. The scanning service receives assessment results via the agents on the assessment machine. Ex.1007 (Veselov), 14:28-34. The assessment machine can implement processing activity in order to demonstrate real-time dynamic or static data, network connections, and communications to other computing resources. Ex.1007 (Veselov), 15:41-59. Additionally, a user or administrator may communicate instructions to the scanning service (shown inside the computing resource service provider's environment as no. 110 in Figures 3A and 3B) via the API (shown as no. 338) to configure and accomplish the security assessment. Ex.1007 (Veselov), 14:36-15:40.

55. In Veselov's second solution, Veselov provides that the scanning system may "mount" the logical volume as a hard drive with data stored thereon or may transform the snapshot into files to be read. Ex.1007 (Veselov), 3:61-66. This technique permits a security assessment on the data at the file system level. Ex.1007 (Veselov), 3:66-67. Figures 5A and 5B are informative.



56. In exemplary Figure 5A, the scanning service analyzes snapshot data in a "test environment" outside of the virtual computing environment. Ex.1007 (Veselov), 17:9-12. In both Figures 5A and 5B, the user communicates directly

with the scanning service outside the target resource computing environment. Ex.1007 (Veselov), Figs. 5A-B. 17:9-67. An API is not depicted in Figures 5A and 5B, and is not used for communication between the user and scanning service. The scanning service interfaces with VM resources through provisioning services such as a hypervisor. Ex.1007 (Veselov), Figs. 5A-B. 17:9-67.

57. In both solutions, Veselov's scanning system stores the assessment results or provides them to "an owner or administrator of the virtual machine in order to evaluate the vulnerabilities of the target resources." Ex.1007 (Veselov), 4:5-8. After storing or sharing the results, the scanning system may delete the snapshot or store it so that "the user can perform additional analysis on the snapshot." Ex.1007 (Veselov), 4:8-10.

58. In both solutions, Veselov's snapshots are "prepared so that the security assessment produces the results that the scanning system otherwise would have obtained if the security assessment were performed on the original virtual machine instance." Ex.1007 (Veselov), 4:1-4. Accordingly, as detailed in this section, Veselov's techniques purport to preserve the activity of the original virtual machine instance and focus on analyzing deployed and executing virtual machines.


B. Overview of Mohanty

59. Mohanty is titled “Method and System to Prioritize Vulnerabilities Based On Contextual Correlation” and issued on June 27, 2017, from an application filed on November 11, 2014. Ex.1075 (Mohanty), Cover.

60. Mohanty includes embodiments for “prioritizing vulnerabilities of an asset in a virtual computing environment” by determining vulnerability scores via scanned vulnerability data. Ex.1075 (Mohanty), Abstract, Fig. 2. Mohanty does not teach or discuss snapshot based analyses. Instead, Mohanty implements agent-based scanning techniques to address what it indicates are drawbacks with security assessment techniques that provide vulnerability assessments of network systems in the past, such as snapshot-based systems. For example, Mohanty states that a “snapshot of a vulnerability assessment report of a system” will become “obsolete within hours or minutes as virtual machines or workloads change positions within a virtualized environment.” Ex.1075 (Mohanty), 1:14-17. Mohanty further advises that “security operations team[s] need[] a strong and continuous prioritization system to track critical vulnerabilities and take actions as changes occur” (Ex.1075 (Mohanty), 1:19-22) and highlights its techniques for “rapid[] updat[es].” Ex.1075 (Mohanty), 1:3-8. Accordingly, Mohanty relies on scanning technologies that “perform vulnerability checks (i.e., scans) **on** the virtual machines,” such as “Qualys” and “Rapid7” to perform vulnerability checks on

virtual machines and virtual applications to report vulnerability data. Ex.1075 (Mohanty), 6:54-58, Fig. 2. A POSA would have understood that, at the time, Qualys and Rapid7 were well known to have agent-based assessment tools. Ex.2173 (Qualys, *Vulnerability Management*); Ex.2174 (Man, *Live Monitoring with Endpoint Agents*). I understand Dr. Stavrou did not investigate or look at whether Mohanty, or the cited Qualys or Rapid7, used agent-based approaches. Ex.2172 (Stavrou 2nd Tr.), 239:3-240:9 (Stavrou “didn’t look at” whether Mohanty uses agent-based approaches).

61. Mohanty relies on a “contextual score,” along with a “vulnerability score” and a “threat score” as part of an analysis to derive a prioritization score for an asset. Ex.1075 (Mohanty), 1:47-63. Mohanty’s contextual score represents static and dynamic tags describing the context for a given asset. Ex.1075 (Mohanty), 3:60-67, 6:12-26. Mohanty’s threat score represents threats correlated with vulnerabilities based on threat impact and threat type. Ex.1075 (Mohanty), 6:12-26. And Mohanty’s vulnerability score represents a CVSS score associated or correlated with a vulnerability or exposure. Ex.1075 (Mohanty), 6:12-26. The prioritization score considers the contextual score, the threat score, and the vulnerability score. Ex.1075 (Mohanty), 6:12-26.



62. Mohanty's Figure 3 sets forth an agent-based scanning environment for reporting vulnerabilities. Ex.1075 (Mohanty), Fig. 3, 2:53-56, 6:45-48. The scanner (314) is a vulnerability analysis (VA) scanner by Qualys or Rapid7. Ex.1075 (Mohanty), 6:54-58. The VA scanners perform scans on the virtual machines (318) and virtual applications (320) to report vulnerability data including a CVSS score. Ex.1075 (Mohanty), 6:54-58. The CVSS score, which, as I note above, is part of the vulnerability score, and may be accompanied by a CVE ID to identify a specific vulnerability or exposure. Ex.1075 (Mohanty), 6:58-61. A vulnerability module tracks specific vulnerabilities for each asset and generates vulnerability scores based on CVSS scores. Ex.1075 (Mohanty), 6:57-63.

Mohanty uses the vulnerability scores to determine prioritization scores, which provide for “action” based on the prioritization score meeting the threshold. Ex.1075 (Mohanty), at 10:25-27. “In the action [(418)], a remediation for the asset is determined.” Ex.1075 (Mohanty), 10:30-34.

C. Overview of Czarny


63. Czarny is titled “Computer Security Vulnerability Assessment” and describes that an assessment “is performed with product binary data and product vulnerability data that correspond with product identification data.” Ex.1084 (Czarny) Abstract. I understand that Czarny is relied upon by Dr. Stavrou and Petitioner for its alleged disclosures of “determining and reporting a risk level based on the detected vulnerabilities” and “matching application files using binary comparison.” *See, e.g.*, Ex.1002 (Stavrou Decl.) ¶¶270, 272.

D. Overview of Hutchins

64. Hutchins is titled “Offloading Operations to a Replicate Virtual Machine.” Ex.1070 (Hutchins) Cover. Hutchins discloses a “method for detecting malicious code within a first virtual machine comprising creating a snapshot of the first virtual machine and transferring the snapshot to a second machine.” Ex.1070 (Hutchins) Abstract. “A scan operation is run on the snapshot using resources of the second machine.” *Id.* “In response to detecting malicious code during the scan

operation, action is taken at the first virtual machine to address the detection of the malicious code.” *Id.*

65. Figure 6 of Hutchins, reproduced below, is a flowchart showing a disclosed “process for creating and running a replicated virtual machine.” Ex.1070 (Hutchins) ¶39. In step 602, a “synchronous snapshot of the virtual machine device state is taken, excluding the main memory.” *Id.* A replica of the virtual machine is made, which includes copying the device state and configuration file corresponding to the original VM to a new, second snapshot image in step 605. *Id.* ¶40. The main memory is then cloned in step 607, and then both virtual machines can run from the cloned state. *Id.* ¶¶40-41. A script can then be run on the replicated VM to perform, e.g., data mining, backup, virus scanning, etc. *Id.* ¶44. The replicated VM is then “usually discarded.” *Id.* ¶45.

**FIG. 6**

66. Hutchins further discloses that if a virus is detected in the replicated VM, the original VM can be quarantined. *Id.* ¶53.

VIII. The Claims Of The '809 Patent Would Not Have Been Obvious

67. Petitioner and Dr. Stavrou challenge claims 1-23 of the '809 patent based on two grounds of purported obviousness, each of which depends on the

combination of U.S. Pat. No. 11,216,563 (“Veselov,” Ex.1007) and U.S. Patent No. 9,692,778 (“Mohanty,” Ex.1075).

68. In my opinion, none of the claims of the ’809 patent would have been obvious over Petitioner’s and Dr. Stavrou’s prior art, for reasons I explain below.

A. Ground 1: Claims 1-10 and 12-23 Would Not Have Been Obvious Over Veselov and Mohanty

69. In Ground 1 of the Petition, Petitioner and Dr. Stavrou challenge claims 1-10 and 12-23 of the ’809 patent based on the combination of Veselov and Mohanty. *See* Ex.1002, Stavrou Decl. §XII.

70. In my opinion, Petitioner and Dr. Stavrou fail to show that any of claims 1-10 and 12-23 of the ’809 patent would have been obvious over the combination of Veselov and Mohanty.

1. A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov and Mohanty and Would Not Have Had a Reasonable Expectation of Success in Doing So

71. Dr. Stavrou asserts that a person of ordinary skill in the art would have been motivated to make the combination of Veselov and Mohanty on which all of Wiz’s grounds depend. *See* Ex.1002 (Stavrou Decl.), ¶¶107-116. I disagree. To a person of ordinary skill in the art, Veselov and Mohanty have unrelated goals and contradictory teachings, and a person of ordinary skill in the art with Veselov in hand would not have been motivated to combine Mohanty’s teachings that Wiz and Dr. Stavrou rely on for their challenges to the ’809 patent claims. Indeed, in

view of these references' distinctions, which I discuss in more detail below, it appears to me that Wiz and Dr. Stavrou are working backwards from the '809 patent's claims, looking for disclosures in disparate references, without meaningfully considering whether the references they cite would, in fact, be combined in the first place. As noted above, I understand that obviousness arguments based on this type of hindsight are improper.

a. Dr. Stavrou's analysis relies on Veselov's duplicate virtual machine examples that do not analyze a snapshot

72. First, I note that Dr. Stavrou's analysis mixes disclosures from Veselov's Figures 1, 3A, 3B, 5A, and 5B, which describe and relate to different implementations of Veselov's analysis techniques, without explaining why a person of ordinary skill in the art would have wanted to combine or modify those varied implementations. *See, e.g.*, Ex.1002 (Stavrou Decl.), ¶149 (opining about "scanning service" in Figures 1, 3A-3B, and 5A-5B). Dr. Stavrou's analysis conflates the different analysis implementations of Veselov, and as a result, incorrectly asserts that a POSA would have combined Veselov and Mohanty. Specifically, Dr. Stavrou's analysis aims to relate Veselov's Figures 3A-3B duplicate VM assessments to the challenged claims, but the analysis in Veselov's Figures 3A-3B do not analyze a snapshot, as that phrase is properly used in the '809 patent claims. *See* § VI.B.1.

73. Veselov describes two different types of analysis: analysis of a duplicate VM instantiated from a snapshot, as described with reference to Figures 3A and 3B, and direct analysis of a snapshot, as described with reference to Figures 5A and 5B.

74. Figure 3A is reproduced below.




FIG. 3A

Ex.1007 (Veselov) Fig. 3A

As shown above, in Figure 3A the scanning service is part of the computing resource service provider. In the Figure 3A embodiment of Veselov, Veselov states that the virtual machine may “use a logical volume 306, which is composed

of a fixed or variable allocation of physical data storage resources, such as data storage blocks selected from one or more hard disk drives of one or more server computers.” Ex.1007 (Veselov) 11:5-9. The API 338 is used by a user device 302 to “access the scanning service 310.” Ex.1007 (Veselov) 14:37-41. The scanning service 310 “may obtain the volume image 308 of the virtual machine 312, **use the volume image 308 to launch the assessment virtual machine 312A, B, and install the security agent program 330 in the assessment virtual machine’s 312A, B second logical volume 316 as described above** and in accordance with the user input.” Ex.1007 (Veselov) 15:29-34.


75. In explaining Figure 4, Veselov describes “an example method 400 of performing a security assessment of a target virtual machine (VM)” in the computing environment of Figures 3A-3B. Ex.1007 (Veselov) 16:15-18. Veselov explains that, as part of this process, “the scanning service may install the security agent program into the assessment VM instance” and “the scanning service may perform the security assessment on the assessment VM instance, such as by activating the security agent program to execute one or more tasks of the security assessment.” Ex.1007 (Veselov) 16:45-53.

76. In this example, Veselov also explains that “the scanning service 310 can cause processing activity to occur within the assessment virtual machine 312A-B in order to apply tasks of the security assessment to real-time dynamic or static

data, network connections, communications to computing resources, and other computational activities of the assessment virtual machine 312A-B.” Ex.1007 (Veselov) 15:41-47.

77. As I explained above in the Claim Construction section, assessing a virtual machine instantiated from a snapshot, as in this example of Veselov, does not teach analyzing the snapshot as required by the claims. *See* Section VI.B.1. As detailed in that section, a VM instantiated from a snapshot is an instantiated VM, and analyzing that instantiated VM is an analysis of a VM, not an analysis of a snapshot.

78. In contrast, in the Figs. 5A-5B example, Veselov discusses assessing snapshots, but in the context of a different system than the system presented in Figure 3A and 3B.



79. In this example of Veselov, Veselov states that “rather than recreating and instantiating a duplicate VM, the scanning service 510 is configured to analyze the volume image 508 as a data file.” Ex.1007 (Veselov) 17:1-3. Thus, in the Figures 5A-5B examples, Veselov analyzes a snapshot. In this example, Veselov states that “the scanning service 510 can access the volume image 508 (e.g., can read files and other data such as the OS 560A, runtime environment 562A, software

applications 564A, and files/file system 566A) as if it were a disk drive.” Ex.1007 (Veslov) 17:24-29. Veslov further explains that “once the file system of the mounted volume image can be read, the scanning service may compare the names of directories and/or files in the image to character patterns representing known directory and/or filenames to identify "signature" directories and files that identify the software application, version, installation attributes, etc.” Ex.1007 (Veslov) 18:38-44.




