

Filed on behalf of: Orca Security Ltd.

Entered: May 9, 2025

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

WIZ, INC.,
Wiz,

v.

ORCA SECURITY LTD.,
Patent Owner.

IPR2024-01190
Patent 11,740,926

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

Table of Contents

I.	Introduction	4
II.	Qualifications	5
III.	Summary Of Materials Reviewed And Considered	8
IV.	Legal Standards.....	9
V.	Level Of Skill In The Art And Perspective Applied In This Declaration	13
VI.	Background	16
A.	Overview of the '926 patent	16
B.	Claim Construction.....	23
1.	“Analyzing the Snapshot”	23
VII.	Overview of the Asserted References.....	27
A.	Overview of Veselov	27
B.	Overview of Mohanty.....	36
C.	Overview of Ranum.....	39
D.	Overview of Seo	40
E.	Overview of Hutchins.....	42
VIII.	The Claims Of The '926 Patent Would Not Have Been Obvious ...	44
A.	Ground 1: Claims 1, 5-10, and 12-15 Would Not Have Been Obvious Over Veselov and Mohanty	45
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	45
2.	Even If Combined, Veselov and Mohanty Do Not Disclose or Suggest Multiple Limitations in the Claims	73
B.	Ground 2: Claims 1-10 and 12-15 Would Not Have Been Obvious Over Veselov, Mohanty, and Ranum.....	111

1.	A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov, Mohanty, and Ranum and Would Not Have Had a Reasonable Expectation of Success in Doing So	112
2.	Even If Combined, the Veselov/Mohanty/Ranum Combination Does Not Render Obvious Limitations 1.4, 14.4, or 15.4.....	114
3.	Even If Combined, the Veselov/Mohanty/Ranum Combination Does Not Render Dependent Claims 2 and 4 Obvious.....	119
C.	Ground 3: Claim 3 Would Not Have Been Obvious Over Veselov, Mohanty, and Seo	121
1.	A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov, Mohanty, and Seo and Would Not Have Had a Reasonable Expectation of Success in Doing So	122
D.	Ground 4: Claim 3 Would Not Have Been Obvious Over Veselov, Mohanty, Seo, and Ranum	124
1.	A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov, Mohanty, Seo, and Ranum and Would Not Have Had a Reasonable Expectation of Success in Doing So	124
E.	Ground 5: Claim 11 Would Not Have Been Obvious Over Veselov, Mohanty, and Hutchins	125
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	126
2.	Even If Combined, Veselov, Mohanty, and Hutchins Do Not Render Claim 11 Obvious	128
F.	Ground 6: Claim 11 Would Not Have Been Obvious Over Veselov, Mohanty, Hutchins, and Ranum.....	128
IX.	Objective Indicia of Non-obviousness.....	128
A.	Orca’s SideScanning™ Uses The ’926 Patent Inventions ..	129
B.	Long-Felt, Unresolved Need	129
C.	Orca’s claimed inventions overcame industry skepticism to become widely accepted--and widely praised—in the industry.	135

D. Extensive Industry Praise Supports Non-Obviousness of the Claims 140

E. Evidence Shows Wiz Copied Orca’s SideScanning™ 146

1. Evidence shows Wiz had access to Orca’s invention 148

2. Evidence shows Wiz’s Workload Scanning is substantially similar to Orca’s SideScanning™ 172

F. Commercial Success 200

X. Conclusion..... 201

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 1-15 of U.S. Patent No. 11,740,926 (the “’926 patent,” Ex.1001), in support of Patent Owner’s Response to the Petition for *Inter Partes* review of the ’926 patent, which I understand was filed by Wiz, Inc. (“Wiz”) and accompanied by a Declaration of Dr. Angelos Stavrou (the “Stavrou Declaration”), among other exhibits. I further understand that on January 22, 2025, this *Inter Partes* review (IPR) was instituted (Paper 8).

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 '926 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, probative of 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 '926 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 against cyber threats, 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 '926 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 '926 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 ’926 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 '926 patent

35. The '926 patent is titled “Techniques for Securing Virtual Machines by Analyzing Data for Cyber Threats” and generally relates to innovative systems and methods for securing virtual cloud assets, such as VMs. Ex.1001 ('926 patent), Abstract, cls. 1, 14, 15.

36. The '926 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 ('926 patent), 1:20-30. The '926 patent also provides context for “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 ('926 patent), 1:31-33.

37. The '926 patent discusses problems that existed with trends in the cybersecurity field. For example, the '926 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 ('926 patent), 1:42-46. And the '926 patent also

explains ways the industry had tried to deal with such issues. For example, the '926 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 ('926 patent), 1:46-2:14. As the '926 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 ('926 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 ('926 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 ('926 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 ('926 patent), 2:14-16. "For example, installing agents in

a large data center may take months.” Ex.1001 ('926 patent), 2:16-17. The '926 patent's systems and method “overcome the[se] deficiencies.” Ex.1001 ('926 patent), 2:18-19.

38. The '926 patent aims to solve such problems. Ex.1001 ('926 patent), 3:35-36, 3:47-63. Figure 1B of the '926 patent, reproduced below, illustrates a diagram of a cloud computing platform to which the '926 patent's techniques for scanning for potential cyber threats that do not represent vulnerabilities in machines:




FIG. 1B

39. The cloud computing platform of the '926 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 ('926 patent), 3:48-51. 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 ('926 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 ('926 patent), 3:55-63. In the example shown in Figure 1B, a “high-speed connection” such as an “optic fiber” connects the storage with the server. Ex.1001 ('926 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 ('926 patent), 4:16-21. The security system is configured to perform the steps of the '926 patent's invention and may be separate from the client environment, as shown in Figure 1B, or part of the client environment. Ex.1001 ('926 patent), 4:1-17. The challenged claims of the '926 patent demonstrate the embodiment where the security system interfaces with the client environment through APIs. Ex.1001 ('926 patent), cls. 1, 14, 15.

40. The claims of the '926 patent are specifically directed to “analyzing the snapshot of at least one virtual disk to determine the existence of a plurality of potential cyber threats, each cyber threat based on data stored on the virtual disk, wherein the data includes at least one of:

- unencrypted sensitive data,
- unencrypted system credentials,
- weak passwords,

- weak encryption schemes,
- disabled Address Space Layout Randomization,
- boot record manipulation,
- suspicious definitions,
- services to be run on startup,
- personally identifiable information,
- data in application logs indicating that the protected virtual cloud asset accessed personally identifiable information,
- data in application logs indicating that the protected virtual cloud asset accessed a computer containing personally identifiable information, or
- at least one change in at least one area of the virtual disk, as compared to an earlier point in time”

Ex.1001 ('926 patent), cls. 1, 14, 15. I note that these types of potential cyber threats based on data that are expressly recited in the claims of the '926 patent are described in the specification as part of particular embodiments where the security system “is configured to detect cyber threats that do *not* represent vulnerabilities.” '926 Patent 6:21-23, 6:23-32 (“For example, the security system 140 may detect and alert on sensitive data not being encrypted on the logical disk, private keys found on the disks, system credentials stored clearly on the disk, risky

application features¹ (e.g., support of weak cipher suites or authentication methods), weak passwords, weak encryption schemes, a disable address space layout randomization (ASLR) feature, suspicious manipulation to a boot record, suspicious PATH, LD_LIBRARY_PATH, or LD_PRELOAD definitions, services running on startup, and the like”), 4:45-47 (describing how the security system of the disclosed embodiments can be used to detect “[1] vulnerabilities,” and separately “[2] *other* cyber threats related to execution VM 119.”).

41. 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 (’926 patent), 4:50-56), generate reports on detected cyber threats, including, “for example, as alerts” (Ex.1001 (’926 patent), 7:53-55), and mitigate potential cyber threats based on data stored on the virtual disk. Ex.1001 (’926 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 (’926 patent), 8:1-6.

42. In order to accomplish the scanning and detection operations I discuss above, the ’926 patent provides a series of steps. First, “a request” to “scan a VM”

¹ I note that “risky application features” is not one of the potential cyber threats recited in the ’926 patent claims.

for “cyber threats” “is received.” Ex.1001 ('926 patent), 7:14-15, 7:19-20. Then, the security system queries for and receive a location of a snapshot of a virtual disk. Ex.1001 ('926 patent), 4:57-60. For example, the security system queries a cloud management console for the location of a snapshot. Ex.1001 ('926 patent), 4:57-60. The snapshot can be a copy of the VM's virtual disk or a copy of the VM's disk file at a point in time. Ex.1001 ('926 patent), 4:58-59. After the snapshot is located, it is accessed by the security system. Ex.1001 ('926 patent), 5:3-4. The security system then analyzes the snapshot to determine the existence of potential cyber threats, without installing agents on the server or the VM, and without interaction with or information from the VM. Ex.1001 ('926 patent), 5:20-25, 7:37-41. The potential cyber threats are based on data stored on the virtual disk such as unencrypted sensitive data, unencrypted system credentials, weak passwords, weak encryption schemes, disabled Address Space Layout Randomization, boot record manipulation, suspicious definitions, services to be run on startup, personally identifiable information, data in application logs indicating that the protected virtual cloud asset accessed personally identifiable information, data in application logs indicating that the protected virtual cloud asset accessed a computer containing personally identifiable information, or at least one change in at least one area of the virtual disk, as compared to an earlier point in time. Ex.1001 ('926 patent), 6:21-55, cls. 1, 14, 15. The security system

may then determine a risk associated with each of the determined potential cyber threats. Ex.1001 ('926 patent), 7:53-67. This risk may include the criticality of the virtual machine in the cloud computing environment in which it exists. Ex.1001 ('926 patent), 7:53-67. The security system then prioritizes potential cyber threats associated with the protected virtual cloud asset based on the determined risk and generates reports on detected potential cyber threats and sensitive data, including, “for example, as alerts” that can be prioritized based on risk. Ex.1001 ('926 patent), 6:40-48, 7:53-67.

B. Claim Construction

43. In addition to the explicit construction discussed below, I have applied the plain and ordinary meanings of the claim terms in the '926 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”

44. Claims 1 and 14 recite “analyzing the snapshot of the at least one virtual disk to determine the existence of a plurality of potential cyber threats, each cyber threat based on data stored on the virtual disk, wherein the data includes at least one of ...” Similarly, claim 15 recites “analyze the snapshot of the at least one virtual disk to determine the existence of a plurality of potential cyber threats, each cyber threat based on data stored on the virtual disk, wherein the data includes

at least one of...” Dr. Stavrou labels these as elements 1.4, 14.4, and 15.4. *See, e.g.,* Ex.1002 (Stavrou Decl.), ¶147.

45. 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 14, as well as “analyzing the snapshot” in claim 15, as encompassing both direct analysis of the snapshot data and analysis of a VM instantiated from a 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.

46. 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,” 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.

47. A person of ordinary skill in the art would not understand the “[analyze/analyzing] the snapshot” limitations in the ’926 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 the 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 ’926 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 (’926 patent), 4:60-61-59. The specification also distinguishes analyzing a snapshot by a “security system 140 to detect vulnerabilities” (Ex.1001 (’926 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 (’926 patent), 5:56-63).

48. In fact, as I discuss further below, it appears that Dr. Stavrou “assume[d] that ‘analyzing the snapshot’ also encompasses analysis of a VM instantiated from the snapshot,” as he was “asked” to (Ex.1002 (Stavrou Decl.), ¶82), 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. *See* Section VIII.A.1.

VII. Overview of the Asserted References

A. Overview of Veselov

49. 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.

