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.,
Petitioner,

v.

ORCA SECURITY LTD.,
Patent Owner.

IPR2024-01191
Patent 11,775,326 B2

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

Orca Security Ltd.
Exhibit 2073
Wiz v. Orca
IPR2024-01191

Table of Contents

I. Introduction And Qualifications 3

II. Qualifications 4

III. Summary Of Materials Reviewed And Considered 7

IV. Legal Standards 8

V. Level Of Skill In The Art And Perspective Applied In This Declaration 12

VI. Background 15

A. Overview of the '326 Patent 15
B. Claim Construction 21
1. "Analyzing the Snapshot" 21

VII. Overview of the Asserted References 24

A. Overview of Veselov 24
B. Overview of Basavapatna 32
C. Overview of Czarny 36
D. Overview of Giakouminakis 37

VIII. The Claims Of The '326 Patent Would Not Have Been Obvious ... 37

A. Ground 1: Claims 1-21 and 28 Would Not Have Been Obvious Over Veselov and Basavapatna 38
1. A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov and Basavapatna and Would Not Have Had a Reasonable Expectation of Success in Doing So 38
2. Even If Combined, Veselov and Basavapatna Do Not Disclose or Suggest Multiple Limitations in the Claims 71
B. Ground 2: Claims 4-5 and 17 Would Not Have Been Obvious Over Veselov, Basavapatna, and Czarny 85
C. Ground 3: Claims 22-27 Would Not Have Been Obvious Over Veselov, Basavapatna, and Giakouminakis 86

PROTECTIVE ORDER MATERIAL
REDACTED

IX. Objective Indicia of Non-obviousness 89

A. Orca's SideScanning™ Uses The '326 Patent Inventions 89
B. Long-Felt, Unresolved Need 90
C. Orca's claimed inventions overcame industry skepticism to become widely accepted--and widely praised—in the industry. 95
D. Extensive Industry Praise Supports Non-Obviousness of the Claims 99
E. Evidence Shows Wiz Copied Orca's SideScanning™..... 107
1. Evidence shows Wiz had access to Orca's invention 109
2. Evidence shows Wiz's Workload Scanning is substantially similar to Orca's SideScanning™..... 126
F. Commercial Success 153

X. Conclusion..... 154

PROTECTIVE ORDER MATERIAL
REDACTED

I. Introduction And Qualifications

1. I, David Kaeli, submit this declaration on behalf of Orca Security Ltd. ("Patent Owner") to provide my opinions concerning the validity of claims 1-28 of U.S. Patent No. 11,775,326 (the "'326 patent," Ex.1001), in support of Patent Owner's Response to the Petition for *Inter Partes* review of the '326 patent, which I understand was filed by Wiz, Inc. ("Petitioner") and accompanied by a Declaration of Dr. Angelos Stavrou (the "Stavrou Declaration"), among other exhibits. I further understand that on January 22, 2025, this *Inter Partes* review (IPR) was instituted (Paper 9).

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

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

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

PROTECTIVE ORDER MATERIAL
REDACTED

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.

PROTECTIVE ORDER MATERIAL
REDACTED

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.

PROTECTIVE ORDER MATERIAL
REDACTED

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

PROTECTIVE ORDER MATERIAL
REDACTED

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 '326 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."

PROTECTIVE ORDER MATERIAL
REDACTED

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.

PROTECTIVE ORDER MATERIAL
REDACTED

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

PROTECTIVE ORDER MATERIAL
REDACTED

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

PROTECTIVE ORDER MATERIAL
REDACTED

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

PROTECTIVE ORDER MATERIAL
REDACTED

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.

PROTECTIVE ORDER MATERIAL
REDACTED

29. With respect to the '326 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 '326 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 '326 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

PROTECTIVE ORDER MATERIAL
REDACTED

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, 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 '326 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.

PROTECTIVE ORDER MATERIAL
REDACTED

VI. Background

A. Overview of the '326 Patent

35. The '326 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 ('326 patent), Abstract, cls. 1, 15, 18.

36. The '326 patent explains that, around the time of its priority date (January 2019), public cloud service providers, such as Amazon, Microsoft, and Google, have "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 ('326 patent), 1:20-30. The '326 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 ('326 patent), 1:31-33.