FIG. 5B

Ex.1007 (Veslov) Fig. 5B

80. The difference between these implementations, where Figures 3A-3B instantiate from snapshots duplicate/assessment virtual machines and install agents on those assessment virtual machines, but do not analyze a snapshot, and where Figures 5A-5B analyzes a snapshot by reading files from the snapshot as if it were a disk drive, affects whether Veselov would have been combined with Mohanty with a reasonable expectation of success. As detailed below, a person of ordinary skill in the art would not have been motivated to combine the teachings of Mohanty that Wiz and Dr. Stavrou rely upon and would not have reasonably expected success in doing so.

b. A Person of Ordinary Skill in the Art Would Not Have Been Motivated To Combine Veselov with Mohanty With a Reasonable Expectation of Success

81. As discussed above in Section VII.A, Veselov's security systems and methods are specifically designed "to avoid consuming processing time, processing power, and storage space" of target VMs, so the active VMs that Veselov's systems and methods target can continue to "servic[e] requests from [] customers or perform[] other useful work." Ex.1007 (Veselov), 12:21-43. Thus, Veselov's security systems either assess a duplicate of a target VM asset instantiated from a snapshot of the target VM asset, or assess a snapshot; in either case, Veselov's security systems do not assess the target VM asset itself. Mohanty, in contrast, discloses a security process where active target VM assets are scanned

by agents. *See* Ex.1075 (Mohanty), 6:54-58; Section VII.A.2. And, notably, Mohanty's system criticizes security assessment techniques that provide vulnerability assessments of network systems in the past, such as snapshot-based technologies, explaining that "[a] snapshot of a vulnerability assessment report of a system provided in the past becomes obsolete within hours or minutes as virtual machines or workloads change positions within a virtualized environment." Ex.1075 (Mohanty), 1:14-17. A person of ordinary skill in the art would have therefore not been motivated to look to combine Mohanty with snapshot-based systems.

82. Additionally, Veselov teaches that its system is designed to overcome "drawbacks of existing security systems, such as allocation and overuse of processing power, unavailability of resources services that are part of the assessment, and complexity of installation and monitoring of the assessment system." Ex.1007 (Veselov), 2:40-46. Veselov also explains that "Performing a robust security assessment on a target computing device, such as a server, may reduce the target computing device's capabilities. For example, a server that is the target of a security assessment may not be capable of responding to a client as quickly as normal, or may not be capable of providing content as quickly as possible because processor time and memory of the target is used to perform the assessment instead of running its normal workload." Ex.1007 (Veselov), 2:56-63.

A POSA reading Veselov would have therefore not looked to systems that rely on agent-based assessments, like Mohanty, that “perform vulnerability checks (i.e., scans) **on** the virtual machines” (Mohanty 1:11-14, 4:59-62, 6:54-58) because it would run counter to Veselov’s express goals. *See, e.g.*, Ex.1007 (Veselov), 2:58-63, 12:21-43.²

² *See also* Ex.2003 [Trend Micro Enterprise Security, *Virtual Datacenter*], 3 (“Resource-intensive operations such as regular anti-virus scans and pattern file updates can quickly result in an extreme load on the system. When anti-virus scans or scheduled updates simultaneously kick into action on all virtual machines on a single physical system, the result is an ‘anti-virus storm’. This ‘storm’ is like a run on the bank, where the ‘bank’ is the underlying virtualized resource pool of memory, storage and CPU. Server applications and virtual desktop/VDI environments are hampered by this performance impact.”); *cf.* Ex.2004 [Red Hat Virtualization 4.0], 46 (same); Ex.2026 [Wiz Marketing], 5, 7 (cautioning against the use of agents in a cybersecurity system, providing that agent-based systems “lower system performance and lead to increased resource consumption,” and can push systems “into higher-priced deployment tiers, causing unplanned cost increases that lead to budget overruns”); Ex.2019 [Uluski Thesis], 38 (“Previous studies [] have taken measures to quantify how much overhead is introduced by

83. Dr. Stavrou asserts that Veselov and Mohanty perform the “same, predictable function, with predictable results” and that a person of ordinary skill in the art “would have recognized that Mohanty’s techniques would have their same predictable function and achieve predictable results in the context of snapshot-based analysis.” Ex.1002 (Stavrou Decl.), ¶110. I disagree. A person of ordinary skill in the art would recognize material differences between Veselov’s snapshot-based techniques and Mohanty’s agent-based risk assessments and would not endeavor to combine the two references.

84. As detailed above, Veselov performs security assessments in two ways. First, Veselov captures a snapshot to facilitate cloning of the target virtual machine in an assessment machine to perform a security analysis. Cloning active processes of the target machine permits security assessments to be performed without interrupting activity of the target device. Indeed, Veselov recognizes that:

the execution of a real-time anti-virus scanner and have indicated that significant amounts of pressure is place on not only the central processor, but the memory subsystem. ... The issue of pressure on the physical system becomes more immense in the ESX server environment, where multiple virtual machines are executed on one physical machine, and each virtual machine running the windows operating system requires its own instance of an anti-virus program.”).

“[I]n some embodiments it may be suitable or advantageous to perform the security assessment on a different but similar virtual machine than the customer's virtual machine because the assessment can be performed without impacting the customer's virtual machine, which may be servicing requests from its customers or performing other useful work. For example, the virtual computing environment may be a production environment in which allocating virtual computing resources to the security assessment may degrade operation of the virtual machine and negatively impact functionality for end users. In another example, the virtual computing environment may have computing limitations, while a separate test environment of the computing resource service provider may lack these limitations and the security assessment may be performed faster in the test environment.”

Ex.1007 (Veselov), 12:21-36. And Veselov emphasizes that “[p]erforming a robust security assessment on a target computing device, such as a server, may reduce the target computing device’s capabilities.” Ex.1007 (Veselov), 2:56-58.

“For example, a server that is the target of a security assessment may not be capable of responding to a client as quickly as normal, or may not be capable of responding to a client as quickly as normal, or may not be capable of providing content as quickly as possible because processor time and memory of the target is used to perform the assessment instead of running its normal workload.”

Ex.1007 (Veselov), 2:58-63. Accordingly, as described in Veselov’s Figures 3A and 3B embodiments, Veselov uses VM duplication to avoid the slowdowns associated with active assessments on target machines (e.g., assessments which employ agents on the target machines, like Mohanty). In my opinion, these disclosures in Veselov would have led a POSA away from the challenged claims, as “analyzing a snapshot” does not include analyzing a duplicated VM instantiated by a snapshot as discussed in the claim construction section above.

85. Second, Veselov explains that “rather than recreating and instantiating a duplicate VM, the scanning service 510 is configured to analyze the volume image 508 as a data file.” Ex.1007 (Veselov), 17:2-8. This is what is described with reference to Veselov’s Figures 5A and 5B.

86. Mohanty involves agent-based scanners to monitor and collect vulnerability data from target assets for detecting and determining vulnerabilities and setting forth remediations. Ex.1075 (Mohanty), 2:13-35, 6:45-63. Mohanty’s agent-based scanners are implemented *on* the target assets – a practice that is disparaged in Veselov’s systems. *See, e.g.*, Ex.1007 (Veselov), 12:21-43. Specifically, Veselov describes “installation of secure software modules in the target resources” (Ex.1007 (Veselov), 3:16-19)) as a drawback, which is what Mohanty uses to perform vulnerability scans via “VA scanner[s]” and report vulnerability data. Ex.1075 (Mohanty), 45-63. Mohanty relies on “QUALYS” and “RAPID7” scanners, which a person of ordinary skill would recognize involve agent-based scanning that cause the very system slowdowns that Veselov cautions against. *See* Ex.2090 (Muthukrishnan, Automating Cloud Security Governance) , 28 (“deployment of Qualys scanners within each VPC”); Ex.2091 (Sönmez et al., Evaluation of Security Information) (“Rapid7 InsightIDR Collector has been installed on a Windows machine for data collection”); *see also* Ex.2043 [Agent-Based vs. Network Based Internal Vulnerability Scanning] (recognizing even post-

priority date that agent-based scanning is “resource intensive and end[s] up hogging [the] computing and memory space.”).

87. Mohanty’s VA scanners collect data on a continuous basis and “monitor[]” systems through the use of “dynamic tags” in a manner that is materially distinct from the snapshot-based features of Veselov. Ex.1075 (Mohanty), 5:38-58; Ex.2044 [Runecast] (recognizing post-priority date that agent-based scanning involves “real-time monitoring of the host system” and “consume[s] system resources, which can impact performance” while agentless scanning involves “no additional resource usage.”). Mohanty makes clear that its systems are not compatible with and not meant to utilize snapshot-based technology, noting that a “snapshot of a vulnerability assessment report of a system provided in the past becomes obsolete within hours or minutes as virtual machines or workloads change positions within a virtualized environment.” Ex.1075 (Mohanty), 1:14-17. While Veselov uses agent-based scanning on duplicate VMs instantiated from a snapshot (*see, e.g.*, Ex.1007 (Veselov) 15:29-38 (“install the security agent program 330 in the assessment virtual machine’s 312A, B”), 14:22-36 (similar), 15:41-47), for example, in the Figures 3A-3B implementation, consistent with the criticism of snapshots in Mohanty, Veselov **does not** use agent-based scanning when analyzing a snapshot itself. This is logical; as Dr. Stavrou acknowledged, it is only possible to install an agent on a virtual machine (e.g., in

Veselov's Figures 3A-3B), and it is not possible to install an agent on a snapshot (e.g., as in Veselov's Figures 5A-5B) (Ex.2172 (2nd Stavrou Tr.) 53:14-20). Thus, the only implementation in Veselov that could be combined with Mohanty's teachings is the implementation in Veselov that does not analyze a snapshot. Accordingly, a person of ordinary skill in the art would understand that the snapshot-analysis teachings of Veselov and the teachings of Mohanty would not have been combined, as the implementations of each reference do not suggest combining their teachings.

88. Dr. Stavrou also asserts that "the type of snapshot-based analysis discussed in Veselov typically examines the same types of data (e.g., applications, configurations, files, etc.) and identifies the same types of security issues as an assessment performed directly on the target VM [as in Mohanty]." Ex.1002 (Stavrou Decl.) ¶110. I disagree. A snapshot-based analysis, as described in Veselov, does not analyze the same types of data and security issues as an agent-based analysis. The references cited by Dr. Stavrou support differences between snapshot-based and agent-based analyses. Suneja, for example, distinguishes snapshotting from "[a]gent assisted access" (as described in Mohanty). Ex.1038 (Suneja), 134. As Suneja states, snapshotting does not require guest cooperation, whereas agent assistance access requires guest cooperation to facilitate runtime state extraction. Ex.1038 (Suneja), 134. Similarly, Jiang provides distinctions

between agent-based and “out-of-the-box” approaches, highlighting a “gap” between the two approaches. Ex.1043 (Jiang), 1. For example, Jiang explains that agent-based or “in-the-box” approaches are “traditional host-based antimalware systems” that “run inside the very hosts they are protecting,” making the hosts “vulnerable to counter detection and subversion by malware.” Ex.1043 (Jiang), 1. Jiang then contrasts this approach with “recent solutions” that place “malware detection facilities outside of the protected VM.” Ex.1043 (Jiang), 1. Dr. Stavrou’s evidence therefore makes clear that a person of ordinary skill would not have combined Mohanty’s agent-based analyses when considering Veselov’s snapshot techniques. Accordingly, I disagree that in a combination of Veselov and Mohanty’s techniques, the references would “perform their same, predictable function, with predictable results” as Dr. Stavrou asserts. Ex.1002 (Stavrou Decl.), ¶ 110. Veselov’s analysis of snapshots as a data file would not perform their same function when Mohanty’s teachings would be combined, as Mohanty’s agents could not be installed on Veselov’s snapshots. A POSA would also not have expected Mohanty’s agent-based teachings to have the same data necessary if using Veselov’s snapshots. A person of ordinary skill in the art would thus recognize material differences between Veselov’s snapshot-based analysis of VMs and Mohanty’s agent-based monitoring of active virtual assets and would not have combined the two references.

89. Additionally, I note that a VM's snapshot is, as the '809 patent states, "a copy of the machine's virtual disk (or disk file) at a given point in time." Ex.1001 ('809 patent) 4:58-62. Generally, snapshots are a "storage function" to allow administrators "to create a restore point of a virtual machine's operating system, applications, and data at a certain point in time." Ex.2006 [Red Hat Virtualization 4.3 – *Virtual Machine Snapshots*]. And snapshots "provide a powerful tool with which a virtual machine can be cloned by one click including also the running system's memory." Ex.2007 [Birk, *Technical Challenges*], 4. That is, a person of ordinary skill in the art would have recognized snapshots to act as a storage function for data to be read from, but are not an active, executing computing device, like a virtual machine. And Veselov's snapshots are not a computing device upon which an agent can be installed, as reflected in the fact that Veselov installs "security agents" only on a duplicate VM (Figs. 3A-3B) and not in a snapshot (Figs 5A-5B).

90. Additionally, as I detail more below in the section detailing objective indicia of non-obviousness, before the '809 patent, those in the industry commented that only agent-based security analysis implementations, like those described in Mohanty, would provide sufficient data to enable detailed monitoring and protection, and by contrast, snapshot-based approaches were "lightweight" and

provided “rudimentary security capabilities.” Ex.2106 (Agentless Approaches Sideline Cloud Workload Protection Strategies), 1-2.

91. Indeed, Veselov itself is consistent with this knowledge in the industry, as it makes clear that agent-based security analyses (Veselov’s Figs. 3A-3B) result in significantly different data compared to analyzing a snapshot(Veselov’s Figs. 5A-5B). In Veselov’s “security agent” based assessment described in Figures 3A-3B, Veselov’s agent can “apply tasks of the security assessment to real-time dynamic and static data, network connections, communications to computing resources, and other computational activities of the assessment virtual machine 312A-B.” Ex.1007 (Veselov), 15:41-47. For example. Veselov explains that the system may install a security agent program in an assessment virtual machine to monitor the activity of the assessment virtual machine without intruding on the target virtual machine. Ex.1007 (Veselov), 15:29-38 (“install the security agent program 330 in the assessment virtual machine’s 312A, B”), 14:22-36 (“installs a security agent program 330 into the assessment virtual machine 312A-B”). Veselov further explains that its agent-based analysis of a duplicate or assessment VM may also include “intercept[ing] traffic” and “stimulat[ing] activity in the assessment virtual machine” in order to generate a traffic profile for the target virtual machine based on the monitored activity of the assessment virtual machine. Ex.1007 (Veselov), 15:60-14.