50. Veselov provides five embodiments in Figures 1, 3A, 3B, 5A, and 5B, each having distinct features and techniques. For example, Figure 1 implements security assessments of VMs using a snapshot storage service to store and “obtain the snapshot data,” where the snapshot storage is located separately from the virtual machines being assessed. Ex.1007 (Veselov), Fig. 1, 7:34-8:15. In the Figure 1 embodiments, Veselov uses “API calls” to specify the “requestor, the user associated with the virtual machine 112, the on-demand storage service 108, and/or

the particular virtual machine 112.” Ex.1007 (Veselov), 6:34-38. In Figures 3A and 3B, Veselov’s security assessments do not include a “snapshot storage service” and implement a “duplicate” or “assessment” VM instantiated with a snapshot of target VM data to detect vulnerabilities. Ex.1007 (Veselov), 11:5-9, Figs. 3A-3B. Veselov’s 3A and 3B embodiments include embodiments where the target and assessment VM are in the same environment (Fig. 3B) and in different environments (Fig. 3A). And Veselov’s Figures 5A and 5B embodiments, in contrast to Figures 3A and 3B, analyze a snapshot as a “data file” rather than implement an instantiated VM. Ex.1007 (Veselov), Figs. 5A-5B, 16:56-18:15. Veselov’s Figures 5A and 5B comprise embodiments that make use of a test environment (Fig. 5A) and embodiments that do not use a test environment (Fig. 5B). All of Veselov’s embodiments are directed towards detecting *vulnerabilities*, which, as I discuss below, are not the same as the cyber threats required by the ’926 patent claims.




FIG. 1

51. As seen in Figure 1 above, Veselov discloses 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, and provide input. Ex.1007 (Veselov), 4:29-41.

52. 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.

53. 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 counsels 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.


54. 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.

55. 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.



56. 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 provide input and communicate instructions to the scanning service (shown inside the computing resource service provider's environment as no. 310 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.

57. 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.



58. 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.

59. 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.

60. 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

61. 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.

62. 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, or detect cyber threats based on the particular types of data recited in the '926 patent claims. 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).

63. 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.





FIG. 3

64. 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 (vulnerability assessment) 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 Ranum

65. Ranum is titled “System and Method for Strategic Anti-Malware Monitoring” and describes “active network scanning and passive network monitoring to provide strategic anti-malware monitoring in a network.” Ex.1093 (Ranum) Abstract. I understand that Ranum is relied upon by Dr. Stavrou and Wiz for its alleged disclosures of using “disclosure of threats based on particular types of data, reporting threat risk levels, and comparing computed cryptographic hashes.” *See, e.g.*, Ex.1002 (Stavrou Decl.) ¶¶246-265 (relying on Ex.1093 (Ranum), 4:36-40, 4:48-57, 5:49-57, 6:63-7:2, 10:42-63, 13:41-14:3, 17:21-30, 17:37-48, 19:1-9, 19:59-62).



Ex.1093 (Ranum), Fig. 2.


66. Ranum’s techniques involve active scanners, as shown in Figure 2 above (element 210), that “communicate packets or other messages within the network to detect new or changed information describing various routers, hosts, servers, or other devices in the network.” Ex.1093 (Ranum), 4:31-36. Ranum’s active scanners perform “credentialed audits or uncredentialed scans” and “obtain information that may then be analyzed” to “identify potential vulnerabilities in the network.” Ex.1093 (Ranum), 4:36-40.

D. Overview of Seo

67. Seo is titled “Security Enhancement Method and Electronic Device Therefor.” Ex.1094 (Seo) Cover. I understand that Seo is relied upon by Dr.

Stavrou and Wiz for its alleged disclosure of “identifying files without a corresponding installation process.” Ex.1002 (Stavrou Decl.) ¶¶266-280 (relying on Ex.1094 (Seo), ¶¶2, 11, 27, 125, 127, 134, 169, 172, claim 10).

FIG. 3



Ex.1094 (Seo), Fig. 3

68. Seo appears to be primarily directed to “an artificial intelligence (AI) system which mimics recognition and judgment of the human brain.” Ex.1094 (Seo), ¶2; *see also* Ex.1094 (Seo), ¶¶3-8. For example, Seo explains that “rule-based smart systems” are “gradually being replaced with deep-learning-based AI systems” and that “the aid of artificial intelligence technology” is needed “to block

encryption before [a] malicious program encrypts data.” Ex.1094 (Seo), ¶¶3-9. Accordingly, Seo describes an “electronic device and method” that are “robust against attacks on encryption-related vulnerabilities as detection of an encryption algorithm based on if artificial intelligence technology is enabled are provided.” Ex.1094 (Seo) Abstract.

69. Seo teaches that the “malicious program[s]” it purports to guard against are “ransomware or malware.” I note, however, that “ransomware” and “malware” are not the same as the claimed cyber threats based on specific types of data stored on a virtual disk, which I enumerate above.


70. Seo’s AI-driven invention includes “receiving vulnerability information of an encryption algorithm from an external device” and then “storing the vulnerability information” in a “vulnerability list,” shown in Figure 3 above. Ex.1094 (Seo), ¶26. Seo also describes an “encryption code identification model” where “output data” may identify “a vulnerability in the vulnerability list.” Ex.1094 (Seo), ¶26. Upon making this identification, Seo teaches that the system denies the loading of executable code. Ex.1094 (Seo), ¶26.

E. Overview of Hutchins

71. 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.” Ex.1070 (Hutchins) Abstract. “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.” Ex.1070 (Hutchins) Abstract.

72. 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.” Ex.1070 (Hutchins) ¶39. 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 in step 605. Ex.1070 (Hutchins) ¶40. The main memory is then cloned in step 607, and then both virtual machines can run from the cloned state. Ex.1070 (Hutchins) ¶¶40-41. A script can then be run on the replicated VM to perform, e.g., data mining, backup, virus scanning, etc. Ex.1070 (Hutchins) ¶44. The replicated VM is then “usually discarded.” Ex.1070 (Hutchins) ¶45.

**FIG. 6**

Hutchins further discloses that if a virus is detected in the replicated VM, the original VM can be quarantined. Ex.1070 (Hutchins) ¶53.

VIII. The Claims Of The '926 Patent Would Not Have Been Obvious

73. Wiz and Dr. Stavrou challenge claims 1-15 of the '926 patent based on three 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).

74. In my opinion, none of the claims of the '926 patent would have been obvious over Wiz's and Dr. Stavrou's prior art, for reasons I explain below.

A. Ground 1: Claims 1, 5-10, and 12-15 Would Not Have Been Obvious Over Veselov and Mohanty

75. In Ground 1 of the Petition, Wiz and Dr. Stavrou challenge claims 1, 5-10, and 12-15 of the '926 patent based on the combination of Veselov and Mohanty. *See* Ex.1002, Stavrou Decl. §XII.

76. In my opinion, Wiz and Dr. Stavrou have failed to show that any of claims 1, 5-10, and 12-15 of the '926 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

77. 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.), ¶¶107-117. 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 '926 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 '926 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.**

78. 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 used in the '926 patent claims. *See* § VI.B.1.

79. 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.

80. 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.


81. 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.

82. 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.

83. 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.

84. 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.



85. 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

86. 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 POSA 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 POSA Would Not Have Been Motivated To Combine Veselov with Mohanty With a Reasonable Expectation of Success**

87. 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 disparages security assessment techniques that provide vulnerability assessments of network systems in the past, such as snapshot-based systems, 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.

88. 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 the execution of a real-time anti-virus scanner and have indicated that significant

89. Dr. Stavrou asserts that Veselov and Mohanty perform the “same, predictable function, with predictable results” and that a POSA “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.), ¶111. 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.

90. 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:

“[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

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.”).

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, 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.

91. 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.

92. 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), 6:45-63. 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.”).

93. 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 disparagement 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 agent-reliant teachings of Mohanty would not have been combined, as the implementations of each reference do not suggest combining their teachings.

94. 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.) ¶111. 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.), ¶ 111. 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.

95. Additionally, I note that a VM's snapshot is, as the '926 patent states, "a copy of the machine's virtual disk (or disk file) at a given point in time." Ex.1001 ('926 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).

96. Additionally, as I detail more below in the section detailing objective indicia of non-obviousness, before the '926 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, 1-2.

97. 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.

98. 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 determining to "use 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 VM 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 opinion that agent-based and

³ Dr. Stavrou appears to indicate that Veselov's "duplicate VM" analysis does not "use[] an agent installed on the target VM, nor do they require an agent installed on the assessment VM." Ex.1002 ¶153. 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

snapshot-based 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.) ¶111, because it is contradicted by Veselov and other evidence.

99. 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.1036 (Wysopal); Ex.1040 (Ammons); Ex.1044 (Giakouminakis); Ex.1045 (Canzanese); Ex.1046 (Crabtree); Ex.1048 (Price); Ex.1056 (Chari); Ex.1063 (Hibbert); Ex.1076 (Milner); Ex.1077 (Khalid); Ex.1093 (Ranum); Ex.1095 (Burdett). I have

(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.

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, Canzanese, Crabtree, and Khalid do not implement Mohanty's agent-based scanning and therefore provide no support for why a person of ordinary skill in the art would combine Mohanty's agent-based teachings with the teachings in Veselov of analyzing a snapshot. *See generally* Ex.1014, Ex.1034, Ex.1045, Ex.1046, Ex.1075; Ex.1077. 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 analysis⁴ techniques. 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 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.

⁴ To be clear, where I refer to Veselov's snapshot analysis techniques or Veselov's techniques of analyzing a snapshot, that is meant to refer to Figures 5A-5B of Veselov. Veselov's Figures 3A-3B and Veselov's disclosures about analyzing a duplicate VM, including with security agents, do not analyze snapshots.

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. Dr. Stavrou thus has not met his burden to demonstrate that its purported combination renders the challenged claims obvious. I address Dr. Stavrou's additional references below.

100. 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's image-based approach is therefore not applicable to Veselov's analysis of a snapshot as a data file and further does not support combining Mohanty's scanning techniques with those of Veselov.

101. Chari (Ex.1056) which Dr. Stavrou refers to for its "labels [of] sensitive data" (Ex.1002 (Stavrou Decl.), ¶109) relies on interactions with the target VM to label and sanitize sensitive data but does not ever *detect* sensitive data or cyber threats (Ex.1056 (Chari), ¶37) and in any event suffers the same intrusive-based drawbacks as seen in Mohanty. Chari also fails to inform as to why a person of ordinary skill would combine the agent-based analysis of Mohanty with Veselov's snapshot-based techniques.

102. The techniques of Khalid, 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 similar to 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.

103. Ranum (Ex.1094), as I discuss further below, relies on network monitors and active scanning, which requires active processes to for its sensitive data disclosures (Ex.1093 (Ranum), 13:54-14:3) on which Dr. Stavrou relies for his analysis (Ex.1002 (Stavrou Decl.) ¶110) and therefore is not applicable to Veselov. Further, Ranum provides no indication as to why a person of ordinary skill would modify Veselov’s analysis of a snapshot as a data file with Mohanty’s VA scanning techniques.

104. Burdett (Ex.1095), which Dr. Stavrou relies on for its discussion of “content that may compromise customer’s infrastructure” (Ex.1002 (Stavrou Decl.) ¶110), does not implement any snapshot-based techniques. Similarly, Milner, which Dr. Stavrou relies on for its discussion of “verifying the security of software packages” (Ex.1002 (Stavrou Decl.) ¶109), lacks any mention or implementation of snapshot analyses. Burdett and Milner are therefore

inapplicable to Veselov's processes and further fail to inform a person of ordinary skill as to how or why Mohanty would benefit Veselov or why a person of ordinary skill in the art would have combined Mohanty and Veselov.

105. 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 not have combined their teachings. Dr. Stavrou does not address this.

106. 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. Mohanty's techniques work 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.