37. The '326 patent discusses problems that existed with trends in the cybersecurity field. For example, the '326 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 ('326 patent), 1:42-46. And the '326 patent also

PROTECTIVE ORDER MATERIAL
REDACTED

explains ways the industry had tried to deal with such issues. For example, the '326 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 ('326 patent), 1:46-2:14. As the '326 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 ('326 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 ('326 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 ('326 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 ('326 patent), 2:14-16. "For example, installing agents in

PROTECTIVE ORDER MATERIAL
REDACTED

a large data center may take months." Ex.1001 ('326 patent), 2:16-17. The '326 patent's systems and method "overcome the[se] deficiencies." Ex.1001 ('326 patent), 2:18-19.

38. The '326 patent aims to solve such problems. Ex.1001 ('326 patent), 3:35-36, 3:47-63. Figure 1B of the '326 patent, reproduced below, illustrates a diagram of a cloud computing platform to which the '326 patent's techniques for scanning for embedded vulnerabilities in machines:

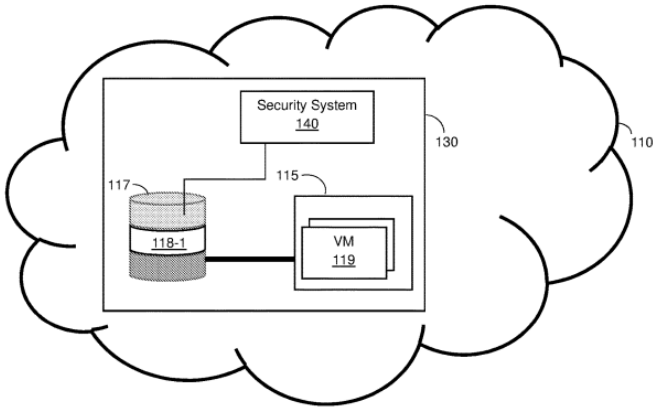

FIG. 1B

39. The cloud computing platform of the '326 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 ('326 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 ('326 patent), 3:50-51, 3:55-60. "The storage may be part

PROTECTIVE ORDER MATERIAL
REDACTED

of the server or, as shown in Figure 1B, separated from the server. Ex.1001 ('326 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 ('326 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 ('326 patent), 4:16-21. The security system is configured to perform the steps of the '326 patent's invention and may be separate from the client environment, as shown in Figure 1B, or part of the client environment. Ex.1001 ('326 patent), 4:1-17. The challenged claims of the '326 patent demonstrate the embodiment where the security system interfaces with the client environment through APIs. Ex.1001 ('326 patent), cls. 1, 15, 18.

40. The security system of the '326 patent interfaces with the client environment so that it can perform scanning operations and detect any "vulnerable software, non-secure configurations, exploitation attempts, compromised assets, data leaks, data mining" and other risks or vulnerabilities. Ex.1001 ('326 patent), 4:50-56. For example, the security system can detect vulnerable code (Ex.1001 ('326 patent), 5:31-33), cyber threats that do not represent vulnerabilities such as unencrypted sensitive data (Ex.1001 ('326 patent), 6:21-32), and changes in sensitive machine areas such as changed application files without installation

PROTECTIVE ORDER MATERIAL
REDACTED

(Ex.1001 ('326 patent), 6:33-39). The security system is also configured to provide "incident response, anti-ransomware, and cyber-insurance" through access to the system's security posture (Ex.1001 ('326 patent), 4:50-56), generate reports on detected vulnerabilities and cyber threats, including, "for example, as alerts" (Ex.1001 ('326 patent), 7:53-55), and mitigate detected vulnerabilities and cyber threats. Ex.1001 ('326 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 ('326 patent), 8:1-6.