92. Veselov's explanations for what data is detected when analyzing a snapshot, such as in Figures 5A-5B, however, are limited in comparison to the dynamic and wide-ranging data available in Veselov's agent-based duplicate virtual machine implementation in Figures 3A-3B. Specifically, Veselov teaches a limited set of data that can be evaluated when analyzing a snapshot as a data file. For example, Veselov contemplates "us[ing] string or other pattern comparisons to identify files, software packages, application versions, and the like, for software installed on the target virtual machine" when analyzing a snapshot as a data file. Ex.1007 (Veselov), 18:34-38. Persons of ordinary skill in the art would understand that determinations and security assessments based on this basic and rudimentary snapshot information are distinguishable from the dynamic data discerned from Veselov's agent-based analyses. And Veselov does not teach that analyzing a snapshot as a data file would permit any analysis of "dynamic" data; nor does Veselov demonstrate that such analyses would result in understanding the "network connections" or "communications to computing resources," configurations, or any of the other information available when using a security agent as in the Figure 3A-3B embodiments.³ Thus, I disagree with Dr. Stavrou's

³ Dr. Stavrou appears to indicate that Veselov's duplicate-VM analyses does not "use[] an agent installed on the target VM, nor do they require an agent installed

opinion that agent-based (Mohanty) and snapshot-based (Veselov Figures 5A-5B) analyses “examine[] the same types of data” and “identif[y] the same types of security issues as an assessment performed directly on the target VM” (Ex.1002 (Stavrou Decl.) ¶110) because it is contradicted by Veselov and other evidence. .

93. Dr. Stavrou also cites to a laundry list of background references to assert his rationale for combining Veselov and Mohanty. *See* Ex.1002 (Stavrou Decl.), ¶¶109-116 (citing Ex.1008 (Basavapatna); Ex.1014 (Klemperer); Ex.1078 (Hufsmith); Ex.1025 (NIST Incident Handling); Ex.1034 (Tang); Ex.1044

on the assessment VM.” Ex.1002 ¶152. I disagree. Veselov expressly discloses that in the context of analyzing a duplicate VM, an “assessment agent” or “security agent” is always installed “in” or “into” the assessment VM. *See* Veselov, Fig. 4 (box 412 “Install Assessment Agent in Assessment VM”), 3A, 3B (“security agent”), 14:22-25 (“the scanning service 310 installs **a security agent program 330 into the assessment virtual machine 312A-B**”). The security agent program 330 may be a software application or software package that can be activated by the scanning service 310 to monitor activity of the assessment virtual machine 312A-B.”), 15:31-33 (“and install the security agent program 330 in the assessment virtual machine's 312A, B”), 16:45-52 (“install the security agent program into the assessment VM instance”), 21:3-8, 21:57-60.

(Giakouminakis); Ex.1045 (Canzanese); Ex.1046 (Crabtree); Ex.1048 (Price); Ex.1063 (Hibbert); Ex.1076 (Milner); Ex.1077 (Khalid); Ex.1085 (Hovor)). I have reviewed the cited portions of these, and none of them appear to explain why persons of ordinary skill in the art would make the purported combination of Veselov and Mohanty. Klemperer, Tang, Giakouminakis, Canzanese, Crabtree, Milner, Hufsmith, and Hovor do not implement Mohanty's agent-based scanning and therefore provide no support for why a person of ordinary skill in the art would consider or combine Mohanty when analyzing the techniques of Veselov. *See generally* Ex.1014, Ex.1025, Ex.1034, Ex.1044, Ex.1045, Ex.1046, Ex.1076, Ex.1078, Ex.1085. Hibbert and Basavapatna, like Mohanty, are directed to agent-based collection of vulnerability data and do not support combining agent-based techniques with Veselov's snapshot analysis techniques.⁴ Ex.1063; Ex.1008. And the NIST Guide broadly discloses prioritizing incidents based on business factors and does not demonstrate that a person of ordinary skill in the art would have combined Mohanty's agent-based techniques with the teachings of Veselov, or that

⁴ To be clear, where I refer to Veselov's "snapshot analysis techniques" or Veselov's snapshot analysis, that is meant to refer to Figures 5A-5B of Veselov, which demonstrate embodiments for analyzing a snapshot. These references to Veselov's "snapshot analysis techniques" do not apply to Veselov's Figures 3A-3B, which, as I describe throughout this declaration, do not analyze a snapshot.

the features of Mohanty were well-known or would be insightful to a person of ordinary skill in the art considering Veselov's snapshot scanning technology. Ex.1025. At base, Dr. Stavrou's other cited references similarly fail to bridge the gap between Veselov's and Mohanty's inventions and do not demonstrate any rationale for the purported combination of Mohanty's agent-based scanning and Veselov's snapshot-based scanning of virtual machines. Accordingly, Dr. Stavrou has not shown that his purported Veselov-Mohanty combination would have rendered the challenged claims obvious. I address Dr. Stavrou's additional references below.

94. Price (Ex.1045) which Dr. Stavrou refers to for its "indication of a seriousness or priority of [an] issue" (Ex.1002 (Stavrou Decl.), ¶109) relies on offline *image* collection, unlike either Veselov or Mohanty, and agent-based data collection, which suffers the same deficiencies as described above with respect to Mohanty. Price is therefore inapplicable to Veselov's analysis of a snapshot as a data file and further does not support combining Mohanty's agent-based scanning techniques with those of Veselov.

95. The techniques of Khalid (Ex.1077), which Dr. Stavrou refers to for "comparing characteristics of suspicious objects to CVE database" (Ex.1002 (Stavrou Decl.) ¶109) use a "mobile device[] ... deploying an endpoint application." Ex.1077 (Khalid) 17:63-66. Thus, Khalid's techniques are again like

Mohanty's agent-based techniques. But Khalid does not discuss snapshots, and does not support combining the agent-based analysis of Mohanty with Veselov's snapshot-based analysis in Figures 5A-5B.

96. In sum, though Dr. Stavrou cites to numerous references with disparate teachings throughout his arguments that a person of ordinary skill in the art would have been motivated to combine Veselov and Mohanty, examining those references reveals that none of them, whether taken alone or in combination, support such a motivation to combine. Many of the references discourage combining Veselov and Mohanty, or simply further support that there are differences between the applied teachings of Veselov and Mohanty, such that a person of ordinary skill in the art would have not combined their teachings. Dr. Stavrou does not address this.

97. Dr. Stavrou does not address the issues with his combination identified above, nor does Dr. Stavrou address how or why Mohanty's agent-based analysis would be applied to Veselov's snapshot-based analysis. It is therefore my opinion that persons of ordinary skill in the art would not have been motivated to apply Mohanty's agent-based techniques to Veselov's snapshot-based techniques. To be sure, Mohanty's techniques are squarely against Veselov's efforts to reduce the complexity of security assessments and are inapposite to Veselov's caution of

“overuse of processing power[] and resource unavailability. Ex.1007 (Veselov), 2:42-46.

98. Additionally, I disagree with Dr. Stavrou’s position that “a POSA would have been motivated to combine the teachings of Veselov and Mohanty to provide more accurate assessment of risks posed by particular threats and the asset as a whole.” Ex.1002 (Stavrou Decl.), ¶112. Indeed, as Dr. Stavrou notes, “Mohanty recognizes that effective risk assessment is important in security assessments of virtual assets whose risk profiles may change frequently.” Ex.1002 (Stavrou Decl.), ¶112. And, as I discuss above, Mohanty sets forth an agent-based scanning system that does not rely on point-in-time captures and in fact actively discourages such point-in-time captures because they might become “obsolete.” Ex.1075 (Mohanty), 1:14-17. But Veselov’s snapshots are point-in-time captures that would, in Mohanty’s view, become “obsolete”, and thus, a person of ordinary skill in the art would not have sought to combine Veselov’s point-in-time captures having static data about the past, with Mohanty’s focus on avoiding such potentially obsolete assessments by using agent-based techniques on live computing devices. Additionally, Dr. Stavrou’s position does not address a key flaw with the contextual data that he touts for his motivation. Specifically, the contextual data of Mohanty relies on the agent-based scanners I discuss above to collect information regarding vulnerabilities, context, and threats. Veselov’s

snapshot-based analysis does not mention or suggest agents for its target machines. And, as Dr. Stavrou stated during his deposition, agents could not be installed on a snapshot. Ex.2172 (2nd Stavrou Tr.) 53:14-20 (“I don’t think that an agent can be installed on a snapshot. I think an agent can only be installed on a virtual machine.”). Veselov also acknowledges the drawbacks of, and counsels against, using or adding complex monitoring or assessment systems (*see, e.g.* Ex.1007 (Veselov), 2:42-46), but those are the precise types of assessment systems that Mohanty relies upon. Persons of ordinary skill in the art would have therefore understood that Mohanty’s techniques for running security assessments on active VMs are materially different from Veselov’s snapshot-based techniques to analyze VMs. Persons of ordinary skill in the art thus would not be motivated to make Dr. Stavrou’s combination.

99. Dr. Stavrou further asserts that a person of ordinary skill would have been motivated to make this combination because “[i]t was well-known that prioritizing alerts was beneficial for security assessments in several ways, including rapid identification of which cyber threats are most at risk and which threats and vulnerabilities are the most dangerous.” Ex.1002 (Stavrou Decl.), ¶113. But this statement, even given the numerous background references Dr. Stavrou cites, does not support that a person of ordinary skill would have combined Veselov’s snapshot-based analysis with Mohanty’s agent-based techniques, or

how a person of ordinary skill in the art would have accomplished such a combination and how such a combination would work. As discussed, Mohanty provides for its security screening in a context that is completely divorced from the goals set forth in Veselov. And contrary to Dr. Stavrou's assertions, Mohanty fails to describe the prioritization features that the challenged claims require. *See* Sections VII.B, VIII.A. Moreover, to persons of ordinary skill in the art, Mohanty's Qualys and Rapid7 agent-based scanners are vastly different from the security assessment systems and methods disclosed by Veselov. Persons of ordinary skill in the art, for instance, would not be motivated to combine teachings in support of agent installations to effect security assessments on snapshots of target resources.

100. Relatedly, Dr. Stavrou's opinion that "applying Mohanty's prioritization techniques would similarly benefit snapshot-based assessments such as those described in Veselov" (Ex.1002 (Stavrou Decl.), ¶114) does not address the underlying differences between Mohanty's agent-based techniques and Veselov's analysis of a snapshot. As detailed above, a person of ordinary skill in the art would have understood that Mohanty's prioritization techniques are reliant upon its agent-based techniques, which analyze real-time, dynamic data from executing processes on a virtual machine. *See, e.g.*, Ex.1075 (Mohanty), 1:36-40 (CVSS is not "sufficient in a dynamic environment such as in virtualization

space”), 4:15-19 (“Dynamic tags 106 have information that is subject to change during the lifespan of the asset. Dynamic tags 106 can also be referred to as security tags, since the changing information is of interest regarding the security of the asset.”). A person of ordinary skill in the art would have understood that, in the context of a snapshot-based analysis, like in Veselov, there is not “information that is subject to change” or “changing information” – the information is static as of the point-in-time capture represented by the snapshot. Thus, a person of ordinary skill in the art would not have recognized that Mohanty’s prioritization techniques would benefit Veselov’s snapshot-based assessments, or how such prioritization techniques would have been incorporated with or adapted to fit with Veselov’s snapshot-based techniques.

101. Additionally, I note that Dr. Stavrou’s opinion that “Mohanty’s teachings regarding prioritizing/reporting based on a determined risk were directly applicable to Veselov’s [alleged] general teachings regarding providing assessment results” or “alerts” “to a user” (Ex.1002 (Stavrou Decl.), ¶115) ignores the teachings of Veselov. Specifically, as I discuss below, Veselov does not teach alerts. And, modifying Veselov to include alerts would not be obvious to persons of ordinary skill in the art; in fact, a person of ordinary skill in the art reading Veselov would not understand Veselov to include or necessitate alerts in view of the finite universe of findings generated by Veselov’s assessment. *See, e.g.,*

Ex.1007 (Veselov), 10:18-24. Given this failure in Veselov, there would be no improvement “regarding prioritizing/reporting based on a determined risk” as Dr. Stavrou asserts.

102. For similar reasons, it is also my opinion that Dr. Stavrou’s contention that a person of ordinary skill in the art would have reasonably expected to succeed (Ex.1002 (Stavrou Decl.), ¶116) in making Wiz’s and Dr. Stavrou’s proposed Veselov-Mohanty combination is incorrect. As described above, nothing in Dr. Stavrou’s tertiary references demonstrates that a person of ordinary skill would have a reasonable expectation of success in combining Veselov’s snapshot-based techniques with Mohanty’s agent-based techniques.

103. As Dr. Stavrou himself admitted, it is not possible to install an agent on a snapshot. Ex.2172 (2nd Stavrou Tr.) 53:14-20 (“I don’t think that an agent can be installed on a snapshot. I think an agent can only be installed on a virtual machine.”). I agree with his testimony, because, for example, it is consistent with Veselov’s description of a snapshot as “a state of the resource at a point in time.” Ex.1007 (Veselov) 3:23-27. A person of ordinary skill in the art would not have reasonably expected techniques that rely on agent-based data to be successfully combined with the different data that is determined from a snapshot-based approach as in Veselov.

104. So, in Veselov's Figures 5A-5B, the only implementations of Veselov that analyze a snapshot, installation of an agent is not possible, which means that Mohanty's techniques (all of which rely on agents) could not be combined with Veselov's analysis of a snapshot. Given these differences between the teachings of Veselov and Mohanty, it cannot be said that there would have been any expectation of success, much less a reasonable expectation of success, in combining the teachings.