107. 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." Ex.1002 (Stavrou Decl.), ¶113. 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.), ¶113. And, as I discuss above, Mohanty sets forth an agent-based scanning system such that it need not rely on point-in-time captures that might become "obsolete." Ex.1075 (Mohanty), 1:14-17. But Veselov's snapshots are point-in-time captures, 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, for its part, does not mention or suggest agents for its target machines. 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.

108. 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 the most dangerous." Ex.1002 (Stavrou Decl.), ¶114. 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.I. 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.

109. Relatedly, Dr. Stavrou's opinion that "applying Mohanty's prioritization techniques would similarly benefit snapshot-based assessments such as those described in Veselov" does not address the underlying differences between Mohanty's agent-based techniques and Veselov's analysis of a snapshot. Ex.1002 (Stavrou Decl.) ¶ 115. 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, as described 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 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.

110. 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.), ¶117) 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.

111. 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.

112. 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.

113. 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 Claims

114. 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 1.4, 14.4, and 15.4

115. Limitations 1.4, 14.4, and 15.4 recite “[analyzing/analyze] the snapshot of the at least one virtual disk to determine the existence of a plurality of potential cyber threats, each cyber threat based on data stored on the virtual disk,

wherein the data includes at least one of: unencrypted sensitive data, unencrypted system credentials, weak passwords, weak encryption schemes, disabled Address Space Layout Randomization, boot record manipulation, suspicious definitions, services to be run on startup, personally identifiable information, data in application logs indicating that the protected virtual cloud asset accessed personally identifiable information, data in application logs indicating that the protected virtual cloud asset accessed a computer containing personally identifiable information, or at least one change in at least one area of the virtual disk, as compared to an earlier point in time.” Dr. Stavrou fails to show that Veselov and Mohanty disclose or suggest these limitations. Ex.1002 (Stavrou Decl.), ¶¶147-168.

116. I first note that in the following portion of the claim language, the reference to “the data” refers to the immediately preceding “based on data stored, such that limitations 1.4, 14.4 and 15.4 require analyzing a snapshot to detect cyber threats based specifically on “at least one of” the “listed types of data” in the claim: “the snapshot of the at least one virtual disk to determine the existence of a plurality of potential cyber threats, each cyber threat based on **data stored on the virtual disk**, wherein **the data** includes at least one of ...” I understand Dr. Stavrou agreed to this in his deposition too:

3 Q. Okay. In that first part of the limitation
 4 where it says "the data," you understand that to
 5 refer to data stored on the virtual disk; right?
 6 MR. ARGENTI: Objection. Objection. Form.
 7 A. Where do you see that?
 8 Q. So, like, the third -- so the third line of
 9 your paragraph 147 is quoting -- is quoting the
 10 claim language; right?
 11 A. Yes.
 12 Q. Okay. And it says, "each cyber threat based
 13 on data stored on the virtual disk, wherein the data
 14 includes."
 15 And so I'm just asking, you understand that
 16 the word "the data" -- like, I don't think "data"
 17 appears anywhere else in the claim. It refers to --
 18 A. That's why I got confused.
 19 (Unreportable crosstalk.)
 20 A. (Inaudible.) Yes. (Inaudible.)

See Ex.2172, 158:3-158:20

21 Q. Okay. Got it.
 22 And so to satisfy the claim language, the --

1 each of the detected cyber threats must be based on
 2 one of the listed types of data. For example,
 3 unencrypted sensitive data, unencrypted system
 4 credentials, et cetera, et cetera; right?
 5 A. It says "at least one" of them, yes.

See Ex.2172, 158:21-159:5

117. Based on my review of Dr. Stavrou's declaration, he only opines that the prior art purportedly taught detecting cyber threats based on one of the following five types of data in the claim: (1) unencrypted sensitive data, (2) weak encryption schemes, (3) personally identifiable information, (4) weak passwords, or (5) services run at startup. Ex.1002¶161. Dr. Stavrou's declaration does **not** address any alleged disclosure of any of the other specifically listed types of data,

i.e., unencrypted system credentials, disabled Address Space Layout Randomization, boot record manipulation, suspicious definitions, services to be run on startup, personally identifiable information, data in application logs indicating that the protected virtual cloud asset accessed personally identifiable information, data in application logs indicating that the protected virtual cloud asset accessed a computer containing personally identifiable information, or at least one change in at least one area of the virtual disk, as compared to an earlier point in time.

118. Dr. Stavrou first argues that this limitation is met by Veselov on its own. I disagree.

119. First, Dr. Stavrou states that “Veselov’s snapshot-based security analysis can be performed in multiple ways” including by “analyzing an assessment VM instantiated based on the snapshot.” Ex.1002 (Stavrou Decl.) ¶

150. As detailed above, analyzing an assessment VM instantiated based on the snapshot is not analyzing a snapshot, accordingly, those aspects of Veselov do not disclose the claims, and in fact would have led a POSA away from what is claimed.

120. With respect to the specific types of data recited in this limitation, initially, Dr. Stavrou contends that Veselov includes an allegedly broad disclosure of “a wide variety of potential security problems.” (Ex.1002 (Stavrou Decl.) ¶ 155). I disagree. Veselov’s disclosures only discuss “vulnerabilities” and “host

configuration assessments.” *See* Ex.1007 (Veselov), 1:46-49 (“vulnerabilities, bad configurations”), 2:35-41 (“security assessments that are defined by rules packages, such as Common Vulnerabilities and Exposures (CVEs), Center for Internet Security (CIS) benchmarks, ‘best practices’ packages, static or runtime behavior analysis, host configuration Assessments”), 4:5-8 (“evaluate the vulnerabilities”), 4:14-18 (“identify changes in the virtual computing environment that may have caused or resolved certain vulnerabilities”), 7:1-5 (“security vulnerability”), 10:24-36 (discussing “vulnerabilities”), 14:41-50 (discussing “vulnerabilities”), 14:63-65 (discussing a CVEs assessment), 15:1-3 (discussing “vulnerabilities”), 21:1-2 (“describing a security vulnerability”), 23:30-33 (“describing security vulnerabilities”), 23:40-47 (describing “security vulnerabilities”), 23:56-60 (describing “security vulnerabilities”).

121. Veselov’s “vulnerabilities” and “host configuration assessments” are not the expressly recited types of data in limitations 1.4, 14.4 and 15.4. *See* ’926 patent, cls. 1, 14, 15 (listing unencrypted sensitive data, unencrypted system credentials, weak passwords, weak encryption schemes, disabled Address Space Layout Randomization, boot record manipulation, suspicious definitions, services to be run on startup, personally identifiable information, data in application logs indicating that the protected virtual cloud asset accessed personally identifiable information, data in application logs indicating that the protected virtual cloud

asset accessed a computer containing personally identifiable information, or at least one change in at least one area of the virtual disk, as compared to an earlier point in time).

122. I note that persons of ordinary skill in the art would have recognized differences between detecting vulnerabilities, such as through CVEs assessments, and detecting the claimed cyber threats based on virtual disk data. “CVE” or “Common Vulnerabilities Exposure,” refers to a database of *vulnerabilities* that can be detected as part of a security assessment. Ex.2188 (*Vulnerabilities*, NIST Vulnerability Database) (reported in the “national vulnerability database (NVD) of United States NIST, the “Common Vulnerabilities and Exposures (CVE) Program’s primary purpose is to uniquely identify *vulnerabilities* and to associate specific versions of code bases (e.g., software and shared libraries) to those vulnerabilities.”); *see also* Ex.1001 (’926 patent), 6:21-32 (“the security system 140 is configured to detect cyber threats that do not represent vulnerabilities”). The ’926 patent specification similarly discusses “common vulnerabilities and exposures (CVE®) databases” and “detect known vulnerabilities identified in, for example a CVE database”).” Ex.1001 (’926 patent), 4:38-43. While vulnerabilities are a type of cyber threat, they are different in kind to the “potential cyber threats based on data” expressly required in elements 1.4, 14.4, and 15.4, which, consistent

with the '926 patent specification, are “cyber threats that do not represent vulnerabilities.” Ex.1001 ('926 patent), 6:21-32.

123. The '926 patent contrasts analyzing snapshots to detect *vulnerabilities* through a snapshot analysis and detecting potential cyber threats based on specific types of data on a virtual disk, which are “cyber threats that do not represent vulnerabilities.” *Compare, e.g.,* Ex.1001 ('926 patent), 6:20-32 (“the security system 140 is configured to detect cyber threats that do not represent vulnerabilities”). I note that Orca filed additional patents with the US Patent and Trademark Office expressly claiming the detection of vulnerabilities, in contrast to the '926 patent which claims the detection of “potential cyber threats” instead of “vulnerabilities.” *See* Exs.2183-2187 (Orca's 031, 032, 809, 685, and 326 patents) at claim 1; *see also* Ex.1001 ('926 patent) 6:21-32 (potential cyber threats that do not represent vulnerabilities include, as examples, “sensitive data not being encrypted on the logical disk, private keys found on the disks, system credentials stored clearly on the disk, risky application features⁵ (e.g., support of weak cipher suites or authentication methods), weak passwords, weak encryption schemes, a disable address space layout randomization (ASLR) feature, suspicious

⁵ As noted above, “risky application features” is not one of the potential cyber threats recited in the '926 patent claims.

manipulation to a boot record, suspicious PATH, LD_LIBRARY_PATH, or LD_PRELOAD definitions, services running on startup, and the like.”). Persons of ordinary skill in the art would have therefore understood that detecting vulnerabilities, as in Veselov’s CVEs assessments, does not render obvious detecting “potential cyber threats” as claimed in the ’926 patent.

124. Dr. Stavrou opines that “CVEs were associated with services that were commonly run when the computer system starts up.” Ex.1002 (Stavrou Decl.) ¶ 156. That does not explain how or why a person of ordinary skill in the art would have understood any such specific CVEs to be detected in Veselov’s snapshot-based analyses. In fact, a particular security assessment will only detect the CVE’s that correspond to the data that is detected by the security assessment. Indeed, the CVE database is what collected data is *compared to* in order to determine vulnerabilities. And Dr. Stavrou’s reliance on selective CVE database entries describing vulnerabilities does not demonstrate how Veselov (or Mohanty) would determine potential cyber threats based on services to be run on startup. Ex.1002 (Stavrou Decl.) ¶156 (discussing Exs. 1090-1091, 1100-1103). I have reviewed Exs.1090-1091, and 1100-1103, and none of them reference “services to be run on startup,” provide any indication of what services are run on startup, and do not explain what types of security assessments will detect the particular vulnerabilities. For example, Ex.1090 is a vulnerability associated with “[T]he iax_net_read

function in the iaxclient open source library,” which is listed as a “vulnerability type(s) execute code overflow.”