41. In order to accomplish the scanning and detection operations I discuss above, the '326 patent provides a series of steps. First, the security system queries for and receive a location of a snapshot of a virtual disk (Ex.1001 ('326 patent), 4:57-60). For example, the security system queries a cloud management console for the location of a snapshot. Ex.1001 ('326 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 ('326 patent), 4:58-59. After the snapshot is located, it is accessed by the security system. Ex.1001 ('326 patent), 5:3-4. The security system then analyzes the snapshot to determine the existence of potential cyber vulnerabilities, without installing agents on the server or the VM, and the snapshot analysis does not involve interaction with or information from the VM. Ex.1001 ('326 patent), 5:20-25, 7:36-40. The security system may determine the risk level associated with

the potential cyber vulnerabilities. Ex.1001 ('326 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 ('326 patent), 7:53-67. The security system then generates reports on detected vulnerabilities and cyber threats and sensitive data, including, "for example, as alerts." Ex.1001 ('326 patent), 6:40-48.

42. The '326 patent provides multiple techniques for analyzing snapshots to detect different types of vulnerabilities and cyber threats. Ex.1001 ('326 patent), 7:59-8:6. For example, the snapshot may be parsed and analyzed by the security system, the system may match installed application lists to a known list of vulnerable applications, and the system may check configuration files of applications. Ex.1001 ('326 patent), 5:20-25, 5:34-36, 5:48-55. The scanning service then generates reports on detected vulnerabilities and cyber threats, including, "for example, as alerts" that may be filtered and prioritized. Ex.1001 ('326 patent), 7:53-55. The security systems and methods also "may cause an execution of one or more mitigation actions," including disabling a VM from execution, updating a VM with patches, blocking traffic from untrusted networks, quarantining an infected VM, and preventing a VM from launching. Ex.1001 ('326 patent), 6:61-67, 8:1-6.

PROTECTIVE ORDER MATERIAL
REDACTED

B. Claim Construction

43. In addition to the explicit constructions discussed below, I have applied the plain and ordinary meanings of the claim terms in the '326 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. Claim 1 recites "analyzing the snapshot of the at least one virtual disk to determine an existence of potential cyber vulnerabilities." Similarly, claims 15 and 18 recite "analyzing the snapshot of the at least one virtual disk to determine existence of a plurality of potential cyber vulnerabilities." Dr. Stavrou labels these as elements 1.2.c, 15.2.c, and 18.2.c. *See, e.g.*, Ex.1002 (Stavrou Decl.), ¶153.

45. Dr. Stavrou states that he was "asked to assume that '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 'analyzing the snapshot' in claims 1, 15, and 18 as encompassing both direct analysis of the snapshot data and analysis of a VM instantiated from snapshot."). Thus, Dr. Stavrou treats limitations in the challenged claims that recite "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.

PROTECTIVE ORDER MATERIAL
REDACTED

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.

Q. Here you state, "A POSA would have understood the ordinary and customary meaning of '[analyze/analyzing] the snapshot' as at least encompassing direct analysis of the snapshot data"; is that right?
A. That is correct.
Q. And your opinion is that that would be sort of the correct meaning of the term "analyze/analyzing the snapshot" in the context of the '926 patent?
MR. ARGENTI: Objection. Form.
A. I believe that I -- they were consistent. So I say, "This understanding is consistent with the '926 specification."

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

PROTECTIVE ORDER MATERIAL
REDACTED

Q. And I think in -- in your prior deposition in the context of the '031, '032, and '685 patents, you said that you were just asked to assume that by counsel and you didn't form an opinion about whether it's correct. Is that also true here?
A. I don't know if this is exactly what I said in my prior deposition, but I was asked to assume that analyzing and analyze the snapshot also encompasses analysis of a VM instantiated from the snapshot. That is correct.
Q. But you don't have an opinion as a technical expert that that is a correct -- a correct way of interpreting analyzing the snapshot consistent with the '926 patent specification; right?
A. I was not asked to opine on that, so I don't 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 "analyzing the snapshot" limitations in the '326 patent to encompass "analysis of a VM instantiated from a snapshot." As an initial matter, there appears to be no disagreement that an instantiated VM is not a snapshot—otherwise Dr. Stavrou would not have enumerated it as separate from "direct analysis of snapshot data" in his assumed construction. Rather, an instantiated VM is a VM, and analysis of a VM instantiated from a snapshot is an analysis of a VM—not a snapshot. Once a VM is instantiated from a snapshot, e.g., when a duplicate VM is instantiated