105. For the foregoing reasons, it is my opinion that Wiz and Dr. Stavrou have failed to show—because they cannot show—that a person of ordinary skill in the art would have made the Veselov-Mohanty combination on which every ground in the Petition depends. A person of ordinary skill in the art would not have made and would not have reasonably expected to succeed in making, the Veselov-Mohanty combination.

2. Even If Combined, Veselov and Mohanty Do Not Disclose or Suggest Multiple Limitations in the Independent Claims

106. Even if a person of ordinary skill in the art would have made the Veselov-Mohanty combination that Dr. Stavrou and Wiz urge, it is my opinion that the proposed combination does not disclose or suggest multiple limitations in the claims, as analyzed below.

a. Limitations 1.5, 16.5, 19.5⁵

107. Limitations 1.5, 16.5, and 19.5 recite “[correlating/correlate] the determined potential cyber vulnerabilities with a network location of the protected virtual cloud asset.” Dr. Stavrou fails to show that Veselov and Mohanty disclose or suggest these limitations. Ex.1002 (Stavrou Decl.), ¶¶161-164.

108. Dr. Stavrou does not allege that Veselov discloses correlating vulnerabilities with a network location of a protected virtual cloud asset. Ex.1002 (Stavrou Decl.), ¶¶161-164. Instead, Dr. Stavrou asserts only that “Veselov teaches *consideration of network information* as part of its security analysis of the target VM.” Ex.1002 (Stavrou Decl.), ¶162 (emphasis added). But considering network information is not enough to meet this claim limitation. Further, the network information or “configuration of network protocols” on which Dr. Stavrou relies upon is described in Veselov as the “properties of the target resources...*that are relevant to the security assessment.*” Ex.107 (Veselov), 9:55-58. And, it is only in Veselov’s Figures 3A-3B embodiments for agent-based assessments of duplicate or assessment VMs that Veselov describes network configurations or

⁵ In this declaration, I use the claim limitation labels that Dr. Stavrou and Wiz used. A listing of the claims of the ’809 patent, with those labels, is presented at the beginning of the Petition, in a “Listing of Challenged Claims.” Pet. at iii-viii.

protocols, such as connections to other resources in a computing environment, as relevant to a security assessment. Ex.1007 (Veselov), 15:41-16:14, 21:3-23. For example, these embodiments of Veselov involve “network connections” and “communications to computing resources” along with the interception of network traffic to create a “profile” to determine assessments to perform. Ex.1007 (Veselov), 15:41-16:14. In such assessments, Veselov teaches that a security agent may collect telemetry data from a duplicated machine, the telemetry data describing, for example, “network activity.” Ex.1007 (Veselov), 21:3-23. But as I discuss above, these embodiments are not applicable to the challenged claims because analyzing a duplicate VM instantiated from a snapshot is not analyzing a snapshot and therefore does not disclose or render obvious the claimed “analyz[e/ing] a snapshot.” Veselov’s Figures 5A and 5B embodiments, however, do include the analysis of a snapshot as a data file. *See, e.g.*, Ex.1007 (Veselov), 17:1-3. But Veselov does not disclose using, analyzing, or determining network protocols and services or connections to other computing resources in these embodiments and thus do not teach “consideration of network information” as Dr. Stavrou contends. Ex.1007 (Veselov), 16:56-18:56, Figs. 5A, 5B, 6.

109. I also note that it would not have been obvious to a person of ordinary skill at the time of filing the ’809 patent invention that analyzing a snapshot as a data file would include determining network location. I have reviewed industry

documentation that demonstrates a widespread belief at the time of filing the '809 patent invention that snapshot-based analyses only provided “lightweight” and “rudimentary security capabilities” that were “mostly useful for compliance” rather than the “dynamic evolution of workloads.” Ex.2106 (Agentless Approaches Sideline Cloud Workload Protection Strategies), 1-2.

110. Dr. Stavrou further relies on Mohanty to address these limitations. More specifically, Dr. Stavrou relies on Mohanty’s “workload context” that includes network location “tags” correlated with vulnerabilities determined from agent-based vulnerability assessment tools. Ex.1002 (Stavrou Decl.), ¶163. But, as I describe above in detail, these teachings in Mohanty are based on Mohanty’s agent-based techniques, which do not apply to and would not have been combined with Veselov’s snapshot analyses. Dr. Stavrou does not explain how Mohanty’s “workload context” would apply in Veselov’s system, or how Veselov would have the same “workload context” to form the claimed correlation.

111. I also note that Dr. Stavrou stated in his deposition that he did not offer any opinions in his Declaration that context “tags” would be applicable to a system based on analyzing a snapshot of a virtual cloud asset, such as the system in Veselov, or that such a snapshot would include any of the tags that are discussed in Mohanty. Ex.2172 (2nd Stavrou Tr.) 241:16-242:10 (“I’m not claiming either that it does or it doesn’t, because I don’t know what the snapshot data would look

like.”). This further supports that a person of ordinary skill in the art would not have combined Mohanty’s teachings with Veselov’s snapshot-based analysis teachings, and would not have had a reasonable expectation of success in doing so.

112. Therefore, it is my opinion that Dr. Stavrou fails to show that limitations 1.5, 16.5, and 19.5 would have been obvious over Veselov and Mohanty.

b. Limitations 1.6, 16.6, and 19.6

113. Limitations 1.6, 16.6, and 19.6 recite “using the determined plurality of potential cyber vulnerabilities and the network location of the protected virtual cloud asset to determine a risk of the protected virtual cloud asset to the cloud computing environment.” Dr. Stavrou does not rely on Veselov for these limitations (Ex.1002 (Stavrou Decl.) ¶¶165-172) and fails to show that Mohanty discloses or suggests these limitations.

114. For limitations 1.6, 16.6, and 19.6, Dr. Stavrou relies on the same analysis for a correlated network location as seen with respect to limitations 1.5, 16.5, and 19.5. Accordingly, my explanation addressing limitations 1.5, 16.5, and 19.5 applies equally to limitations 1.6, 16.6, and 19.6.

115. Even where Dr. Stavrou asserts that “[a]t a minimum” these limitations would have been obvious, his analysis ignores the distinction between agent-based analysis in Mohanty and snapshot-based analysis. Ex.1002 (Stavrou

Decl.) ¶ 169. For example, Dr. Stavrou's citation to Basavapatna and Giakouminakis does not support his conclusions, because both references rely on agent-based techniques that would not have been combined with Veselov's snapshot analysis. Hufsmith relates to container images, which similarly is inapplicable to Veselov's snapshot analysis.

116. Therefore, it is my opinion that Dr. Stavrou fails to show that limitations 1.6, 16.6, and 19.6 would have been obvious over Veselov and Mohanty.

c. Limitations 1.7, 16.7, and 19.7

117. Limitations 1.7, 16.7, and 19.7 recite "[prioritize/prioritizing], by the determined risk, the plurality of potential cyber vulnerabilities." Dr. Stavrou does not rely on Veselov and fails to show that Mohanty discloses or suggests these limitations. Ex.1002 (Stavrou Decl.) ¶¶173-176.

118. For limitations 1.7, 16.7, and 19.7, Dr. Stavrou relies on the same analysis for a correlated network location as seen with respect to limitations 1.5, 16.5, and 19.5. Accordingly, my explanation addressing limitations 1.5, 16.5, and 19.5 applies equally to limitations 1.7, 16.7, and 19.7.

119. And, as with limitations 1.6, 16.6, and 19.6, Dr. Stavrou's assertion that this limitation would have been obvious "[a]t a minimum" relies again on

Basavapatna, Giakouminakis, and Hufsmith, which do not support his combination or his assertions that this limitation would have been obvious.

120. Therefore, it is my opinion that Dr. Stavrou fails to show that limitations 1.7, 16.7, and 19.7 would have been obvious over Veselov and Mohanty.

d. Limitations 1.8, 16.8, 19.8

121. Limitations 1.8, 16.8, and 19.8 recite: “[report/reporting] the determined plurality of potential cyber vulnerabilities as alerts prioritized according to the determined risk.” Dr. Stavrou relies on Veselov in view of Mohanty to teach these limitations. Ex.1002 (Stavrou Decl.), ¶¶177-184. It is my opinion that Dr. Stavrou fails to show that Veselov and Mohanty disclose or suggest these limitations.

122. Contrary to Dr. Stavrou’s position that “Veselov teaches reporting ‘the determined plurality of potential cyber vulnerabilities as alerts,’” (Ex.1002 (Stavrou Decl.) ¶178), Veselov does not suggest its techniques involve alerts. Instead, Veselov teaches that the “scanning service 110 may be configured to perform security assessments and produce assessment results based at least in part on snapshot data.” Ex.1007 (Veselov), 5:9-13. And Veselov provides that the scanning service may “process the snapshot data” and “generate one or more assessment results.” Ex.1007 (Veselov), 8:19-21. Additionally, Veselov states

that its service can perform actions such as “storing or delivering the results.” Ex.1007 (Veselov), 18:53-56. But a person of ordinary skill in the art would understand that storing and producing or delivering results is not the same as an alert. Instead, Dr. Stavrou’s arguments conflict with the plain reading of the limitations 1.8, 16.8, and 19.8. A person of ordinary skill in the art would recognize that these limitations do not end after reciting “reporting the determined plurality of potential cyber vulnerabilities,” but continue and require that the determined plurality of potential cyber vulnerabilities be reported “as alerts.” *See* Ex.1001 (’809 patent), cls. 1, 16, 19. And, the specification of the ’809 patent confirms my understanding, as it also demonstrates that reporting results “as alerts” involves techniques beyond simply reporting results. *See* Ex.1001 (’809 patent), 7:53-54 (“[T]he detected cyber threats and/or vulnerabilities are reported, for example, as alerts”), 6:40-47 (disclosing that “detected cyber threats” that are “reported” to a security information and event management (SIEM) system “may be filtered or prioritized” to “reduce[] the number of alerts reported to the user”).

123. A person of ordinary skill in the art would have understood that communicating results, as in Veselov, does not require or suggest reporting detected cyber vulnerabilities “as alerts,” as claimed. *See, e.g.,* Ex.1001 (’809 patent), 7:53-54 (“[T]he detected cyber threats and/or vulnerabilities are reported, *for example, as alerts*”) (italics added), cls. 1, 16, 19. Indeed, Dr. Stavrou agrees

that “results” and “alerts” are different. Ex.2077, 287:4-21 (“I would say that alerts are different than reports, and I agree with you on that.”). Dr. Stavrou also stated that an alert “potentially makes it stand out.” Ex.2077, 287:4-21. But Veselov’s disclosure of simply “communicating” results does not make any particular result stand out, and is therefore not enough to constitute an alert. For example, Veselov’s “assessment results” include “findings” resulting from “applying the rules of the assessment” that do not suggest the urgency implied by an alert. Ex.1007 (Veselov), 10:18-22. Rather, simply “communicating” results in Veselov suggests a standard report. Persons of ordinary skill in the art would have understood that alerts and reports are different and can involve different system components, capabilities, and purposes. *Compare* Ex.2010 [U.S. Pat. No. 9,467,465 (“Hibbert”)], 20:34-44 (“The **alert** module 516 may **generate an alert** based on the assessment. If one or more applications or files identified within the records of the assessment request are classified as vulnerable or highly suspicious, the alert module 516 may generate an alert. The alert may flag the vulnerability, potentially identify the exploit, indicate the degree of danger, or provide any other information. The alert may be provided in any manner. For example, the alert module 516 may provide one or more alerts via email, SMS text message, MMS, web page, intranet alert, extranet alert, or any other way.”) *with* Ex.2010 [U.S. Pat. No. 9,467,465 (“Hibbert”)], 20:12-21 (“The **report** module 514 may **generate any**

kind of report indicating the results of one or more assessments. Information contained within the report may include the digital device, one or more users of the digital device, applications and files identified by the records which may be safe, vulnerable, suspicious, or unknown, or any other kind of information. The report module 512 may provide the report to the digital device associates with the assessment, the security administrator system 220, a network administrator, or any digital device(s) or individual(s).”).⁶

⁶ See also, e.g., Ex.2011 [US Pat. No. 10,298,599 (“Zhang”)], 9:50-67 (providing that a system may “*report the results without* causing the browser to notify, request permission from, or otherwise *alert*, a user using the browser” because “[p]resenting an alert [] interrupts the user’s experience with the web page”); Ex.2012 [US Pat No. 9,594,912 (“Thioux”)], 10:41-48 (illustrating a specific “alert generation logic” as a subset of “reporting logic” to “generate an alert for the client device”); Ex.1023 (ITL Bulletin), at 4 (“Agencies should use CVE in their internal reports of vulnerability scans, notifications to system owners of observed vulnerabilities, and alerts identifying the vulnerabilities targeted by active exploits.”); Ex.1017 (US Pat. No. 11,106,784), 13:12-19 (“generate a report to be shown in the report interface”), 13:26-38 (“transmitting an alert (e.g., to a system admin)”).