Vulnerability Details : [CVE-2006-2923](#)

The `iax_net_read` function in the `iaxclient` open source library, as used in multiple products including (a) LoudHush 1.3.6, (b) IDE FISK 1.35 and earlier, (c) Kiax 0.8.5 and earlier, (d) DIAX, (e) Ziaxphone, (f) IAX Phone, (g) X-lite, (h) MediaX, (i) Extreme Networks ePhone, and (j) iaxComm before 1.2.0, allows remote attackers to execute

– CVSS Scores & Vulnerability Types

CVSS Score	6.4
Confidentiality Impact	Partial (There is considerable informational disclosure.)
Integrity Impact	Partial (Modification of some system files or information is possible, but the attacker does not have control over what can be modified, or the scope of what the attacker can affect is limited.)
Availability Impact	None (There is no impact to the availability of the system.)
Access Complexity	Low (Specialized access conditions or extenuating circumstances do not exist. Very little knowledge or skill is required to exploit.)
Authentication	Not required (Authentication is not required to exploit the vulnerability.)
Gained Access	None
Vulnerability Type(s)	Execute Code Overflow
CWE ID	119

Ex. 1091 is a vulnerability for “Stack-based buffer overflow in the NCTAudioFile2.AudioFile ActiveX control (NCTAudioFile2.dll),” also classified as a “vulnerability type(s) Execute Code Overflow”:

Vulnerability Details : [CVE-2007-0018](#)

Stack-based buffer overflow in the NCTAudioFile2.AudioFile ActiveX control (NCTAudioFile2.dll), as used by multiple products, allows remote attackers to execute arbitrary code via a long argument to the SetFormatLikeSample function. NOTE: the products include (1) NCTsoft NCTAudioStudio, NCTAudioEditor, and NCTDialogicVoice; (2)

– CVSS Scores & Vulnerability Types

CVSS Score	9.3
Confidentiality Impact	Complete (There is total information disclosure, resulting in all system files being revealed.)
Integrity Impact	Complete (There is a total compromise of system integrity. There is a complete loss of system protection, resulting in the entire system being compromised.)
Availability Impact	Complete (There is a total shutdown of the affected resource. The attacker can render the resource completely unavailable.)
Access Complexity	Medium (The access conditions are somewhat specialized. Some preconditions must be satisfied to exploit)
Authentication	Not required (Authentication is not required to exploit the vulnerability.)
Gained Access	Admin
Vulnerability Type(s)	Execute Code Overflow

Ex.1100 is a vulnerability for “system-tmpfiles in system through 237 mishandles symlinks present in non-terminal path components” which “allows local users to obtain ownership of arbitrary files”:

Vulnerability Details : [CVE-2018-6954](#)

systemd-tmpfiles in systemd through 237 mishandles symlinks present in non-terminal path components, which allows local users to obtain ownership of arbitrary files via vectors involving creation of a directory and a file under that directory, and later replacing that directory with a symlink. This occurs even if the fs.protected_symlinks sysctl is turned on.

Publish Date : 2018-02-13 Last Update Date : 2018-11-20

Ex.1101 is a vulnerability for a “race condition” that “allows an attacker to cause system to set arbitrary permissions”:

CVE-ID	
CVE-2018-15687	Learn more at National Vulnerability Database (NVD). • CVSS Severity Rating • Fix Information • Vulnerable Software Versions • SCAP Mappings • CPE Information
Description	
A race condition in chown_one() of systemd allows an attacker to cause systemd to set arbitrary permissions on arbitrary files. Affected releases are systemd versions up to and including 239.	

Ex.1102 is a vulnerability in “Apache httpd 2.2.0 to 2.4.29” where “when generating an HTTP Digest authentication challenge” a “nonce sent to prevent reply attacks was not correctly generated”:

CVE-ID	
CVE-2018-1312	Learn more at National Vulnerability Database (NVD). • CVSS Severity Rating • Fix Information • Vulnerable Software Versions • SCAP Mappings • CPE Information
Description	
In Apache httpd 2.2.0 to 2.4.29, when generating an HTTP Digest authentication challenge, the nonce sent to prevent reply attacks was not correctly generated using a pseudo-random seed. In a cluster of servers using a common Digest authentication configuration, HTTP requests could be replayed across servers by an attacker without detection.	

Ex.1103 is a vulnerability that “allows remote attackers to cause a denial of service” via “a POST request”:

Vulnerability Details : CVE-2013-2765
The ModSecurity module before 2.7.4 for the Apache HTTP Server allows remote attackers to cause a denial of service (NULL pointer dereference, process crash, and disk consumption) via a POST request with a large body and a crafted Content-Type header. Publish Date : 2013-07-15 Last Update Date : 2013-11-18

125. As discussed above, these references are simply excerpts from the CVE database, and thus describe “vulnerabilities,” not “potential cyber threats” that are “based on data stored on the virtual disk” as recited in the ’926 patent’s claims. There is no support that they are “associated with services that were commonly run when the computer system starts up” but an examination of these references indicates no relation to services that run when a computer system starts

up, and Dr. Stavrou cites no other documentation in support of his conclusion. Ex.1002 (Stavrou Decl.) ¶156.

126. Dr. Stavrou also argues that detection of nonsecure system configurations includes the detection of “potential cyber threats.” Ex.1002 (Stavrou Decl.) ¶157. But this position ignores that nonsecure system configurations are not “at least one of” the types of “cyber threats based on data” in limitations 1.4, 14.4, and 15.4.

127. Dr. Stavrou’s reliance on nonsecure system configurations (Ex.1002 (Stavrou Decl.), ¶157) therefore does not establish that Veselov teaches this claim limitation. Assessments of system configurations do not include any of the specific types of data in the ’926 patent claims and do not obviously or necessarily require detection of the claimed potential cyber threats in elements 1.4, 14.4 and 15.4. In fact, the tertiary references Dr. Stavrou relies on in other portions of his declaration refute his conclusion here, as those references explain, for example, that “the configuration of an asset is a hardware and/or software configuration,” and could include assessing “particular network port settings.” Ex.1008 (Basavapatna), ¶ 37. As another example, a “system configuration xml file” may include host names and component names but would not generally be associated with potential cyber threats such as passwords or encrypted data. Ex.1020 (Sancheti), 12:40-13:4. Therefore, an assessment of hardware configurations or port settings or other

purported “non-secure configurations” would not obviously or “necessarily” include determining “weak passwords” or “encryption schemes” but would instead focus on, for example, host data and/or network and connectivity settings.

128. Dr. Stavrou also cites Veselov’s statement that “contemplates implementation of any *suitable* security assessment, including security assessments that are defined by rules packages, such as Common Vulnerabilities and Exposures (CVEs), Center for Internet Security (CIS) benchmarks, “best practices” packages, static or runtime behavior analysis.” (Ex.1002 (Stavrou Decl.) ¶155.) From this, Dr. Stavrou appears to contend that Veselov can be used with *any* security assessment. I disagree. In particular, Dr. Stavrou fails to give any thought to what is “suitable” versus not “suitable” in Veselov. Veselov itself indicates to a POSA that, for example, analyzing a snapshot is not “suitable” with an agent based vulnerability or configuration assessment, which is why Veselov describes them as entirely distinct embodiments that collect different data. It is my opinion that a POSA reading Veselov, which discusses vulnerability and configuration assessments as its examples, would not have understood Veselov to be suitable with detecting cyber threats that do not represent vulnerabilities like in the challenged claims.

129. Dr. Stavrou’s position that Veselov’s security assessment identifies potential cyber threats “based on data stored on the virtual disk” (Ex.1002 (Stavrou

Decl.) ¶158) also fails. Dr. Stavrou provides myriad citations to Veselov's volumes and snapshots, but none of the examples that Veselov provides correspond to the specific types of data enumerated in limitations 1.4, 14.4, and 15.4.

130. Separately from Veselov, Dr. Stavrou also relies on disclosures of Mohanty. *See, e.g.*, Ex.1002 (Stavrou Decl.) ¶ 159. But Mohanty does not perform any detection via a snapshot analysis and is inapplicable to the claimed invention for at least this reason.

131. Dr. Stavrou relies initially on Mohanty's CVEs assessments (Ex.1002 (Stavrou Decl.), ¶159), but Mohanty's CVE analysis does not necessarily determine at least one of the claimed potential cyber threats based on data that are required by limitations 1.4, 14.4 and 15.4 for the same reasons as detailed above with respect to Veselov's CVE assessments. Further, Mohanty's data analysis is based on input from publicly available vulnerability assessment (VA) products, such as from Qualys or Rapid7, which are agent-based approaches to analyze a virtual cloud asset. Ex.2173 (Qualys, *Vulnerability Management*) (relying on "[s]ensors that provide continuous visibility" and "Agent technology" that lets users "proactively address potential threats whenever new vulnerabilities appear"); Ex.2174 (Man, *Live Monitoring with Endpoint Agents*) ("Our agents automatically collect data from your endpoints"); *see* Ex.2172 (2nd Stavrou Tr.) 239:3-240:9 (testifying that he "didn't look at" whether Mohanty uses agent-based approaches

in preparing his declaration). Such scanners do not analyze or evaluate snapshots, as required by the analysis of the '926 patent claims.

132. Additionally, Dr. Stavrou's reliance on Mohanty's "threat information" and "threat module" fails. Ex.1002 (Stavrou Decl.) ¶¶159-160. Mohanty's threat module does not "determine" cyber threats "based on data stored on the virtual disk" as claimed. Instead, Mohanty explains that its "threat module" refers to an external "threat intelligence system" such as DEEPSIGHT that provides "information 208 *about external threats*" and "reports the latest lists" of such threats. Ex.1075 (Mohanty), 3:60-67, 4:33-36, 6:64-7:5. A POSA would have understood that references to things like an "external threat," would be in reference to risks that originate from outside of an organizations network like a known hacking group is currently targeting machines with a particular operating system (OS) using a specific exploitable vulnerability. *See* Ex.1075 (Mohanty), 6:64-7:5. In other words, Mohanty's threat information stems from risk external to the asset, rather than potential cyber threats present on a virtual disk, and Mohanty combines this external threat data with vulnerability data collected from its agent-based scanners in order to form its security assessment. Ex.1075 (Mohanty), 7:5-9.

133. And, I note that Dr. Stavrou's position that Mohanty addresses the claimed sensitive data detection (Ex.1002 (Stavrou Decl.) ¶¶164-165) is incorrect.

Although Mohanty discusses “sensitive data,” Mohanty does not suggest “analyzing” a virtual disk of a virtual cloud asset to *determine* sensitive data as required by claims 1.4, 14.4 and 15.4, and instead acknowledges that sensitive data may be tagged upon **creation** of an asset. More specifically, Mohanty’s data is a “static tag” that can be set by a virtual cloud environment (such as VMWare) when “an asset (e.g., a virtual machine or a virtual application) is created and deployed in a virtual environment.” Ex.1075 (Mohanty), 4:4-8, 4:13-15. The “static tag information” may “include indication of whether there is sensitive data” or “CRITICAL DATA” “in or handled by the asset.” Ex.1075 (Mohanty), 5:38-43. Dr. Stavrou also acknowledges this type of static tag in Mohanty and that these static tags are set upon asset deployment and are not changed. *See* Ex.2172 (2nd Stavrou Tr.)199:10-12 (“Q. And there are two types of contextual tags in Mohanty: static and dynamic; right? A. Correct.”), 209:12-19 (“static tags” are when asset or virtual environment is created). Ex.2172 (2nd Stavrou Tr.) 212:6-10 (explaining that “static tags” are set when the asset is created or deployed and do not change during the lifespan of an asset).

134. Dr. Stavrou is also incorrect in his assertion that “the combination of Veselov and Mohanty teaches or suggests determining the existence of potential cyber threats based on stored data that includes ‘services to be run at startup.’” (Ex.1002 (Stavrou Decl.) ¶166). Specifically, Dr. Stavrou relies on disclosures of

CVEs assessments, but neither Veselov nor Mohanty discloses the *detection* of cyber threats based on service to be run on startup. Dr. Stavrou instead appears to confuse detecting vulnerabilities via CVEs assessments with the claimed detection of potential cyber threats based on data stored on a virtual disk. Put simply, that CVEs may relate to services commonly run at startup does not teach the claimed requirements of limitations 1.4, 14.4, and 15.4. Likewise, Dr. Stavrou's assertion that "A POSA would have expected many of the VMs assessed per the teachings of Veselov and Mohanty to have startup services that were the basis of potential cyber threats and that were flagged as part of typical CVE assessments" (Ex.1002 (Stavrou Decl.) ¶167) is inapposite to the challenged claims. Dr. Stavrou does not provide any citations in the prior art disclosing that analyzing *a snapshot* would determine cyber threats based on the particular types of data claimed.