124. Dr. Stavrou, however, incorrectly argues that Veselov teaches alerts by “provid[ing] the assessment results, including the identified vulnerabilities, to the user through a user device and/or user interface.” Ex.1002 (Stavrou Decl.), ¶178 (quoting Ex.1007 (Veselov), 10:30-36). But, as I describe above, providing assessment results to a user interface is not the same as reporting assessment results *as alerts*. And the ““providing [of] the assessment results” through a user device (Ex.1007 (Veselov), 10:30-36) does not suggest that the results presented in a user-readable format are presented as alerts. Instead, reporting information as an alert requires an additional step beyond simply reporting the information. For example, as stated above, it is my opinion that alerts are notifications triggered to inform stakeholders of a situation that needs attention. *See* Ex.2011 [U.S. Pat. No. 10,298,599 B1 (“Zhang”)], 9:50-67 (providing that a system may “report the results without causing the browser to notify, request permission from, or otherwise alert, a user using the browser” because “[p]resenting an alert [] interrupts the user’s experience with the web page”); Ex.2012 [U.S. Pat. No. 9,594,912 B1 (“Thioux”)], 10:41-48 (illustrating a specific “alert generation logic” as a subset of “reporting logic” to “generate an alert for the client device”). And persons of ordinary skill in the art would recognize that the act of providing or reporting results does not suffice to demonstrate an alert, which requires technical features beyond simply conveying information to an end user. *See, e.g.,* Ex.2010

[U.S. Pat. No. 9,467,465 B2 (“Hibbert”)], 20:34-44; *See also* Ex.2011 [U.S. Pat. No. 10,298,599 B1 (“Zhang”)], 9:50-67; Ex.2012 [U.S. Pat. No. 9,594,912 B1 (“Thioux”)], 10:41-48. Veselov’s Figures 2, 4, and 6 simply confirm – in flow charts – that assessment results are provided to a user and stored. Ex.1007 (Veselov), FIGs. 2, 4, and 6. *See also* Ex.1017 13:12-19 (“generate a report to be shown in the report interface”), 13:26-38 (“transmitting an alert (e.g., to a system admin)”); Ex.1023 at 4 (“Agencies should use CVE in their internal reports of vulnerability scans, notifications to system owners of observed vulnerabilities, and alerts identifying the vulnerabilities targeted by active exploits.”). None of these disclosures teach or suggest the reporting of vulnerabilities as alerts, as is required by the claim language.

125. Dr. Stavrou’s citations to the assessment results described by Veselov similarly fail to demonstrate Veselov teaches the claimed “alerts.” Ex.1002 (Stavrou Decl.), ¶178. For example, Dr. Stavrou relies on select portions of Veselov which detail the substance of the assessment results (i.e., that they “describe the identified risks, including vulnerabilities discovered in the snapshot data.”). Ex.1002 (Stavrou Decl.), ¶178. But Dr. Stavrou’s citations merely highlight the distinction between Veselov’s reporting of results and the ’809 patent’s reporting of results as alerts, and in fact, Dr. Stavrou does not provide a single citation to an alert in Veselov.

126. Additionally, Dr. Stavrou provides no indication as to why a person of ordinary skill in the art would modify Veselov's techniques to report its assessments results as alerts as limitations 1.8, 16.8, and 19.8 require. For example, a person of ordinary skill in the art reading Veselov would not understand Veselov to include or necessitate alerts in view of the finite universe of findings generated by Veselov's assessment. *See, e.g.*, Ex.1007 (Veselov), 10:18-24, 2:42-46, 7:15-16.

127. Dr. Stavrou also relies on select portions of Mohanty. Ex.1002 (Stavrou Decl.), ¶¶179-184. But for the reasons I describe above, a person of ordinary skill would not have been motivated to modify Veselov to incorporate Mohanty's techniques.

128. Additionally, Dr. Stavrou relies on Mohanty's disclosure of "an alert" and "prioritization information." Ex.1002 (Stavrou Decl.), ¶¶179-180. But Mohanty's "alerts" are based on specific prioritization scores that address a need for a "continuous prioritization system to track critical vulnerabilities and take actions as changes occur." Ex.1075 (Mohanty), 1:19-22, 6:54-58 ("a VA scanner by QUALYS or RAPID7, performs vulnerability checks (i.e., scans) on the virtual machines 318."), 9:11-19 ("alert" is based on "vulnerability data 210"). In other words, Mohanty uses the agent-based techniques that I describe above to detect vulnerabilities in real-time, and the "prioritization scores" relied upon by Dr.

Stavrou are generated by Dr. Stavrou's "continuous" monitoring of the target VMs via installed agents. Ex.1002 (Stavrou Decl.), ¶¶180-181; Ex.1075 (Mohanty), 1:19-22, 6:54-58. But Veselov, in contrast, relies on "snapshots" that provide data at a point in time (Ex.1007 (Veselov), 3:23-27) for its assessments, using an approach that Mohanty criticizes, suggesting that such data would become "obsolete" within hours or minutes (Ex.1075 (Mohanty), 1:5-22. Accordingly, a person of ordinary skill would understand that Veselov's system should not issue "alerts" on the vulnerabilities it detects.

129. And, I have reviewed objective evidence pertaining to the nonobviousness of the '809 patent, which demonstrates that Wiz's combination is driven entirely by hindsight with an eye on the challenged claims and lacks any motivation to make the combination.

REDACTED

130. Accordingly, Dr. Stavrou has not shown that independent claims 1, 16, and 19 were obvious as of January 28, 2019.

* * *

131. For at least the foregoing reasons, it is my opinion that Petitioner and Dr. Stavrou have not shown that any of the independent claims of the '809 patent—

claims 1, 16, and 19—would have been obvious. And it follows that Petitioner and Dr. Stavrou have not shown that any of the claims depending on those independent claims, including but not limited to the dependent claims challenged in Petitioner's Ground 1 (claims 2-10, 12-15, 17-18, and 20-23), would have been obvious. Thus, Petitioner and Dr. Stavrou have not shown that any of claims 1-10 and 12-23 of the '809 patent would have been obvious.

e. Dependent claims 15 and 22

132. Dependent claims 15 and 22 depend from claims 1 and 19, respectively, and recite “wherein determining the location of the snapshot of at least one virtual disk further includes querying a cloud management console of the cloud computing environment for the location of the snapshot and the *location of the virtual disk.*”

133. Dr. Stavrou asserts that “an environment API or web interface to communicate with the computing resource service provider hosting the VM” constitutes the “cloud management console of the cloud computing environment.” Ex.1002 (Stavrou Decl.) ¶246. I first note that “computing resource service provider” is not typically used in the art to refer to a cloud computing environment. Nevertheless, although the examples of Figures 1 and 3A-3B include a computing resource service provider and APIs, the examples of Figures 5A-5B do not include such a computing resource service provider or APIs. But with respect to the

Figures 5A-5B environment, where there is no API or computing resource service provider, Veselov discloses that a snapshot of the VM can be “obtained.” Ex.1007 (Veselov) 18:18-21. And in Figures 3A, 3B, 4, 5A, and 5B, Veselov teaches that the volume image may be *provided to* a hypervisor or service that launches an assessment machine, or may be read by the hypervisor to launch the assessment machine. Ex.1007 (Veselov), 12:37-51; 18:16-21. A person of ordinary skill in the art therefore would understand that APIs are not needed to obtain snapshots in Veselov, and therefore can be, and would be, obtained by means other than APIs even in Figures 1 and 3A-3B, and certainly within Figures 5A-5B, where no API even exists.

134. Dr. Stavrou further asserts that determining the snapshot’s location “at a minimum would have suggested that such a query would also determine the virtual disk’s location.” Ex.1002 (Stavrou Decl.) ¶232. For example, Dr. Stavrou asserts that the snapshot and virtual disk can have the “same general location (e.g., being located in the same computing environment).” Ex.1002 (Stavrou Decl.) ¶250. Even if true, that would not suggest, however, actually determining the virtual disk’s location. For example, Veselov explicitly shows in Figure 1 that snapshots are located separately from any VM in a “snapshot storage service” external to the computing resource service provider. And Veselov states that “computing systems may be located in a single geographic location or located in

multiple, distinct geographic locations” (Ex. 1007 (Veslov), 1:12-15) so the same “general location” does not necessarily mean that a determination of a snapshot location obviously leads to the determination of a virtual disk’s location. And, it is especially true that Dr. Stavrou has not shown that this limitation would have been obvious, where Dr. Stavrou fails to provide any reason why the Veslov-Mohanty combination would *benefit* from such information.




FIG. 1

Ex.1007 (Veslov) Fig. 1

135. At best, Veslov identifies a virtual machine but does not teach or suggest determining a location of a virtual disk. Identifying a VM does not

determine a location of a virtual disk. A virtual disk's location cannot be deterministically solely established or determined based on the location of a virtual machine because virtualization allows for decoupling of the location of the virtual machine and a virtual disk. The fact that Veselov's systems and methods work with and from the identification of VMs (not identifying, let alone determining locations of, virtual disks) belies Dr. Stavrou's assertions about the virtual disk location determination of claims 15 and 22. Dr. Stavrou offers no explanation for why a person of ordinary skill in the art would have changed those systems and methods to instead determine locations of *virtual disks* of a VM. Veselov's system has no need to provide a "UI/API" for the user *to locate a virtual disk*. Dr. Stavrou asserts that this was a "common, routine task for a POSA" but none of the references he cites supports such common, routine knowledge. Ex.1002 (Stavrou Decl.) ¶¶252-253.

136. As the '809 patent explains and shows, virtual disks in cloud computing platforms associated with VMs can be located apart from those VMs. *See* Section VI.A, *see also, e.g.*, Ex.1001 ('809 patent), Fig. 1B, 3:55-60 ("[S]torage 117 emulat[ing] virtual discs for the VMs" "may be part of the server" hosting VMs, or, as depicted in Figure 1B, separate from but "connected to" that server). And, as Dr. Stavrou admitted at his deposition, "Virtual disks are different from virtual machines"—"[t]hey are definitely not synonymous." Ex.2077, (1st

Stavrou Tr.) 97:13-98:10. Thus, the identification of a VM does not necessarily or obviously determine the location of a virtual disk of that VM. Further, identifying a VM does not require or render obvious using an API to “query” such location at least because an identification means the user does not need to query for a location.

137. Dr. Stavrou cites Ex.1015 (VMware vSphere) but does not explain why those disclosures are applicable to Veselov, e.g., Dr. Stavrou makes no effort to explain that VMware vSphere is used in Veselov or that its teachings would apply to Veselov. Ex.1002 (Stavrou Decl.) ¶252. Nor does he explain how VMware vSphere makes it apparent that “querying a cloud management console was a common, routine task for a POSA” and “would have been a natural initial step”; nothing in VMware vSphere supports these bare contentions. Instead, Dr. Stavrou points to numerous pages throughout the VMware vSphere document with no effort to link those disparate portions together or how they relate to specific claim limitations at issue or his position that querying was a “natural initial step.” Dr. Stavrou also quotes Exhibit 1031 (Ex.1002 (Stavrou Decl.) ¶252) as emphasizing the importance of knowing the location of “assets” but omits that this reference is referring to “networked devices and software,” not virtual disks in a cloud computing system. *See* Ex.1031 (NIST IT Asset Mgmt), 1 (“track the location and configuration of networked devices and software across an enterprise”), 2 (“Our example solution has the following benefits: enables faster

responses to security alerts by revealing the location, configuration, and owner of a device.”). This reference says nothing about querying a cloud management console of the cloud computing environment for the location of a virtual disk.

138. Dr. Stavrou also cites to his discussion of claim 10, which recites “further comprising mitigating a potential cyber vulnerability posing a risk to the protected virtual cloud asset.” Ex.1002 (Stavrou Decl.), ¶253. But his analysis for that section does nothing to address the issues I discuss above.

139. Dr. Stavrou further cites to his discussion of claim 12, which recites “wherein determining the location of the snapshot of at least one virtual disk further includes determining a virtual disk allocated to the protected virtual cloud asset.” Like his citation to claim 10, his analysis for claim 12 does nothing to address the issues I discuss above. Ex.1002 (Stavrou Decl.), ¶250.

140. Additionally, in his discussion of claim 12, Dr. Stavrou relies on a handful of references other than Veselov, including Ex.1051 (Workstation Pro), Ex.1053 (Prahlad), and Ex.1020 (Sancheti), purportedly to show that “snapshot generation typically included determining virtual disks allocated to the VM that will be included in the snapshot.” Ex.1002 (Stavrou Decl.), ¶¶232-233. But Dr. Stavrou does not identify anything in those references that would explain why a person of ordinary skill in the art would have been motivated to change Veselov’s systems and methods to identify a location of a virtual disk, as claimed, as opposed

to locating a snapshot or identifying a VM generally. Veselov's systems and methods lack such determinations, and I have been informed and understand that, in an obviousness analysis, merely collecting disparate disclosures from a handful of references, to purportedly show that a person of ordinary skill in the art *could have* made the claimed inventions because such disclosures allegedly show that each piece of a claim was known or routine or commonplace is not enough. Rather, I understand that obviousness analyses look to what a person of ordinary skill in the art *would have* been motivated to make or do. In my opinion, nothing in the other references Dr. Stavrou cites show that a person of ordinary skill in the art would have been motivated to change Veselov's systems and methods, at all, let alone to modify those systems and methods as claimed in the '809 patent's claims.⁷

⁷ Dr. Stavrou's tertiary references do not demonstrate that the claimed features are routine or commonplace. Dr. Stavrou cites three pages of Ex.1015 (VMware vSphere) but again does not explain why those disclosures are applicable to Veselov. Similarly, the disclosure Dr. Stavrou cites from the Workstation Pro reference (Ex.1051) simply discloses that data from a disk is included by default in a snapshot of a virtual machine, and makes no mention of a step of determining a location of a virtual disk. Ex.1051 (Workstation Pro), 47, 125. The VirtualCenter reference (Ex.1052), relied on in Ex.1002 (Stavrou Decl.) ¶233, shows that there

f. Summary of Ground 1 Opinions

141. For at least the foregoing reasons, it is my opinion that Wiz and Dr. Stavrou have not shown that any of the independent claims of the '809 patent—claims 1, 16, and 19—as well as dependent claims 15 and 22, would have been

may be an identification for a virtual machine location, but does not demonstrate determining a location of a virtual disk. Ex.1052 (VirtualCenter), 445-46. The disclosures Dr. Stavrou cites from Prahlad (Ex.1053) does not demonstrate determining a location of a virtual disk, nor is it relied upon in Dr. Stavrou's grounds. Ex.1053 (Prahlad), ¶¶77-82, 87-88, 92. Nor does Ex.1020 (Sancheti) demonstrate routineness of determining locations of virtual disks—in fact, its disclosure (which is not part of Dr. Stavrou's proposed grounds), supports that merely identifying a VM does not determine a location of a virtual disk. *See* Ex.1002 (Stavrou Decl.), ¶233 (quoting Sancheti 21:9-41). And Sancheti does not disclose determining a location of virtual disks for the purpose of security assessments. Finally, Dr. Stavrou's citation to Ex.1074 also does not explain how that reference determines a location of virtual disks, or even if it did, how that determination would be incorporated with Veselov to teach this limitation. Rather, Dr. Stavrou's citation to page 12 of that reference suggest it only describes “querying the location of the VM” and not the virtual disk, as required.

obvious. It follows that Wiz and Dr. Stavrou have not shown that any of the claims depending on those independent claims, including but not limited to the dependent claims challenged in Wiz's ground 1 (claims 2-10, 12-15, 17-18, and 20-23), would have been obvious. Thus, Wiz and Dr. Stavrou have not shown that any of claims 1-10 and 12-23 of the '809 patent would have been obvious.

B. Ground 2: Claims 1-10 and 12-23 Would Not Have Been Obvious Over Veselov, Mohanty, and Czarny

142. Dr. Stavrou sets forth Ground 2 to “build[] upon” Ground 1 with Czarny's alleged disclosure of “reporting asset risk levels and using binary and hash comparison for purposes of matching installed application[s] against a list of known vulnerable applications to detect vulnerabilities.” Ex.1002 (Stavrou Decl.), ¶259. Therefore, for at least the same reasons Dr. Stavrou fails to demonstrate obviousness of claims 1-10 and 12-23 in his Ground 1, it is also my opinion that Dr. Stavrou fails to demonstrate obviousness of claims 1-10 and 12-23 in his Ground 2.

143. It is my opinion that Dr. Stavrou has not demonstrated obviousness of Claims 1-10 and 12-23 for the following additional reasons.

1. A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov, Mohanty, and Czarny and Would Not Have Had a Reasonable Expectation of Success in Doing So

144. Dr. Stavrou asserts that a person of ordinary skill would have combined Veselov, Mohanty, and Czarny. Ex.1002 (Stavrou Decl.) ¶¶260-268. A person of ordinary skill would not have been motivated to combine Veselov, Mohanty, and Czarny. As I explain above, Veselov and Mohanty involve security assessments of virtual machines via analyses of snapshots as data files or via analysis of instantiated VMs (Veselov) or via agent-based monitoring (Mohanty). *See* Sections VII.A-VII.B. But in contrast, Czarny is directed to assessments of electronic devices. Czarny (Ex.1084), Abstract, 1:39-56. Czarny does not include any reference or application to virtualized resources and also does not involve either snapshot analyses, instantiated VM analyses, or agent-based analyses of VMs, as described in Veselov and Mohanty, respectively. Czarny (Ex.1084), Abstract, 1:39-56.

145. Czarny performs its security assessments via received “product binary data” that includes “the complete binary level files of [] software products.” Ex.1084 (Czarny), 4:7-13. Dr. Stavrou states that “Veselov and Mohanty both teach matching installed applications against applications in known lists” and that “using binary comparison and hash-based matching in the context of similar matching” would be obvious. Ex.1002 (Stavrou Decl.) ¶264. But Veselov and

Mohanty do not implement any such matching using electronic device “binary data.” Therefore, “using binary and hash-based matching” in Veselov and Mohanty would not “merely involve using [Czarny’s] matching techniques” as Dr. Stavrou states (Ex.1002 (Stavrou Decl.) ¶264) and would instead necessitate technology and techniques that Veselov and Mohanty do not contemplate. And Dr. Stavrou’s belief that Czarny would benefit “Veselov’s disclosure of reporting assessment results and Mohanty’s discussion of determining risk levels” (Ex.1002 (Stavrou Decl.) ¶¶265-266) ignores that, as I describe above, the combination of Veselov and Mohanty do not render obvious alerts. And Dr. Stavrou does not explain how Czarny’s electronic product data would be implemented in Veselov’s analysis of a virtual disk snapshot or duplicate virtual machine or in Mohanty’s agent-based monitoring of a virtual machine.

146. Dr. Stavrou also asserts that “Czarny’s discussion of reporting a risk level for the asset as a whole is consistent with Veselov’s general disclosure of reporting assessment results and Mohanty’s discussion of determining risk levels and sending reports in various formats with various types of information.” Ex.1002 (Stavrou Decl.) ¶265. But this reasoning and paragraph as a whole only allege, at best, that the references are similar, and not that a person of ordinary skill in the art would have been motivated to combine the teachings of Czarny with Veselov and Mohanty. Additionally, Dr. Stavrou does not provide any reasoning

as to why he would incorporate Czarny's discussion, as alleged, when he earlier asserts that Mohanty teaches this subject matter in Ground 1. That is, for example, Dr. Stavrou does not provide any reasoning as to why one would abandon Mohanty's purported teaching of the same subject matter in favor of Czarny's teachings. Likewise, Dr. Stavrou does not establish how the distinct pieces of Veselov, Mohanty, and Czarny would fit together and how such a combination would work, either in theory or in practice.

147. It is accordingly my opinion that a person of ordinary skill would not have been motivated to combine Veselov, Mohanty, and Czarny and would not have a reasonable expectation of success in doing so.

2. Even If Combined, Veselov, Mohanty, and Czarny Do Not Render Claims 1-10 and 12-23 Obvious

a. Limitations 1.8, 16.8, and 19.8

148. Dr. Stavrou asserts that Czarny supplements the combination in Ground 1 "by expressly disclosing determining and reporting a risk level based on the detected vulnerabilities." Ex.1002 (Stavrou Decl.) ¶270. But Dr. Stavrou does not explain how such a risk level would amount to a "risk of the [] asset to the cloud computing environment" as required by the challenged claims. I also note that Czarny's risk assessment does not relate to determining a risk level to a cloud environment for purposes of generating alerts because Czarny instead attempts to determine whether target device should be granted or denied access to a secure

environment based on the target device's "security vulnerability level." Ex.1084 (Czarny), 7:46-67. Dr. Stavrou agrees with this characterization that Czarny determines a risk before the target asset is allowed to become a part of the respective environment. Ex.2172 (2nd Stavrou Tr.), 244:11-245:1, 245:18-246:5. Thus, any information communicated in Czarny does not relate to a "determined risk" that is "of the protected virtual cloud asset to the cloud computing environment" where the "protected virtual cloud asset is instantiated in the cloud computing environment" as the claims require, as the target device in Czarny would not be a "protected virtual cloud asset" and would not be "instantiated in the cloud computing environment."

149. I also note that, even when viewing Czarny in combination with Veselov and Mohanty, there is still a gap between the challenged claims and the asserted combination of prior art references. For example, the references in combination still do not teach "alerts" or alerts that are "prioritized" or "a risk level" to a cloud computing environment. It therefore would not have been obvious to apply Czarny's teachings with the Ground 1 combination, and doing so would not lead a person of ordinary skill to the claimed subject matter.

b. Claims 6-7 and 17-18

150. Dr. Stavrou asserts that claims 6-7 and 17-18 are rendered obvious by the combination of Veselov, Mohanty, and Czarny. Ex.1002 (Stavrou Decl.)

¶¶272-279. But as I explain above, claims 1 and 16, from which claims 6-7 and 17-18 depend, are not obvious. Therefore, claims 6-7 and 17-18 are not obvious for at least this reason.

C. Ground 3: Claim 11 Would Not Have Been Obvious Over Veselov, Mohanty, and Hutchins

151. Dr. Stavrou sets forth Ground 3 to address only the additional limitations of claim 11. Ex.1002 (Stavrou Decl.), ¶280. But Dr. Stavrou challenges claims 1 and 10, on which claim 11 depends, based on only Veselov and Mohanty. *See* Ex.1002 (Stavrou Decl.), ¶104. Accordingly, for at least the same reasons Dr. Stavrou fails to demonstrate obviousness of claims 1 and 10, it is also my opinion that Dr. Stavrou fails to demonstrate obviousness of claim 11.

152. It is my opinion that claim 11 is also not obvious for the following additional reasons.

1. A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov, Mohanty, and Hutchins and Would Not Have Had a Reasonable Expectation of Success in Doing So

153. Claim 11 depends from claim 10 and recites “wherein mitigating a potential cyber vulnerability includes one of: blocking traffic from untrusted networks to the protected cloud asset, halting operation of the protected virtual cloud asset, or quarantining the protected cloud asset.” Dr. Stavrou proposes to

combine Veselov, Mohanty, and Hutchins to render claim 11 obvious. Ex.1002 ¶¶280-289. I disagree with Dr. Stavrou's analysis.

154. A person of ordinary skill in the art would not have been motivated to combine Hutchins with Veselov and Mohanty. As an initial matter, I discussed above in relation to Ground 1 that a person of ordinary skill in the art would not have been motivated to combine Veselov and Mohanty in the proposed manner, and that discussion applies here to Ground 3 because Ground 3 assumes the combination of Veselov and Mohanty. *See* Section VIII.A.

155. Additionally, even assuming that Veselov and Mohanty are combined, a person of ordinary skill in the art would not have combined them with Hutchins to render the claims obvious. Hutchins uses a VM snapshot to replicate a VM to temporarily offload computationally-intensive tasks while both the original VM and replicated VM are active. *E.g.*, Ex.1070 (Hutchins) ¶53. But Hutchins's approach is different than, and would not have been combined with, Veselov's approach.

156. Veselov discloses scanning operations that are focused on a "scannable volume" from a snapshot captured at a specific point in time. *E.g.* Ex.1007 (Veselov) 10:1-2. In contrast, as discussed above, Hutchins is directed towards offloading resource-intensive (i.e., virus scanning) applications to replicated VMs. Ex.1070 (Hutchins) ¶53; *see also* Section VII. Application of

Hutchins to Veselov would require additional resources beyond those provided for in Veselov. At best, Hutchins's approach has similarities to Veselov's Figs. 3A/3B examples, which, as discussed above, does not analyze snapshots themselves, and Dr. Stavrou does not discuss or suggest combining Hutchins's analysis of replicated VMs with Veselov's snapshot-based assessments. To the contrary, a POSA looking at Hutchins and Veselov would have been led away from analyzing a snapshot, because to the extent there is a commonality between the references, the commonality is using a snapshot to instantiate a virtual machine, which does not disclose the claimed "analyzing the snapshot" for the reasons discussed previously. Dr. Stavrou does not identify any reason why, in a combination, a POSA would have or could have applied snapshot analyses with Hutchins techniques that rely on duplicate VM analyses to determine whether to quarantine a VM.

2. Even If Combined, Veselov, Mohanty, and Hutchins Do Not Render Claim 11 Obvious

157. Because claim 11 depends on claims 1 and 10, it is not obvious over Veselov, Mohanty, and Hutchins for the same reasons that claims 1 and 10 are not obvious in view of Veselov and Mohanty. *See* Section VIII.A.

158. I also note that Hutchins does not describe "vulnerabilities" and further does not describe any remedial steps in response to determined vulnerabilities. Ex.1070 (Hutchins). Hutchins' quarantining features, on which

Dr. Stavrou's analysis relies, involve quarantining a VM only upon the detection of a virus. But viruses are not the same as the claimed potential cyber vulnerabilities. Dr. Stavrou's own references make clear that detection of a virus or "malware detection" is in a different "field[] or industry[y]" than a "[v]ulnerability assessment." Ex.1084 (Czarny), 1:6-7. A virus is a significant cyber threat such as "[a] computer program that can copy itself and infect a computer without permission or knowledge of the user" and may "corrupt or delete data on a computer, use e-mail programs to spread itself to other computers, or even erase everything on a hard disk." Ex.2175 (<https://csrc.nist.gov/glossary/term/virus>). A "vulnerability" in contrast is a "[w]eakness in an information system, system security procedures, internal controls, or implementation that could be exploited or triggered by a threat source." Ex.2176 (<https://csrc.nist.gov/glossary/term/vulnerability>). Thus, Dr. Stavrou's statement that "Hutchins teaches mitigating a potential cyber vulnerability" (Ex.1002 (Stavrou Decl.) ¶288) is wrong, as Hutchins only teaches mitigating a potential virus, which his own evidence supports is different than a "vulnerability." Accordingly, Dr. Stavrou has not shown that claim 11 would have been obvious over Veselov, Mohanty, and Hutchins.

D. Ground 4: Claim 11 Would Not Have Been Obvious Over Veselov, Mohanty, Hutchins, and Czarny

159. Dr. Stavrou sets forth Ground 4 to address only the additional limitations of claim 11. Ex.1002 (Stavrou Decl.), ¶280. But Dr. Stavrou challenges claims 1 and 10, on which claim 11 depends, based on Veselov, Mohanty, and Czarny (Ground 2), and based on Veselov, Mohanty, and Hutchins (Ground 3). *See* Ex.1002 (Stavrou Decl.), ¶104. Accordingly, for at least the same reasons Dr. Stavrou fails to demonstrate obviousness of claims 1 and 10, it is also my opinion that Dr. Stavrou fails to demonstrate obviousness of claim 11.

160. It is my opinion that claim 11 is also not obvious for the following additional reasons.

1. A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov, Mohanty, Czarny, and Hutchins and Would Not Have Had a Reasonable Expectation of Success in Doing So

161. Claim 11 depends from claim 10 and recites “wherein mitigating a potential cyber vulnerability includes one of: blocking traffic from untrusted networks to the protected cloud asset, halting operation of the protected virtual cloud asset, or quarantining the protected cloud asset.” Dr. Stavrou proposes to combine Veselov, Mohanty, Czarny, and Hutchins to render claim 11 obvious. Ex.1002 ¶¶280-289. I disagree with Dr. Stavrou’s analysis.

162. A person of ordinary skill in the art would not have been motivated to combine Hutchins with Veselov, Mohanty, and Czarny. As an initial matter, I discussed above in relation to Grounds 1-3 that a person of ordinary skill in the art would not have been motivated to combine Veselov, Mohanty, and Czarny, and Veselov, Mohanty, and Hutchins in the proposed manner, and that discussion applies here to Ground 4 because Ground 4 assumes the combinations of Veselov, Mohanty, and Czarny, and Veselov, Mohanty, and Hutchins. *See* Sections VIII.A-VIII.C.

163. I also note that Czarny discloses that the detection of “malware” requires “immediate response to eliminate a threat” including “to quarantine” but that “[w]ith a positive vulnerability assessment, on the other hand, the computer device can usually continue to operate without concern for a threat to the computer device, since a malicious event is not necessarily imminent.” Ex.1084 (Czarny), 1:8-15. Thus, Czarny does not support or teach “mitigating a potential cyber vulnerability” as recited by quarantining a computer device and does not render obvious claim 11 when taken with Veselov, Mohanty, and Hutchins.