135. Dr. Stavrou's position that Veselov teaches "best practices" in its scanning assessments and that "best practices" generally include "avoiding exposure of unencrypted or weakly-encrypted data" also misses the mark. Ex.1002 (Stavrou Decl.) ¶162. Specifically, broad assertions relating to the coverage of best practices is not sufficient to teach the detection of potential cyber threats based on the '926 patent's enumerated virtual disk data. Dr. Stavrou cites to his overview of cybersecurity to further support his reliance on "best practices." (Ex.1002 (Stavrou Decl.) ¶162 (relying on §VIII.D.1-2 of the Declaration)). But here, Dr.

Stavrou simply cites a swath of references, all with disparate and varied techniques security assessments that detect different types of data, that can detect types of security risks and techniques to detect security risks through a variety of tertiary references that do not explain why a person of ordinary skill would combine Veselov and Mohanty to arrive at the challenged claims, which require the detection of potential cyber threats based on specific types of virtual disk data. Ex.1002 (Stavrou Decl.) ¶¶49-63 (relying on additional tertiary references). None of the references refer to “best practices” for analyzing *a snapshot*, let alone to determine what is claimed.

136. I also note that Wiz’s own documentation demonstrates that analyzing snapshots to detect data-based cyber threats in the cloud were not obvious or well known, even well after the 926 patent’s priority date in January 2019. For example, Wiz explained in 2022 articles that “securing cloud data is a difficult problem” (Ex.2155 (Wiz Becomes the First CNAPP to Deliver Data Security Posture), 2) and other articles from third parties about Wiz’s new product release explained that “less than 1% of companies around the world use a DSPM solution,” i.e., a system that “identif[ies] where critical information – such as ... personal information ... is stored in the cloud,” and was a field that “only emerged in recent years” (Ex.2154 (Wiz Launches New Cloud Data Security technology), 3).

137. And, Dr. Stavrou's statement that Mohanty may apply "stronger encryption" (Ex.1002 (Stavrou Decl.) ¶163) to sensitive data is beside the point. The challenged claims require *analyzing* the snapshot of a virtual disk to *determine* potential cyber threats based on data stored on the virtual disk. Any alleged stronger encryption does not change that Mohanty's "sensitive data" is something a virtual machine is "mark[ed]" with as a static tag, rather than being the subject of being "analyzed" or determined. Ex.1075 (Mohanty), 5:38-43.

138. It is my opinion that a POSA would have understood Mohanty's "static tags," including "sensitive data," are not determined as part of a security assessment particularly in view of Mohanty's discussion of "dynamic tags." Mohanty explains that "dynamic" tags relate to information "that is subject to change during the lifespan of the asset," and can also referred to as "security tags, since the changing information is of interest regarding security of the asset." Ex.1075 (Mohanty), 4:15-23. These "dynamic" or "security" tags include information like "VIRUS FOUND" or "INTRUSION DETECTED," which are determined "by security technologies monitoring the same systems" and are distinct from the static tags that do not change. Ex.1075 (Mohanty), 5:43-48. Mohanty explains that "security technologies" monitoring a virtual cloud asset can set these dynamic tags, but there is no similar disclosure of using any "security technologies" or "monitoring" system to determine "static tags" like "sensitive

data.” *See generally* Ex.1075 (Mohanty). Mohanty does not teach that “dynamic tags” would (or should) include any of the “cyber threats based on data stored on the virtual disk” recited in the challenged claims.

139. Additionally, Mohanty’s disclosure that a “dynamic tag” could include “whether or not suspicious data has been transferred into or out of the asset” does not bridge Mohanty closer to the challenged claims. Ex.1075 (Mohanty), 5:54-59. The ’926 patent claims require determining cyber threats based on data “***stored on*** the virtual disk.” ’926 patent, cls. 1, 14, 15. A person of ordinary skill reading Mohanty would understand that data stored on a virtual disk is distinct from Mohanty’s suspicious “transfers” of data into or out of the asset, which does not necessarily implicate virtual disk data. And, I note that Mohanty makes clear that “suspicious data” transferred in and out of an asset is also distinct from “sensitive data,” as it provides that the former is represented by dynamic tag data or changing data, and the latter is represented by static tag or unchanged data. *See* Ex.1075 (Mohanty), 5:50-58. For example, “suspicious data” being transferred could include transfer of *any* data from a networked computer to a computer located in a foreign country and therefore encompasses data that changes over time rather than remains static from asset deployment. And, I further note that Dr. Stavrou admitted during his deposition that a person of ordinary skill could not conclude that Mohanty’s reference to “suspicious data” is the same as

“sensitive data” without further context. *See also* Ex.2172 (2nd Stavrou Tr.) 200:18-201:6.

140. For these reasons, it is my opinion that Dr. Stavrou fails to show that limitations 1.4, 14.4, and 15.4 would have been obvious over Veselov and Mohanty.

b. Limitations 1.5, 14.5, and 15.5

141. Limitations 1.5, 14.5, and 15.5 recite “[determine/determining] a risk associated with each of the determined plurality of potential cyber threats.” Dr. Stavrou fails to show that Veselov and Mohanty disclose or suggest these limitations. Ex.1002 (Stavrou Decl.), ¶¶169-172.

142. Dr. Stavrou does not rely on Veselov in his analysis of this limitation. With respect to Mohanty, Dr. Stavrou argues that Mohanty “teaches determining the existence of threats by correlating ‘threat information’ with the asset’s vulnerability data and with contextual information describing data stored on the asset and other characteristics of the asset.” Ex.1002 (Stavrou Decl.), ¶170. Dr. Stavrou specifically refers to “calculating a threat score in columns 8-9” (Ex.1002 (Stavrou Decl.), ¶171) and, as he states, in this example the threat module 308 calculates a threat score. Ex.1075 (Mohanty) 7:59-8:14.

143. But Dr. Stavrou omits critical details of this process. As Mohanty explains, the score the “threat module 308” calculates uses “threat information

208.” Ex.1075 (Mohanty) 7:59-8:14. Mohanty explains earlier that its “threat information” represents “external threats.” Ex.1075 (Mohanty), 4:38-58. For example, Mohanty explains that its threat information “comes from one or more threat intelligence systems” such as “DEEPSIGHT” which “provide threat information [] about external threats.” Ex.1075 (Mohanty), 4:33-35. Thus, the information that Mohanty “correlate[s]” is based on risk that exists external to the asset, and it follows that risks “external” to an asset would not be “based on data stored on the virtual disk” as recited in the claims. Ex.1075 (Mohanty), 4:35-39. Mohanty’s claims 14 and 17 further confirm that the threat information is received from a source that is external to the specific or target asset. Ex.1075 (Mohanty), cls. 14, 17. Mohanty’s threats are therefore distinguishable from the ’926 patent’s claimed “cyber threat[s] based on data stored on the virtual disk” that form the basis for the subsequent “determined plurality of potential cyber threats.” Ex.1001 (’926 patent), cls. 1, 14, 15.

144. Accordingly, it is my opinion that Dr. Stavrou fails to show that limitations 1.5, 14.5, and 15.5 would have been obvious over Veselov and Mohanty.

c. Limitations 1.6, 14.6, 15.6

145. Limitations 1.6, 14.6, 15.6 recite “[prioritize/prioritizing] the potential cyber threats associated with the protected virtual cloud asset based on

the determined risk associated with each of the plurality of potential cyber threats.”

Dr. Stavrou fails to show that Veselov and Mohanty disclose or suggest these limitations. Ex.1002 (Stavrou Decl.), ¶¶173-176.

146. I note that for Limitations 1.6, 14.6, and 15.6 Dr. Stavrou’s relies on the same “external threats” as I discuss above with respect to Limitations 1.4-1.5, 14.4-14.5, and 15.4-15.4. Accordingly, my explanation with respect to Limitations 1.4-1.5, 14.4-14.5, and 15.4-15.5 applies equally to Limitations 1.6, 14.6, and 15.6.

147. Therefore, it is my opinion that Dr. Stavrou fails to show that limitations 1.6, 14.6, and 15.6 would have been obvious over Veselov and Mohanty.

d. Limitations 1.7, 14.7, and 15.7

148. Limitations 1.7, 14.7, and 15.7 recite: “reporting at least some of the determined plurality of potential cyber threats as alerts prioritized according to their associated risks.” Dr. Stavrou relies on Veselov in view of Mohanty to teach these limitations. Ex.1002 (Stavrou Decl.), ¶¶177-186. It is my opinion that Dr. Stavrou fails to show that Veselov and Mohanty disclose or suggest these limitations.

149. Contrary to Dr. Stavrou’s position that “Veselov teaches reporting ‘at least some of the determined plurality of potential cyber threats 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 1.7, 14.7, and 15.7. A person of ordinary skill in the art would recognize that these limitations do not end after reciting “reporting at least some of the determined plurality of potential cyber threats,” but continue and require that the determined plurality of potential cyber threats be reported “as alerts.” *See* Ex.1001 (’926 patent), cls. 1, 14, 15. And, the specification of the ’926 patent confirms my understanding, as it also demonstrates that reporting results “as alerts” involves techniques beyond simply reporting results. *See* Ex.1001 (’926 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”).

150. 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 threats “as alerts,” as claimed. *See, e.g.*, Ex.1001 (’926 patent), 7:53-54 (“[T]he detected cyber threats and/or vulnerabilities are reported, *for example, as alerts*”) (italics added), cls. 1, 14, 15. In fact, 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.”). Veselov’s disclosure of simply “communicating” results 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

151. 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” 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)”).

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 potential cyber threats as alerts, as is required by the claim language.

152. 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 ’926 patent’s reporting of results as alerts, and in fact, Dr. Stavrou does not provide a single citation to an alert in Veselov.

153. 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 1.7, 14.7, and 15.7 require.

154. Dr. Stavrou also relies on select portions of Mohanty. Ex.1002 (Stavrou Decl.) ¶¶ 179-182. But Mohanty does not cure the deficiencies of Veselov.

155. As an initial matter, Dr. Stavrou’s reliance on Mohanty includes the “external threats” I discuss above. Accordingly, Mohanty fails to report “potential cyber threats” as prioritized alerts for at least the reasons I discuss with respect to Limitations 1.4-1.6, 14.4-14.6, and 15.4-15.6.

156. Mohanty mentions “alerts” in the context of a remediation response so that appropriate personnel could take appropriate action (Ex.1075 (Mohanty), 9:11-26). Those “alerts,” however, are based on what is detected in real time by an agent, which is entirely distinct from the data provided in Veselov based on a snapshot at a past point in time. Dr. Stavrou asserts that Mohanty meets the claimed alert prioritization because “Mohanty teaches including prioritization information

in the alerts.” Ex.1002 (Stavrou Decl.), ¶180. I disagree. Mohanty teaches that its “prioritization scores” are associated with “a specific **vulnerability** of a specific asset,” rather than “cyber threats based on data stored on a virtual disk” as claimed.

Mohanty, 9:16-19.

REDACTED

158. Therefore, it is my opinion that Dr. Stavrou fails to show that limitations 1.7, 14.7, and 15.7 would have been obvious over Veselov and Mohanty.

159. 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 '926 patent—claims 1, 14, and 15—would have been obvious. And 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 5-10 and 12-13), would have been obvious. Thus, Wiz and Dr. Stavrou have not shown that any of claims 1, 5-10, and 12-15 of the '926 patent would have been obvious.

160. In addition to failing to disclose or suggest the limitations of independent claims 1, 14, and 15, the combination of Veselov and Mohanty also fails to disclose or suggest limitations contained in multiple dependent claims, as analyzed below.

e. **Dependent Claim 13**

161. Dependent claim 13 depends from claim 1 and recites “wherein locating 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* of the respective protected virtual cloud asset.”

162. 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.) ¶227. 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, 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 exists.

163. 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.) ¶231. 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.)

¶231. 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 (Veselov), 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 Veselov-Mohanty combination would *benefit* from such information.




FIG. 1

Ex.1007 (Veselov) Fig. 1

164. At best, Veselov 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 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 claim 13. 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.) ¶¶232-233.

165. As the '926 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 ('926 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, (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.

166. 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.) ¶230. Dr. Stavrou also quotes Exhibit 1031 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) at 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.

167. Dr. Stavrou also 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.” But his analysis for that section does nothing to address the issues I discuss above.

168. 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.), ¶¶223-224. 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 '926 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. ¶224, shows that there 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), ¶¶77-79, 81-82, 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 ¶224 (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.

f. **Summary of Ground 1 Opinions**

169. For at least the foregoing reasons as well as the objective indicia I discuss in Section IX below, it is my opinion that Wiz and Dr. Stavrou have not shown that any of the independent claims of the '926 patent—claims 1, 14, and 15—as well as dependent claims 2-13, would have been 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 and 12-13), would have been obvious. Thus, Wiz and Dr. Stavrou have not shown that any of claims 1-15 of the '926 patent would have been obvious

B. Ground 2: Claims 1-10 and 12-15 Would Not Have Been Obvious Over Veselov, Mohanty, and Ranum

170. Dr. Stavrou sets forth Ground 2 to “build[] upon” Ground 1 and to further apply “Ranum’s [alleged] disclosure of threats based on particular types of data, reporting risk levels, and comparing computed cryptographic hashes of disk files to earlier hashes for security assessments.” Ex.1002 (Stavrou Decl.) ¶236. But Dr. Stavrou’s analysis for Ground 2 relies on Ground 1, which is based on Veselov and Mohanty. Accordingly, for at least the same reasons Dr. Stavrou fails to demonstrate obviousness with respect to Ground 1, it is also my opinion that Dr. Stavrou fails to demonstrate obviousness with respect to Ground 2.

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

171. A person of ordinary skill in the art would not have been motivated to combine Ranum 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 2 because Ground 2 assumes the combination of Veselov and Mohanty. *See* Section VIII.A.

172. Additionally, even assuming that Veselov and Mohanty are combined, a person of ordinary skill in the art would not have combined them with Ranum to render the claims obvious. 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. Ranum does not teach snapshot scanning, which I understand Dr. Stavrou also agreed with in his deposition. *See* Ex.2172, 251:8-252:2 (conceding Ranum uses “active network scanning” which is different than Veselov’s snapshot based assessments), 276:17-20 (Stavrou does “not offer any opinion in [his] declaration that Ranum is analyzing snapshots of a virtual cloud asset”). Instead, Ranum is directed towards active scanning and passive network monitoring and detects risk by “communicat[ing] packets or other messages within the network to detect new or changed information.” Ex.1093

(Ranum) Abstract, 4:36-57, 13:41-44 (“in one implementation, the active scanners 210 may generally scan...”), 13:48-49 (“active scanners 210 may perform the credentialed and/or uncredentialed scans in the network”), 14:3-11 (“... the active scanners 210 may then report the information ... use the information provided by the active scanners...”). A POSA would have understood these teachings in Ranum to be incompatible with Veselov’s analysis of a snapshot because snapshots lack active processes. Accordingly, a POSA would have understood that Ranum’s teachings for how to detect vulnerabilities or other threats cannot be applied to snapshot analyses, because there is no mechanism to communicate packets or other messages to a snapshot. Ex.2172 (2nd Stavrou Tr.) 256:11-257:17 (discussing Ranum 10:34-41). I understand that Dr. Stavrou agreed with this too. Ex.2172 (2nd Stavrou Tr.) 257:20-258:6 (“Q. ... I just want to make sure that I’m right that because a snapshot is a data file, does not have any active processing, it would not be possible to communicate messages or packets to it as described in Ranum? A. **That is correct.**”). Thus, a person of ordinary skill in the art would have recognized that the techniques of Veselov’s analysis of snapshots and Ranum were fundamentally different and would not have combined the teachings of Ranum with Veselov.

2. Even If Combined, the Veselov/Mohanty/Ranum Combination Does Not Render Obvious Limitations 1.4, 14.4, or 15.4

173. Dr. Stavrou opines that Ranum “reinforces the teachings of Veselov and Mohanty” and “provides additional express disclosure of detecting threats that are based on weak encryption schemes, unencrypted sensitive data, personally identifiable information, weak passwords, and services to be run on startup” and of “detecting threats that are based on changes to at least one area of the asset as compared to an earlier point in time.” Ex.1002 (Stavrou Decl.) ¶¶247-248. I disagree that Ranum “reinforces” the teachings in Dr. Stavrou’s purported combination of Veselov and Mohanty, which relies on snapshot scanning, because Ranum teaches incompatible assessment techniques for analyzing a snapshot.

174. I note initially that Ranum’s detection of threats based on data is derived only from data collected from “active scanners.” Ex.1093 (Ranum), 13:41-53. For example, Ranum explains that the “active scanners” “may perform the credentialed and/or uncredentialed scans in the network” and “may then report the information.” Ex.1093 (Ranum), 13:48-49, 14:3-11. Ranum also explains that the system may “use the information provided by the active scanners...” Ex.1093 (Ranum), 14:3-11. Persons of ordinary skill would therefore understand that the data in Ranum is collected via “active scanners” and that any threat information detected in Ranum requires performance of active scans.

175. I also note that persons of ordinary skill would understand that active scanners and passive network monitoring such as what is described in Ranum are incompatible with techniques for analyzing a snapshot such as what is described in the '926 patent and Veselov's embodiments for analyzing a data file. Dr. Stavrou confirmed in his deposition that Ranum's use of "active network scanning and passive network monitoring" are different than Veselov's snapshot-based assessments. *See* Ex.2172 (2nd Stavrou Dep.) 251:8-252:2; *see also* Ex.2172 (2nd Stavrou Dep.) 276:17-20 ("[Dr. Stavrou does] not offer any opinion in [his] declaration that Ranum is analyzing snapshots of a virtual cloud asset"). The '926 patent also describes problems with prior art "network scanners" and "traffic inspection" such as what is described with respect to Ranum's passive network monitoring. Ex.1001 ('926 patent), 1:56-2:9.

176. Ranum performs its "active scan[s]" to detect vulnerabilities and other cyber threats by "communicat[ing] packets or other messages within the network to detect new or changed information describing various routers, hosts, servers, or other devices in the networks." Ex.1093 (Ranum), 4:36-57, 10:34-41. A person of ordinary skill would have understood that Ranum's techniques cannot be combined with Veselov's Figures 5A and 5B embodiments for analyzing a snapshot because a snapshot that is analyzed as a data file does not have active processing. I understand Dr. Stavrou agreed with this during his deposition too:

11 Q. Can we go to column 10, lines 34 through 41?
12 A. Let me do that.
13 Q. And here Ranum describes "active scanners
14 110 that can communicate packets or other messages
15 within the network to detect new or changed
16 information describing various routers ...
17 hosts ... servers ... or other devices 130 in the
18 network."
19 Do you see that?
20 A. Uh-huh.
21 Q. It would not be possible to send or
22 communicate packets or other messages to a snapshot

1 as a data file; right?
2 A. It's not through the data file. It doesn't
3 have active processing.

Ex.2172 (2nd Stavrou Dep. 256:11-257:3.

20 Q. Yeah. And my question is a different one --
21 I think maybe a simpler one, which -- I just want to
22 make sure that I'm right that because a snapshot is

1 a data file, does not have any active processing, it
2 would not be possible to communicate messages or
3 packets to it as described in Ranum?

4 A. That is correct. And, again, I don't -- I
5 don't believe that we're using the teachings of the
6 network here for the combination.

Ex.2172 (2nd Stavrou Dep. 257:20-258:6.

177. Furthermore, Dr. Stavrou appears to contend that because Mohanty, Ranum, and Veselov are all dealing with security assessments of virtual or physical machines, they are all considering the “same data.” I disagree. The way that a machine is inspected determines what data is considered, and it is my opinion that persons of ordinary skill would have understood that Mohanty’s agent-based “vulnerability assessment tools,” Ranum’s “network scanners,” and Veselov’s agentless “snapshot” based analyses inspect virtual cloud assets in fundamentally different ways. Both Ranum and Mohanty’s techniques require interactions with an active virtual cloud asset for analysis, which is directly contrary to Veselov’s stated objective of using a security assessment that “avoid[s] consuming processing time, processing power, and storage space of the target computing resource.” Veselov, Abstract, 12:21-43. Dr. Stavrou at his deposition also agreed that the data considered by Veselov and Ranum both look different and the data they detect are “not identical.”

18 Q. Active network scanning and passive network
19 monitoring are different security assessment
20 techniques than what is disclosed in Veselov; right?

21 A. While they are -- on the face of it, it
22 looks different, they rely on similar information.

252

1 That's important here. So we are not relying on
2 different set of information.

3 So, for example, Veselov does -- would be
4 able -- scanners in Veselov would be able to get --
5 host suspicious files, as we have discussed, or
6 files in general. And malware signatures, again, it
7 talks about malware signatures that various
8 antivirus vendors have cataloged to detect malware
9 infections. All of these agree with me that McAfee
10 would have been the scanner or host.

11 So these things are present in a host. And
12 the data here, the important part is that the data
13 that Veselov analyzes and what Ranum is talking
14 about are similar. Some of them might not be. So
15 they're not identical. And then Ranum might have
16 more information about the network. But there
17 are -- there's an overlap. That's what I'm trying
18 to say.

178. Ex.2172 (2nd Stavrou Dep) 252:11-18.

179. Accordingly, Dr. Stavrou has not shown that Veselov, Mohanty, and
Ranum render obvious Limitations 1.4, 14.4, and 15.4.

3. Even If Combined, the Veselov/Mohanty/Ranum Combination Does Not Render Dependent Claims 2 and 4 Obvious

180. Dr. Stavrou argues that Ranum discloses “detecting the data stored on the virtual disk by computing a cryptographic hash of at least one area of the virtual disk and comparing the computed cryptographic hash of the at least one area of the virtual disk to an earlier compute[d] cryptographic hash of the at least one area of the virtual disk” as is required by dependent claims 2 and 4. Ex.1002 (Stavrou Decl.) ¶¶256-265. I disagree.

181. I initially note that dependent claims 2 and 4 require that the disclosed “virtual disk” be the same “virtual disk” as is described with respect to the snapshot analysis of the independent claims. Ex.1001 ('926 patent), cls. 1, 2, 4, 14, 15. Accordingly, the detection of the “data stored” must take place via a snapshot analysis. Ex.1001 ('926 patent), cls. 1, 2, 4, 14, 15. But Ranum’s detections are not based on data stored on a virtual disk; they instead are based on the identification of “malware signatures.” Ex.1093 (Ranum), 1:15-26. And, the cryptographic hashes of Ranum on which Dr. Stavrou relies are not used as part of a snapshot analysis and instead are used to “index[] one or more trusted file systems known to have no malware infections.” Ex.1093 (Ranum), 17:1-9. For example, Ranum further describes an implementation where the reference systems “include a clean system shipped from a suitable vendor that can be trusted to have no

malware infections.” Ex.1093 (Ranum), 17:9-20; *see also* Ex.1093 (Ranum), 17:21-23 (describing cryptographic has computations for “trusted file systems”), 17:49-54 (describing “file hashes” for “trusted file systems.”). And, Dr. Stavrou confirmed in his deposition that Ranum’s trusted file systems are not the running systems on the network and therefore are not comparable to the security assessment of a virtual disk. 0402Stavrou 261:18-22 (“Q. What are the known reference systems? Those aren't the running systems in the network. A. Yeah.”).I also note that Ranum’s techniques for “detecting data” using “hashes” are only implemented through “active scanners” that “perform a credentialed scan.” Ex.1093 (Ranum), 17:54-63. As Dr. Stavrou agreed at his deposition, Ranum’s discussion in Columns 16 and 17 and in relation to Figure 3 is determining processes currently running on managed hosts in the network using an active scanner. Ex.1093 (Ranum) 16:47-57.