2. Even If Combined, Veselov, Mohanty, Czarny, and Hutchins Do Not Render Claim 11 Obvious

164. I address the merits of Claim 11 in Ground 3 above, and my analysis in Ground 3 applies equally here.

165. Also, because claim 11 depends on claims 1 and 10, it is not obvious over Veselov, Mohanty, Czarny, and Hutchins for the same reasons that claims 1 and 10 are not obvious in view of Veselov and Mohanty with or without Czarny. *See* Sections VIII.A-VIII.C.

IX. Objective Indicia of Non-obviousness

166. I have considered objective evidence that supports that a person of ordinary skill in the art would not have considered the claims of the '809 patent to be obvious, as discussed below. It is my opinion that, even assuming that Petitioner and Dr. Stavrou made a threshold showing of obviousness (they have not for the reasons explained above), there is objective evidence of non-obviousness that shows that the claimed combination of features would not have been obvious to a person of ordinary skill in the art.

A. Orca's SideScanning™ Uses The '809 Patent Inventions

167. I have reviewed Section III of Exhibit 2074, which is the Declaration of Avi Shua. I understand that Mr. Shua is a co-founder of Orca and the named inventor of the '809 patent. In that section of his declaration, Mr. Shua maps Orca's SideScanning™ to the claims of the '809 patent. I have further reviewed the Exhibits cited in that Section of Exhibit 2074. Mr. Shua's mapping is consistent with my review of those Exhibits. Thus, I adopt and incorporate by reference Section III of Exhibit 2074. Based on my review of those Exhibits and

Mr. Shua's mapping, it is my opinion that Orca's SideScanning™ practices the claims of the '809 patent.

B. Long-Felt, Unresolved Need

168. Between 2010-2014, there was a major shift of commercial computing services to cloud computing platforms. Ex.2076, 5-6 (describing the “shift to cloud” beginning in 2010-2014). Amazon introduced the first commercial cloud in 2006, followed by Google in 2008, and Microsoft in 2010. Ex.2134, *The History of Cloud Computing*, Cyber Magazine (Nov. 28, 2021) at 2. By 2014, cloud adoption had “take[n] off” with huge year-over-year increases in adoption. Ex.2134 at 2 (“Cloud adoption takes off” around 2012). As the '809 patent describes, “[o]rganizations [] increasingly adapted their applications to be run from multiple cloud computing platforms.” Ex.1001 ('809 patent), 1:23-24. Virtualization technologies such as virtual machines, software containers, and serverless functions helped enable this shift to commercial cloud computing. See Ex.1001 ('809 patent), 1:39-42.

169. As applications moved to the cloud, security lagged behind. “In the early day of cloud, [security] vendors rushed to deploy their existing solutions into this new environment.” Ex.2076 at 4. But even as vendors realized that cloud computing posed new security challenges, new solutions had a number of drawbacks. See Ex.2136, Cloud Security Report, 2 (2018) (“As more workloads

are shifting to the cloud, cybersecurity professionals remain concerned about security of data, systems, and services in the cloud.”); Ex.1029, *Overview of Attacks on Cloud Computing* (April 2012) at 1 (“[S]ecurity is the biggest problem of” cloud computing systems). As the ’809 patent describes, the cloud security industry developed a number of methods for protecting cloud computing infrastructures, including: (1) “traffic inspection”, (2) using “network scanner[s] deployed out of path,” (3) “traffic monitor[ing],” and (4) using “agents installed in each server,” but these methods suffered from multiple disadvantages. Ex.1001 (’809 patent), 1:47-2:17.

170. Generally, these problems related to (1) limited visibility and (2) limited coverage. The problem of limited visibility includes problems with inspecting encrypted traffic, scanning servers that either don’t respond or risk exposing sensitive data, and inability to detect misconfiguration based on only observing, e.g., network traffic. Ex.1001 (’809 patent), 1:52-2:9. The problem of limited coverage, which was especially true with agent-based systems, includes difficulties associated with installing agents on every VM. Thus, while an agent may have better visibility into a VM’s security posture, the difficulty in installing agents throughout a platform left big gaps in coverage. See Ex.2076, 11 (“Traditional agents struggle with complex deployment and maintenance.”); Ex.1001 (’809 patent), 2:16-17 (“[I]nstalling agents in a large data center may take

months.”); Ex.2103, *Security Agents Don't Belong In Your Cloud* (Aug. 18, 2021) at 3 (“Customers tell me that agent-based approaches are only able to cover up to 50% of the cloud real estate even after several months of deployment.”).

171. Thus, even by 2019, “[a]lthough cloud security has evolved over the years and is becoming more mature, the endless journey to obtain the right level of visibility and control over the cloud workloads [was] still a challenge.” Ex.2137, Diogenes, *The Quest for Visibility and Control in the Cloud*, ISSA Journal, 1 (March 2019); *see also* Ex.2136, Cloud Security Report, 2 (“Organizations continue to adopt cloud computing at a rapid pace,” but “traditional security tools are often not suited for the challenges of dynamic, virtual and distributed cloud environments.”), 3 (“Eighty-four percent [of respondents] say traditional security solutions either don’t work at all in cloud environments or have only limited functionality.”).

172. Accordingly, there was a long-felt, unmet need for a cloud security solution that could provide both visibility and coverage to adequately address “undetected vulnerabilities.” Ex.1001 (’809 patent), 1:49-55, 1:60-62, 2:4-9, 2:18-19.

173. Orca’s claimed inventions solved this need. As discussed in the following sections, Orca’s claimed approach, embodied in its SideScanning™ product, was a “genuinely genius innovation” such that cloud security posture

management “can be defined by pre and post-Orca.” Ex.2076 at 7. I have also reviewed articles by Wiz and Palo Alto Networks that describe agentless scanning, and its benefits, in a manner that closely matches the claimed invention. For example, I have reviewed Exhibit 2138, which appears to be a Wiz article from September 28, 2021, titled “Agents Are Not Enough.” Ex.2138 at 1. In this article, Wiz describes: “[A]gentless deep scanning approaches [] achieve a level of comprehensiveness in terms of coverage and features that is not possible with agent-based approaches due to the tradeoff between performance and impact to the environment.” Additionally, I have reviewed Exhibit 2135, which appears to be a Palo Alto Networks article from November 16, 2021, titled “Palo Alto Networks Shifts Left with Prisma Cloud 3.0.” Ex.2135 at 1. In this article, Palo Alto Networks describes: “agentless scanning... works off cloud provider APIs and disk snapshots” to “extend[] coverage to both active *and dormant* virtual machines”. Ex.2135 at 3.

174. “Orca launched the first version of **agentless side scanning**, which allowed deep visibility into cloud workloads” “without the need to install agents.” Ex.2076 at 11.

Orca Security made a *big leap* in cloud security with the introduction of its patented agentless SideScanning technique in 2019. This innovation introduced workload visibility without needing security teams to install an agent on endpoints or installing software on

workloads for assessments. They operate out-of-band by accessing cloud workloads at the runtime storage level, gathering crucial data through the cloud provider's APIs, while scanning disk snapshots.⁸

Ex.2076 at 31.

175.

REDACTED

⁸ See '809 patent, cls. 1, 16, 19 (limitations 1.1-1.5, 16.1-16.5, 19.1-19.5).

REDACTED

C. Orca's claimed inventions overcame industry skepticism to become widely accepted--and widely praised—in the industry.

176. As discussed below, I have reviewed objective evidence that there was industry skepticism of Orca's claimed inventions, but that they were quickly accepted and widely praised by the industry. This evidence further supports that skilled artisans at the time of the invention would not have viewed Orca's claims as obvious.

177. A 2024 report on “the Future of Cloud Security” states “[t]here is no doubt, the pivotal moment was the introduction of agentless scanning by companies like Orca and Wiz”—of course, “Orca launched the first version of agentless side scanning” and, as shown below, Wiz later copied Orca. Ex.2076 at 8, 11. It was a “turning point – but anyone who tells you they knew it at the time is lying.” *Id.* at 8.

178. After Orca “introduced the concept of agentless side scanning,” it was “difficult” for “older providers to come to terms with [its] success.” Ex.2076 at 7. That 2024 report states that these providers were “skeptics” because “from a security perspective, agentless scanning was nothing compared to what attacks

⁹ See '809 patent, cls. 1, 16, 19 (limitations 1.7-1.8, 16.7-16.8, 19.7-19.8).

they were able to detect and protect.” *Id.* “They missed however that security teams were not yet in the weeds of these architectures,” so before Orca, they “didn’t know what they didn’t know.” *Id.* 7. And in 2021, Palo Alto Networks released an article about how Orca’s solutions “fall short.” Ex.2106, *How Agentless Approaches Sideline Cloud Workload Protection Strategies*, Palo Alto Networks (Feb. 23, 2021) at 1 (“Agentless Solutions Fall Short”). By November, Palo Alto Networks announced their own agentless scanning solution. Ex.2135, *Palo Alto Networks Shifts Left with Prisma Cloud 3.0* at 2 (announcing “Agentless Security for Flexible Cloud Workload Protection” that “Provide[s] visibility into an organization’s cloud workload and application risks”).

179.

REDACTED

REDACTED

180.

REDACTED

181. Eventually, “everyone needed [agentless scanning] in order to stay competitive.” Ex.2076, 7; Ex.2103, *Security Agents Don’t Belong In Your Cloud* (Aug. 18, 2021) (“[A] new ‘agent-less’ approach to cloud security has recently become possible due to cloud-native APIs now available from CSPs. This

approach provides security in a fundamentally different way.”). As I mentioned above, Palo Alto Networks quickly reversed course on its public criticisms of the approach, and Wiz entered the market a year after Orca with the same technology—copied from Orca.

D. Extensive Industry Praise Supports Non-Obviousness of the Claims

182. I have reviewed evidence of industry praise of Orca’s inventions that further shows that skilled artisans did not view (and a person of ordinary skill in the art would not have viewed) the inventions as obvious. This evidence includes numerous industry awards, including:

- CRN’s “20 Coolest Tech Startups of 2020.” Ex.2119, *The 20 Coolest Tech Startups of 2020*, CRN (Nov. 19, 2020) at 9-10. CRN touted “Orca’s SideScanning technology” which “automatically assesses the security state of every discovered asset” to provide “complete visibility.” *Id.*
- Forbes’s “Top 20 Cybersecurity Startups to Watch in 2021.” Ex.2120, *The Top 20 Cybersecurity Startups to Watch in 2021 Based on Crunchbase*, Forbes (Nov. 20, 2020) at 6-7. Forbes found Orca “noteworthy for its innovative approach to providing instant-on, workload-deep security” for cloud computing systems, emphasizing Orca’s “patent-pending SideScanningTM technology” that “reads cloud configuration and workloads’ runtime block storage out-of-band.” *Id.*

- SC Media's Best Cloud Workload Protection Solution.
Ex.2123, *SC Award Winners 2023 Orca Security – Best Cloud Workload Protection Solution*, SC Media (Aug. 21, 2023)
("After years of dealing with agents, then seeing how easy it is to install and use Orca, I knew that its agentless approach was both a major innovation and a game changer")

183.

REDACTED

REDACTED

REDACTED

REDACTED

185. I have also reviewed evidence of public praise of Orca's inventions, including:

- Exhibit 2076, which I have been informed is a 2024 report published on Substack by Software Analyst Cyber Research (founded by Francis Odum) and James Berthoty that heaped praise on Orca's inventions specific to the claimed combination of features of the '809 patent, such as using cloud computing platform APIs to determine locations of virtual disks and virtual disk snapshots and facilitate snapshot analyses. Select quotes include:

Orca: Introduced the concept of **agentless side scanning**, which allowed for deep visibility into cloud environments, and workloads, without the need to install agents on each workload. CSPM can be defined by pre and post-Orca: the ability to view workload data without an agent was a **genuinely genius** innovation - utilizing the nature of cloud APIs being accessible to vendors.

Id. at 7

Orca Security made a big leap in cloud security with the introduction of its patented **agentless SideScanning** technique in 2019. This **innovation** introduced workload visibility without needing security teams to install an agent on endpoints or installing software on workloads for assessments.

Id. at 31

They operate out-of-band by accessing cloud workloads at the runtime storage level, gathering crucial data through the cloud provider's APIs, while scanning disk snapshots. Their technology evaluates block storage bytes and reconstructs file systems, scans operating systems, and applications in a read-only mode, ensuring that virtual assets remain unaffected. This **unique** method allows companies to deploy Orca across their entire cloud environment without worrying about performance impacts or operational disruptions. For security teams, this means getting visibility into engineering workloads, without needing to bother them at all along the way.

***Id.* at 31**

Orca's SideScanning technology has proven **highly effective** for vulnerability assessments and scanning with a fast and unified database for attack paths especially for non-container assets. While Orca's advantages in agentless remains a key strength, it also has shown some limitations as cloud security teams start looking for more runtime data and visibility into their workloads. Critically, agentless scanning has always had difficulties being applied to ephemeral workloads, where what's on disk is changing minute by minute.

***Id.* at 31**

- **Quick deployment and rapid-time to value:** Agentless solutions are easier to deploy since they do not require installing software on each workload. They **leverage cloud APIs to snapshot underlying instance hard disks**, copy them into the security vendors account, where they're hooked up on the vendor's side and scanned. This creates an **elegant "like magic" approach** - where you get full visibility into the workload, without any instrumentation.

***Id.* at 11-12**

- Exhibit 2118, which is titled "Vendor Landscape and TCO Analysis for Achieving Cloud-Wide, Workload-Deep Security and Compliance for Public Cloud Environments." Ex.2118 at 1. This document describes "Orca uses a patent-

pending technique called SideScanning™, ... which unobtrusively scans the run-time block storage layer—out of band—to extract valuable data about the environment and the assets communicating in it.” Ex.2118 at 5. The report states that a Senior Cloud Security Engineer from a company called Sisense described Orca’s solution as “unique in that it locates vulnerabilities with precision and delivers tangible, actionable results—without having to sift through all the noise.” *Id.* at 4.