13 I'm just asking, in the context of figure 3
14 and the discussion on 16 and 17, Ranum is
15 determining processes currently running on managed
16 hosts in the network using an active scanner; is
17 that fair?

18 A. That's what it says, yes.

Ex.2172 (2nd Stavrou Dep.) 259:13-18.

But the purpose of these active scanners is to detect processes that are “**running** on managed hosts.” Ex.1093 (Ranum), 17:54-63. As I describe above, Ranum’s

active scanners for its credentialed scans require active processes that persons of ordinary skill would not combine with Veselov's analysis of snapshots as data files because such snapshots have "no active processes" to detect using the active scanning procedures of Ranum. 04/02Stavrou 256:11-258:6.

182. It is my opinion that persons of ordinary skill reading Veselov's embodiments for analyzing snapshots as data files would not have understood they could be combined with Ranum's descriptions of active scanners because combining Ranum with Veselov would alter the snapshot analysis of Veselov and would interfere with Veselov's stated purpose of developing a security assessment that "avoid[s] consuming processing time, processing power, and storage space of the target computing resource." Ex.1007 (Veselov), Abstract, 12:21-43, 2:54-63, 17:9-12. Accordingly, a person of ordinary skill would not have combined Ranum's teachings with the snapshot-based virtual disk analysis disclosed in Veselov or the proposed Veselov/Mohanty combination.

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

183. Dr. Stavrou sets forth Ground 3 to "build[] upon" Ground 1 and to "further apply Seo's disclosure of well-known techniques to identify illegitimate files stored on an asset." Ex.1002 (Stavrou Decl.) ¶266. But Dr. Stavrou challenges claim 1, on which claim 3 depends, based on only Veselov and Mohanty. Accordingly, for at least the same reasons Dr. Stavrou fails to

demonstrate obviousness of claim 1, it is also my opinion that Dr. Stavrou fails to demonstrate obviousness of claim 3.

184. It is my opinion that claim 3 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 Seo and Would Not Have Had a Reasonable Expectation of Success in Doing So

185. A person of ordinary skill in the art would not have been motivated to combine Seo 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.

186. Additionally, even assuming that Veselov and Mohanty are combined, a person of ordinary skill in the art would not have combined them with Seo to render the claims obvious.

187. Seo's approach is also different than, and would not have been combined with, Veselov's approach. Veselov discloses scanning operations that are focused on a "scannable volume" from a snapshot captured at a specific point in time in order to detect vulnerabilities in VMs. *E.g.* Ex.1007 (Veselov) 10:1-2. In contrast, Dr. Stavrou confirms that Seo instead is directed to *preventing attacks*

(Ex.1002 (Stavrou Decl.) ¶268) in an artificial intelligence context and involves techniques for preventing encryption-related vulnerabilities based on whether artificial intelligence technology is enabled. Ex.1094 (Seo), Abstract. Dr. Stavrou does not provide any indication why a POSA would have looked solely to Seo's alleged teaching of identifying illegitimate files based on "installation records" (Ex.1002 (Stavrou Decl.) ¶269) in a combination with Veselov and Mohanty, where these processes of Seo are based on Seo's artificial intelligence network, which is not used in either Veselov or Mohanty or Dr. Stavrou's proposed combination. It is my opinion that POSA would not have selectively identified only the certain aspects of Seo to combine with the certain portions of Veselov and Mohanty that Dr. Stavrou relies on, unless that POSA had the 926 patent claims as a guide to select the references and portions. I understand that is an improper analysis for obviousness. I also note that Seo would not have led a POSA to the challenged claims because, for example, Seo does not describe or explain using any snapshot scanning or analysis of VMs in a cloud computing environment context for the detection of cyber threats based on data claimed in the patent, nor does Dr. Stavrou opine that it does.

188. Accordingly, I disagree that a POSA would have been motivated to combine the particular portions of Seo with the particular portions of Veselov and Mohanty that Dr. Stavrou has identified.

D. Ground 4: Claim 3 Would Not Have Been Obvious Over Veselov, Mohanty, Seo, and Ranum

189. Dr. Stavrou sets forth Ground 4 to “build[] upon” Ground 2 and to “further apply Seo’s disclosure of well-known techniques to identify illegitimate files stored on an asset” as described in Ground 3. Ex.1002 (Stavrou Decl.) ¶266. But Dr. Stavrou challenges claim 1, on which claim 3 depends, based on only Veselov and Mohanty and Ranum. Accordingly, for at least the same reasons Dr. Stavrou fails to demonstrate obviousness of claim 1 in Ground 2, it is also my opinion that Dr. Stavrou fails to demonstrate obviousness of claim 3.

190. It is my opinion that claim 3 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, Seo, and Ranum and Would Not Have Had a Reasonable Expectation of Success in Doing So

191. A person of ordinary skill in the art would not have been motivated to combine Ranum with Veselov, Mohanty, and Seo. I discussed above in relation to Ground 3 that a person of ordinary skill in the art would not have been motivated to combine Veselov, Mohanty, and Seo in the proposed manner, and that discussion applies here to Ground 4 because Ground 4 assumes the combination of Veselov, Mohanty, and Seo. My discussion above explaining that a person of ordinary skill in the art would not have been motivated to combine Veselov,

Mohanty, and Ranum (*see* Ground 2) also applies equally to Ground 4 because Ground 4 assumes the combination of Veselov, Mohanty, and Ranum. Furthermore, it is my opinion that POSA would not have selectively identified only the certain aspects of Seo to combine with the certain portions of Veselov, Mohanty, and Ranum that Dr. Stavrou relies on, unless that POSA had the 926 patent claims as a guide to select the references and portions. I understand that is an improper analysis for obviousness. I also note that Seo would not have led a POSA to the challenged claims because, for example, Seo does not describe or explain using any snapshot scanning or analysis of VMs in a cloud computing environment context for the detection of cyber threats based on data claimed in the patent, nor does Dr. Stavrou opine that it does.

E. Ground 5: Claim 11 Would Not Have Been Obvious Over Veselov, Mohanty, and Hutchins

192. Dr. Stavrou sets forth Ground 5 to “build[] upon” Ground 1 and for Hutchins’ alleged disclosures of “quarantining VMs.” Ex.1002 (Stavrou Decl.) ¶281. But Dr. Stavrou challenges claim 1, on which claim 11 depends, based on only Veselov and Mohanty. Accordingly, for at least the same reasons Dr. Stavrou fails to demonstrate obviousness of claim 1, it is also my opinion that Dr. Stavrou fails to demonstrate obviousness of claim 11.

193. 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

194. 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 5 because Ground 5 assumes the combination of Veselov and Mohanty. *See* Section VIII.A.

195. 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 arrive at what is claimed in the '926 patent. 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. This approach is fundamentally different than Mohanty's analysis, which focuses on VA scanners to continuously monitor target resources. *Compare* Ex.1070 (Hutchins), Abstract *with* Mohanty, Abstract; *see also* Section VII.A. I explained above with respect to Ground 1 the fundamental differences, known to a person of ordinary skill in the art, between analyzing snapshots and collecting data via agent-based scanners.

196. Hutchins's approach is also different than, and would not have been combined with, Veselov's analysis of a snapshot. 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.A. 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 does not discuss or suggest combining Hutchins's analysis of replicate virtual machines 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

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

198. Accordingly, Dr. Stavrou has failed to show that Claim 11 is obvious over Veselov, Mohanty, and Hutchins.

F. Ground 6: Claim 11 Would Not Have Been Obvious Over Veselov, Mohanty, Hutchins, and Ranum

199. Dr. Stavrou sets forth Ground 6 to “build[] upon” Ground 2 and for Hutchins’ alleged disclosures of “quarantining VMs.” Ex.1002 (Stavrou Decl.) ¶281. But Dr. Stavrou’s Ground 6 relies on his same positions from Grounds 1, 2, and 5. My analyses for Grounds 1, 2, and 5 therefore apply equally to Ground 6; Ground 6 thus fails for at least the same reasons.

200. Accordingly, Dr. Stavrou has failed to show that Claim 11 is obvious over Veselov, Mohanty, Hutchins, and Ranum.

IX. Objective Indicia of Non-obviousness

201. I have considered objective evidence that supports that a person of ordinary skill in the art would not have considered the claims of the ’926 patent to be obvious, as discussed below. It is my opinion that, even assuming that Wiz 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 '926 Patent Inventions

202. 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 '926 patent. In that section of his declaration, Mr. Shua maps Orca's SideScanning™ to the claims of the '926 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 '926 patent.

B. Long-Felt, Unresolved Need

203. 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 '926 patent

describes, “[o]rganizations [] increasingly adapted their applications to be run from multiple cloud computing platforms.” ’926 patent, 1:20-21. Virtualization technologies such as virtual machines, software containers, and serverless functions helped enable this shift to commercial cloud computing. *See* ’926 patent, 1:36-39.

204. 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 ’926 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. ’926 patent, 1:44-2:14.

205. Generally, these problems relate 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. '926 patent, 1:49-2:6. 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."); '926 patent, 2:13-14 ("[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.").

REDACTED

REDACTED

207. Preexisting techniques also failed to provide adequate detection of and protection for sensitive data. For example, a 2024 report describes that “The task of locating and classifying the entirety of an organization’s data across all of the possible locales is no small endeavor and is generally proving unmanageable for legacy solutions such as data loss prevention (DLP) tools.” Ex.2152 at 5.

208. 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.”).

209. Accordingly, there was a long-felt, unmet need for a cloud security solution that could provide both visibility and coverage to adequately address undetected cyber threats.

210. 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.

211. “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.

REDACTED

⁸ See '926 patent, cls. 1, 14, 15 (limitations 1.1-1.6, 14.1-13.6, 15.1-22.6).

REDACTED

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

213. 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.

214. 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

⁹ See '926 patent, cls. 1, 14, 15 (limitations 1.7, 14.7, 15.7).

(Future of Cloud Security), 8, 11. It was a “turning point – but anyone who tells you they knew it at the time is lying.” Ex.2076 (Future of Cloud Security), 8.

215. After Orca “introduced the concept of agentless side scanning,” it was “difficult” for “older providers to come to terms with [its] success.” Ex.2076 (Future of Cloud Security), 7. That 2024 report states that these providers were “skeptics” because “from a security perspective, agentless scanning was nothing compared to what attacks they were able to detect and protect.” Ex.2076 (Future of Cloud Security), 7. “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.” Ex.2076 (Future of Cloud Security), 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”).

REDACTED

REDACTED

REDACTED

REDACTED

219. Eventually, “everyone needed [agentless scanning] in order to stay competitive.” Ex.2076 (Future of Cloud Security), 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

220. 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)), 9-10. CRN touted "Orca's SideScanning technology" which "automatically assesses the security state of every discovered asset" to provide "complete visibility." Ex.2119 (*The 20 Coolest Tech Startups of 2020*, CRN (Nov. 19, 2020)), 9-10.
- 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)), 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." Ex.2120 (*The Top 20 Cybersecurity Startups to Watch in 2021 Based on Crunchbase*, Forbes (Nov. 20, 2020)), 6-7.
- 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”)

REDACTED

REDACTED

REDACTED

REDACTED

223. 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 '926 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.

Ex.2067 (Future of Cloud Security), 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.

Ex.2067 (Future of Cloud Security), 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.

Ex.2067 (Future of Cloud Security), 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.

Ex.2067 (Future of Cloud Security), 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.

Ex.2067 (Future of Cloud Security), 11-12

224. In summary, Orca's extensive praise for the claimed combination of features in the '926 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™

225. 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 '926 patent's claims obvious.

REDACTED

REDACTED

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

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

231. 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.

232. REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

256. I have also analyzed evidence that Orca introduced technology for analyzing snapshots for cyber threats based on virtual disk data in its SideScanning technology. For example, Orca's Ultimate Guide from June 2019 describes how

Orca analyzed the “data stack level” including “sensitive data such as PII or payment info” which is claimed in Limitations 1.4, 14.4, and 15.4 of the '926 patent. Ex.2090 (Ultimate Guide (June 2019)), 4, 19. And, Orca’s SideScanning brief from 2020 explains how Orca’s data collector searched entire file systems for sensitive data. Ex. 2067, (SideScanning Technical Brief (2020)), 14.

REDACTED

REDACTED

REDACTED

259. 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™

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

261. REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

263. 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

274. Exhibit 2155 shows that Wiz added “Data Security Posture Management” as a feature that “now detects [sensitive] data such as PII” in November 2022. Ex.2155 (Wiz Becomes the First CNAPP to Deliver Data Security Posture), 1, 3. After it copied Orca, Wiz touted the importance of detecting sensitive data on its website. *See* Ex.2156 (Wiz CSPM vs DSPM), 4

(“Identifying and securing sensitive data across environments”), 5 (“DSPM prioritizes data security by identifying data-targeted vulnerabilities and enforcing security policies. This is essential for organizations with large amounts of sensitive data and those in regulated industries.”); Ex.2157 (Sensitive Data Discovery), 2 (“[W]e’ll find out why the sensitive data discovery process is so important”), 5 (“With integrated DSPM, Wiz takes a load off your mind with a streamlined approach to data security. It finds and classifies data across all your cloud environments.”).

REDACTED

REDACTED

REDACTED

277. In the discussion that follows, I discuss additional exhibits I have reviewed that further supports the substantial similarity between Wiz and Orca's products as it relates to specific limitations of the claims of the '926 patent. ^{REDACTED}

- a. **Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.1-1.3, 14.1-14.3, 15.1-15.3.**

278. Limitations 1.1-1.3 recite "*receiving a request to scan a protected virtual cloud asset in the cloud computing environment*" and "*locating, using an API or service provided by the cloud computing environment, a snapshot of at least one virtual disk of the protected virtual cloud asset*" and "*accessing, using an API or service provided by the cloud computing environment, the snapshot of the at least one virtual disk.*" Wiz's Workload Scanner is substantially similar to Orca's SideScanningTM in relation to limitations 1.1-1.3, 14.1-14.3, and 15.1-15.3 of claims 1, 14, and 15 of the '326 Patent.

279. For example, similar to Orca's SideScanning™, 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.

280. 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.

REDACTED

282. 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

REDACTED

REDACTED

REDACTED

- b. Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.4, 14.4, 15.4.

284. Limitation 1.4 recites “*analyzing the snapshot of the at least one virtual disk to determine the existence of a plurality of potential cyber threats, each cyber threat based on data stored on the virtual disk, wherein the data includes at least one of: unencrypted sensitive data, unencrypted system credentials, weak passwords, weak encryption schemes, disabled Address Space Layout Randomization, boot record manipulation, suspicious definitions, services to be run on startup, personally identifiable information, data in application logs*

indicating that the protected virtual cloud asset accessed personally identifiable information, data in application logs indicating that the protected virtual cloud asset accessed a computer containing personally identifiable information, or at least one change in at least one area of the virtual disk, as compared to an earlier point in time.” Wiz’s Workload Scanner is substantially similar to Orca’s SideScanning™ in relation to limitations 1.4, 14.4, and 15.4 of claims 1, 14, and 15 of the ’326 Patent (reproduced in full above).

285. For example, Wiz’s workload scanner “detects data such as PII [personally identifiable information], PHI, and PCI and adds this as a new risk factor to the Wiz Security Graph.” Ex.2155 at 3.

Prioritize and stop attack paths targeting your most critical cloud data


Wiz Data Security Posture Management helps organizations discover which data is stored where, who can access what data, how data assets are configured and utilized across human and non-human identities, and how data moves across environments. Wiz now detects data such as PII, PHI, and PCI and adds this as a new risk factor to the Wiz Security Graph to enable:

Ex.2155 at 3

286. Wiz’s Workload Scanner also “focuses on finding and securing sensitive data across your network and cloud environment.”

What is DSPM?

[Data security posture management](#) (DSPM) focuses on finding and securing sensitive data across your network and cloud environments. DSPM identifies and fixes weaknesses like misconfigurations and excessive permissions that could lead to a data breach.



Ex.2156 at 3

With integrated DSPM, Wiz takes a load off your mind with a streamlined approach to data security. It finds and classifies data across all your cloud environments, including PaaS, IaaS, and DBaaS, cutting your risk of shadow data. Wiz lets you enforce data security controls and policies uniformly, regardless of how and where your data is stored.




Figure 2: Wiz's all-in-one dashboard gives you uniform control over all your cloud environments

Ex.2157 at 5-6

Wiz is extending its cloud-native application protection platform (**CNAPP**) with integrated data security posture management (DSPM) capabilities to detect cloud data exposure and prevent data breaches. Customers can now continuously monitor for data exposure before it becomes a costly breach and arm their teams with all the context they need to remediate issues. Wiz improves data security posture efficiently and allows teams to focus on prevention. This lowers the risk of a financial impact due to exposure or leakage of critical data. And, by breaking down silos between cloud security, data governance, and development teams, organizations drive significant productivity gains from reducing friction and the ability to seamlessly collaborate. This is one of the many ways in which Wiz helps its customers transform their **cloud operating model**. Today, our DSPM capabilities are in public preview and available for Wiz customers to begin using.

Ex.2155 at 5

287. And Wiz's workload scanner "gives you visibility into data security risks and compliance."

Adopting DSPM for sensitive data discovery and classification

Data security posture management (DSPM) is one of the best sensitive data discovery tools, letting you continuously manage and secure sensitive data across multiple cloud environments.

DSPM gives you visibility into data security risks and compliance. With customizable rules, it automates discovery and classification. DSPM then provides ongoing monitoring, alerting, and policy enforcement. And with clear, meaningful reporting on your data security posture from DSPM, regulatory compliance becomes much simpler.

Ex.2157 at 4.

- c. **Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.5-1.7, 14.5-14.7, 15.5-15.7**

288. Limitations 1.5-1.7 recite "*determining a risk associated with each of the determined plurality of potential cyber threats*" and "*prioritizing the potential cyber threats associated with the protected virtual cloud asset based on the determined risk associated with each of the plurality of potential cyber threats*" and "*reporting at least some of the determined plurality of potential cyber threats as alerts prioritized according to their associated risks.*" Wiz's Workload Scanner uses the features recited in limitations 1.5-1.7, 14.5-14.7, 15.5-15.7 of claims 1, 14, and 15 of the '926 Patent (reproduced in full above).

289. For example, as detailed above, Wiz's Workload Scanner can "[p]rioritize and stop attack paths" and "provides ongoing monitoring, alerting, and policy enforcement." Additionally, Wiz's Workload Scanner can "build a single prioritized queue" of attack paths.

Prioritize and stop attack paths targeting your most critical cloud data

Wiz Data Security Posture Management helps organizations discover which data is stored where, who can access what data, how data assets are configured and utilized across human and non-human identities, and how data moves across environments. Wiz now detects data such as PII, PHI, and PCI and adds this as a new risk factor to the Wiz Security Graph to enable:

Ex.2155 at 3

- **Continuous detection and prioritization of critical data exposure:** Wiz conducts a deep cloud analysis that automatically correlates data risks with other cloud risks to build a single prioritized queue of attack paths and toxic combinations of risk to reduce noise and focus teams on what is important.

Ex.2155 at 4.

290. Wiz's Workload Scanner also provides ongoing "alerting":

Adopting DSPM for sensitive data discovery and classification

Data security posture management (DSPM) is one of the best sensitive data discovery tools, letting you continuously manage and secure sensitive data across multiple cloud environments.

DSPM gives you visibility into data security risks and compliance. With customizable rules, it automates discovery and classification. DSPM then provides ongoing monitoring, alerting, and policy enforcement. And with clear, meaningful reporting on your data security posture from DSPM, regulatory compliance becomes much simpler.

Ex.2157 at 4.

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

The diagram illustrates the components and application of Data Security Posture Management (DSPM). It features a blue header box labeled 'DSPM', followed by three stacked light blue boxes describing its functions: 'Protecting sensitive data', 'Identifying and securing sensitive data across environments', and 'Secure the entire cloud environment'. Below these is a light blue box listing target users: 'Organizations with large amounts of sensitive data, organizations in highly regulated industries (e.g., healthcare, finance)'. At the bottom, a text block describes a use case where DSPM at a healthcare provider identifies unencrypted 'shadow data' on a cloud server, alerts security, and helps with remediation steps.

DSPM

Protecting sensitive data

Identifying and securing sensitive data across environments

Secure the entire cloud environment

Organizations with large amounts of sensitive data, organizations in highly regulated industries (e.g., healthcare, finance)

DSPM at a healthcare provider discovers that a large amount of patient data is stored on a cloud server without proper encryption. This unknown "shadow data" poses a significant security risk. DSPM alerts security, identifies the data, pinpoints its location, and helps implement risk-remediation steps.

Ex.2156 at 4.

With integrated DSPM, Wiz takes a load off your mind with a streamlined approach to data security. It finds and classifies data across all your cloud environments, including PaaS, IaaS, and DBaaS, cutting your risk of shadow data. Wiz lets you enforce data security controls and policies uniformly, regardless of how and where your data is stored.




Figure 2: Wiz's all-in-one dashboard gives you uniform control over all your cloud environments

Ex.2157 at 5-6

One such customer is Chevron Phillips Chemical Company, which uses Wiz to discover and protect mission-critical data. For Cory Zaner, Cloud Security Manager, "We are not the data governance team, but we want to proactively protect our data in the cloud. The visibility that Wiz gives us into our data and how it maps to external exposure is key as we don't want to be in the news." Asonye Onwudebe, Cloud Security Architect, explains further, "Wiz is our early warning system for critical data risks. Wiz alerts us with laser precision and all the context needed to take fast proactive action to prevent issues from becoming data breaches."

Ex.2155 at 3.

F. Commercial Success

292. 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

293. 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). And after Wiz introduced sensitive data scanning in its "DSPM" feature in November 2022, which as discussed above it appears Wiz

copied from Orca's SideScanning, Wiz's annual recurring revenue purportedly grew from ~\$100M to \$500M in the next two years. Ex.2153

294. It is my opinion that the commercial success of both Orca's and Wiz's products has nexus to the claims of the '926 patent. For example, both systems practice the claims of the '926 patent as discussed above. Additionally, both entities tout their usage of the claimed features in their marketing materials. *See, e.g.,* Ex.2107 at 2-3 (touting Orca's features such as "[w]here is my sensitive data stored" in an announcement about fundraising). REDACTED

X. Conclusion

295. For the foregoing reasons, it is my opinion that claims 1-15 of the '926 patent would not have been obvious over the prior art in Wiz's Petition for *inter partes* review.

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, cursive script.

Dr. David Kaeli