186. In summary, Orca’s extensive praise for the claimed combination of features in the ’809 patent is objective evidence that persons of ordinary skill in the art would not have viewed Orca’s claims as obvious at the time of invention.

E. Evidence Shows Wiz Copied Orca’s SideScanning™

187. I have reviewed evidence that Wiz copied Orca’s SideScanning™. This evidence, which I discuss below, further demonstrates that a person of ordinary skill in the art, at the time of the invention, would not have found the ’809 patent’s claims obvious.

188. REDACTED

REDACTED

191. In the following section I discuss additional evidence that Wiz copied Orca's product.

1. Evidence shows Wiz had access to Orca's invention

192. I have been informed that Wiz had access to Orca's invention at least because Orca was founded before Wiz and released its SideScanning™ before

Wiz. Specifically, I have been informed that Orca and Wiz were both startup companies, that Orca was formed before Wiz, and that Orca presented its SideScanning™ technology to Wiz's founder when he (Wiz's founder) was still at Microsoft in May 2019.

193.

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

212. As the above-discussed evidence shows, Wiz accessed and copied Orca's claimed inventions.

2. Evidence shows Wiz's Workload Scanning is substantially similar to Orca's SideScanning™

213. Wiz's Cloud Security Platform uses what Wiz calls "workload scanning," which is substantially similar to Orca's SideScanning™.

214. REDACTED

215. REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

216. I have also reviewed some of Wiz's public documents, which show a substantial similarity with Orca's SideScanning™, including:

- Exhibit 2026, which I am informed is a webpage on Wiz's website. This webpage touts the benefits of agentless scanning and states that agentless analysis is "based on two fundamental principles" including "Privileged access to customer's cloud environment *via APIs*" and "*Snapshot scanning*." Ex.2026 at 3. "By connecting to cloud APIs, they can discover new resources as they're created, without requiring manual installation of an agent process. This maximizes security

coverage from day one, improving the visibility of security issues.” Ex.2025 at 8.

What is agentless security?

Agentless security provides visibility into the threats in your environments, without requiring the installation of software-driven agents.

Under the agent-based model, each of your hosts must run a monitoring process that collects data from the host's environment and sends it to your security service. Agentless security removes this requirement by having the service collect data itself, using cloud provider APIs and metadata.

Agentless analysis approaches are based on two fundamental principles:

- Privileged access to customer's cloud environment via APIs
- Snapshot scanning

Privileged access assigned to the security vendor enables discovery of all resources and services used in the cloud environment. This information is then processed to determine the list of workloads to be scanned.

Ex.2026 at 3

Advantages of agentless security

Agentless security solves most of the problems associated with agents. Instead of running an agent in each of your environments, agentless services sit *outside* your resources. They collate security information by monitoring data provided by cloud APIs and infrastructure services. This model presents several compelling advantages for security teams and administrators.

Simple, automatic coverage

Agentless platforms automatically monitor the resources in your cloud provider accounts. By connecting to cloud APIs, they can discover new resources as they're created, without requiring manual installation of an agent process. This maximizes security coverage from day one, improving the visibility of security issues.

Ex.2026 at 8

- Exhibit 2115, which I have been informed is a webpage on Wiz's website, and is titled "Securing clouds, securely." Ex.2115 at 1. This page describes "agentless scanning approaches" as "cloud native and API-based." Ex.2115 at 1. This page also describes the Wiz approach to agentless scanning, which it calls "workload scanner." Ex.2115 at 3. For example, in a "Full SaaS Wiz-hosted workload scanner," the scanner has "limited permissions to run the scan – including creating, sharing, and deleting snapshots." Ex.2115 at 3. I note that the "limited permissions" here refer at least in part to a policy file I discussed in the previous section.

Full SaaS Wiz-hosted workload scanner — In this model, the Wiz workload scanner is hosted in a Wiz cloud environment, with limited permissions to run the scan — including creating, sharing, and deleting snapshots. During the scan, the scanner has read-only access to the scanned disk, so some environments are a better fit for the hosted model (highly regulated, sensitive data, etc). This model offers several benefits including ease of deployment, scalability, and lack of ongoing management.

Ex.2115 at 3

This is one of the reasons that agentless scanning approaches to securing the cloud — approaches that are cloud native and API-based — have taken over the cloud security market, replacing legacy agent-based practices. At Wiz, we see this firsthand: from the pain of customers who want certainty and context, to the relief they feel when they gain that much-needed visibility. As the CISO of Fox said, "Wiz came into the picture to allow us to feel secure and confident in how fast we're moving, even as our cybersecurity challenges keep changing."

Ex.2115 at 1

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

225. In the discussion that follows, I discuss additional exhibits I have reviewed that further support the substantial similarity between Wiz and Orca's

products as it relates to specific limitations of the claims of the '809 patent. ^{REDACTED}

Notably, the below is not an infringement analysis, but rather a discussion of how Wiz's product is substantially similar to Orca's product with respect to the below limitations.

a. Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.1-1.2, 16.1-16.2, 19.1-19.2.

226. Limitations 1.1-1.2 recite *“determining, using an API or service provided by the cloud computing environment, a location of a snapshot of at least one virtual disk of a protected virtual cloud asset, wherein the protected virtual cloud asset is instantiated in the cloud computing environment”* and *“accessing, based on the determined location and using an API or service provided by the cloud computing environment, the snapshot of the virtual disk.”* Wiz's Workload Scanner is substantially similar to Orca's SideScanningTM in relation to limitations 1.1-1.2, 16.1-16.2, and 19.1-19.2 of claims 1, 16, and 19 of the '809 Patent.

227. For example, similar to Orca's SideScanningTM, Wiz's Workload Scanner is based on assigning privileges to access and scan an environment using cloud APIs. For example, in Exhibit 2115, which I discussed above, Wiz characterizes “Agentless scanning approaches” as being “based on assigning

privileges to security vendors to access and scan an environment using cloud APIs.” Ex.2115 at 1. Wiz then “outline[s]” its “approach to secure agentless scanning” by describing the “Wiz workload scanner.” Ex.2115 at 2. I therefore understand Wiz’s description of “Agentless scanning approaches” in this document to encompass its own Workload Scanner.

Agentless scanning approaches are based on assigning privileges to security vendors to access and scan an environment using cloud APIs. When selecting an agentless scanning technology, it’s important to assess the privileges assigned to the security vendor, to ensure **least privileges** and secure access to your data. In addition, it’s important to understand exactly which data is being shared with the vendor and how that data is shared across cloud environments.

Ex.2115 at 1.

228. Wiz’s Workload Scanner thus includes an onboarding process where a client environment provides Wiz’s Workload Scanner with “privileges” that allow access to an environment and to scan the environment using cloud APIs. *Id.* This establishes an interface between a client organization’s cloud computing environment implemented by one or more cloud computing platforms, such as AWS, Azure, or Google Cloud Platform (GCP), and Wiz’s Workload Scanner.

229. **REDACTED**

REDACTED

230. Wiz's Workload Scanner uses the interface between a client's environment and Wiz's security components to use an API to determine a location of a snapshot of a virtual disk of a protected virtual cloud asset.

The Wiz workload scanner

Traditional cloud security posture management tools are focused on the configuration layer only. When we started Wiz, we understood that in order to analyze an end-to-end risk, we would need to analyze *multiple* layers, including network, identity, data, and the workloads themselves. This is why we built a workload scanner that is designed for the cloud and uses cloud APIs, instead of relying on older agent-based approaches.

Ex.2115 at 2

231. REDACTED

REDACTED

REDACTED


- b. Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.3-1.5, 16.3-16.5, 19.3-19.5.**

232. Limitations 1.3-1.5 recite “*analyzing the snapshot of the at least one virtual disk by matching installed applications with applications on a known list of vulnerable applications*” and “*determining, based on the matching, an existence of a plurality of potential cyber vulnerabilities*” and “*correlating the determined potential cyber vulnerabilities with a network location of the protected virtual cloud asset.*” Wiz's Workload Scanner is substantially similar to Orca's

SideScanning™ in relation to limitations 1.3-1.5, 16.3-16.5, and 19.3-19.5 of claims 1, 16, and 19 of the ’809 Patent (reproduced in full above).

233. For example, Wiz’s workload scanner “analyze[s] an end-to-end risk” by “analyz[ing] multiple layers, including network, identity, data, and the workloads themselves.” Ex.2115 at 2. Additionally, Wiz performs “vulnerability analysis” based on the snapshot.

Stage 6: Vulnerability analysis



Example vulnerability dashboard that prioritizes issues by contextual severity

After scanning, the tool will generate a comprehensive list of identified vulnerabilities, order them based on severity, filter away false positives, and provide options for remediation.

Ex.2165 at 8

234. Additionally, Wiz detects and prioritizes a catalog of “known exploited vulnerabilities.” See Ex.2147.

What is the CISA KEV Catalog?


CISA, the Cybersecurity & Infrastructure Security Agency, released a **binding operational directive** to reduce the significant risk of known exploited vulnerabilities. In recognition that vulnerabilities that have been previously exploited are a frequent vector for malicious cyber actors, CISA prioritizes these vulnerabilities as the biggest threats that must be aggressively remediated. Since November 2021, CISA has managed a catalog of known exploited vulnerabilities that carry significant risk, along with remediation requirements. Federal agencies are compelled to patch these vulnerabilities within the timeframe defined by CISA.

Currently, there are more than 352 vulnerabilities in the **catalog**, covering vulnerabilities in operating systems (e.g. Windows, Linux), proprietary applications (e.g. SolarWinds products, Zoho ManageEngine ServiceDesk), and open source projects (e.g. Apache HTTP Server, Log4j).

The **latest vulnerability**, a Microsoft Win32k privilege escalation vulnerability (CVE-2022-21882), was added on February 4th, 2022 and required remediation action within 14 days. Even though the CISA directive is binding for US federal agencies specifically, it is quickly becoming a best security practice for public and private sectors as well.

Ex.2147 at 1-2

235. And, Wiz's platform can "[i]dentify and stop lateral movement" because it "provides direct visibility into your cloud environment, prioritizes risks, and offers remediation guidance." Ex.2177 (Wiz, What Is Lateral Movement, <https://www.wiz.io/academy/what-is-lateral-movement#identify-and-stop-lateral-movement-with-wiz-82>). For example, Wiz uses a "Security Graph" to "correlate dozens of disparate signals into a single prioritized risk assessment." Ex.2178 (Wiz, Attack Path Analysis, <https://www.wiz.io/blog/uniting-builders-and-defenders-a-new-vision-for-cloud-security>).

**Ex.2178.**

- c. **Wiz’s Workload Scanner is substantially similar to Orca’s SideScanning in relation to limitations 1.6-1.8, 16.6-16.8, 19.6-19.8**

236. Limitations 1.6-1.7 recite “*using the determined plurality of potential cyber vulnerabilities and the network location of the protected virtual cloud asset to determine a risk of the protected virtual cloud asset to the cloud computing environment*” and “*prioritizing, by the determined risk, the plurality of potential cyber vulnerabilities*” and “*reporting the determined plurality of potential cyber vulnerabilities as alerts prioritized according to the determined risk.*” Wiz’s Workload Scanner uses the features recited in limitations 1.6-1.8, 16.6-16.8, and 19.6-19.8 of claims 1, 16, and 19 of the ’809 Patent (reproduced in full above).

237. For example, Wiz's Workload Scanner performs "Risk-based prioritization" of vulnerabilities that purports to "mitigate the amount of alert fatigue you experience."

3. Risk-based prioritization




Example visualization of the risk-based context that a vulnerability scanner should provide

Choose a tool that provides risk-based prioritization of vulnerabilities, considering factors such as severity, exploitability, and asset criticality alongside elements such as external exposure, **cloud entitlements**, secrets, misconfigurations, and malware presence. A tool with this functionality can correlate your vulnerabilities that have numerous risk factors to mitigate the amount of alert fatigue you experience.

Ex.2165 at 14-15

238. Wiz's Workload Scanner reports issues as alerts.

Stage 6: Vulnerability analysis



Example vulnerability dashboard that prioritizes issues by contextual severity

After scanning, the tool will generate a comprehensive list of identified vulnerabilities, order them based on severity, filter away false positives, and provide options for remediation.

Ex.2165 at 8

F. Commercial Success

239. I have reviewed evidence that Orca's SideScanning™ was commercially successful, including investment funding of more than \$600 million and REDACTED Ex.2140, *Orca 2020 Recap*, 2; Ex.2112, *Upstart Orca Security Raises \$210M In Google-Backed Funding*, CRN, 1-2; Ex.2139, *Orca Security Raises \$340M To Scale Platform R&D, Channel*, CRN, 2; Ex.2312. I have also reviewed evidence that Orca's innovations created a new market for agentless Cloud Security Posture Management. *See* Ex.2076, 7

(“CSPM can be defined by pre and post-Orca”); *see also* REDACTED

240. I have also reviewed evidence that Wiz’s Cloud Security Platform, which uses what it calls “workload scanning” and which practices the claims (as shown above) has been commercially successful as well, including claims that it became the fastest-growing software company in history, and articles describing valuations of the company in the tens of billions of dollars. Ex.2116, *\$100M ARR in 18 months: Wiz becomes the fastest growing software company ever*, 2 (Aug. 10, 2022); Ex.2141, *Cloud Security Startup Wiz, Now Valued at \$10B, Raises \$300M*, 2; Ex.2142, *Wiz walks away from \$23 billion deal with Google, will pursue IPO* (July 22, 2024).

241. It is my opinion that the commercial success of both Orca’s and Wiz’s products has nexus to the claims of the ’809 patent. For example, both systems practice the claims of the ’809 patent as discussed above. REDACTED

X. CONCLUSION

All statements made herein of my own knowledge are true, and all statements made on information and belief are believed to be true. Further, I am aware that these statements are made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001. I declare under penalty of perjury that the foregoing is true and correct.

Executed on Friday, the 9th day of May, 2025, in Boston, Massachusetts.

A handwritten signature in black ink, appearing to read "D. Kaeli", with a stylized, overlapping initial "D" and "K".

Dr. David Kaeli