PROTECTIVE ORDER MATERIAL
REDACTED

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 '326 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 ('326 patent), 4:60-61-59. The specification also distinguishes analyzing a snapshot by a "security system 140 to detect vulnerabilities" (Ex.1001 ('326 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" (*id.* 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

PROTECTIVE ORDER MATERIAL
REDACTED

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

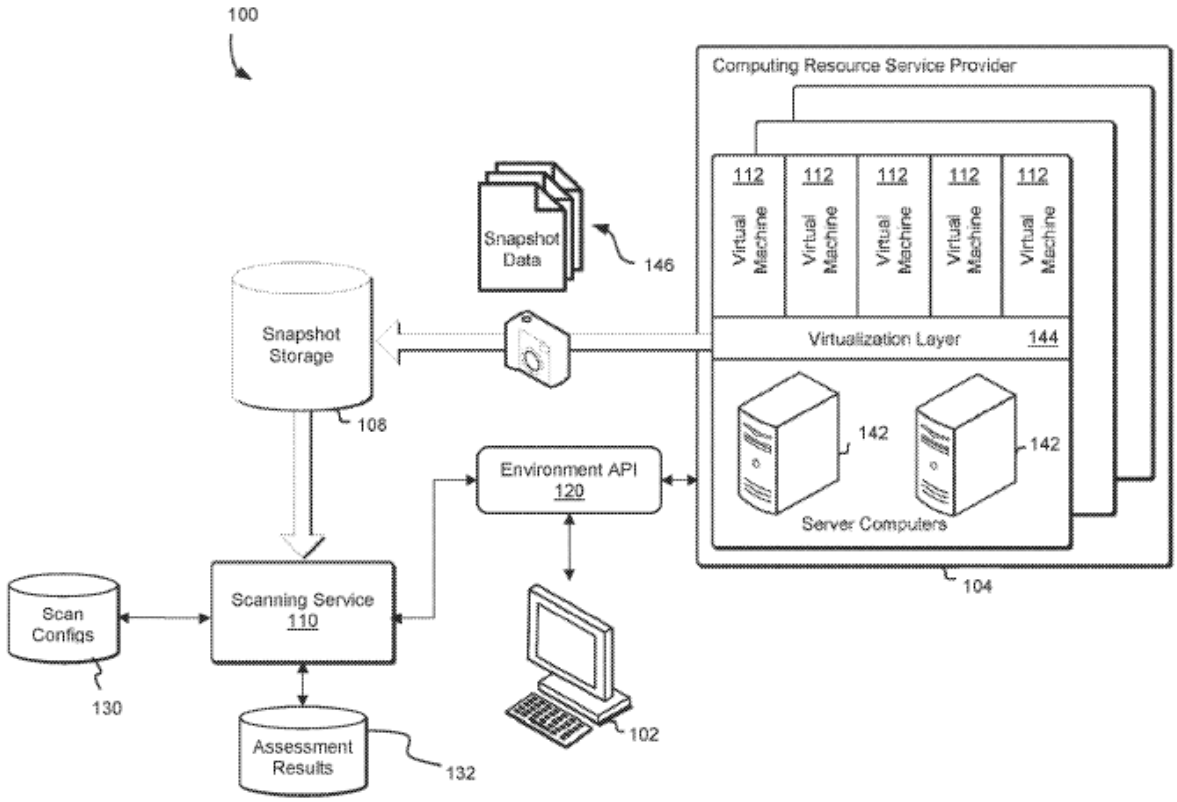

FIG. 1

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

PROTECTIVE ORDER MATERIAL
REDACTED

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.

51. 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.

PROTECTIVE ORDER MATERIAL
REDACTED

52. 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.

53. 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.

54. 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.

PROTECTIVE ORDER MATERIAL
REDACTED

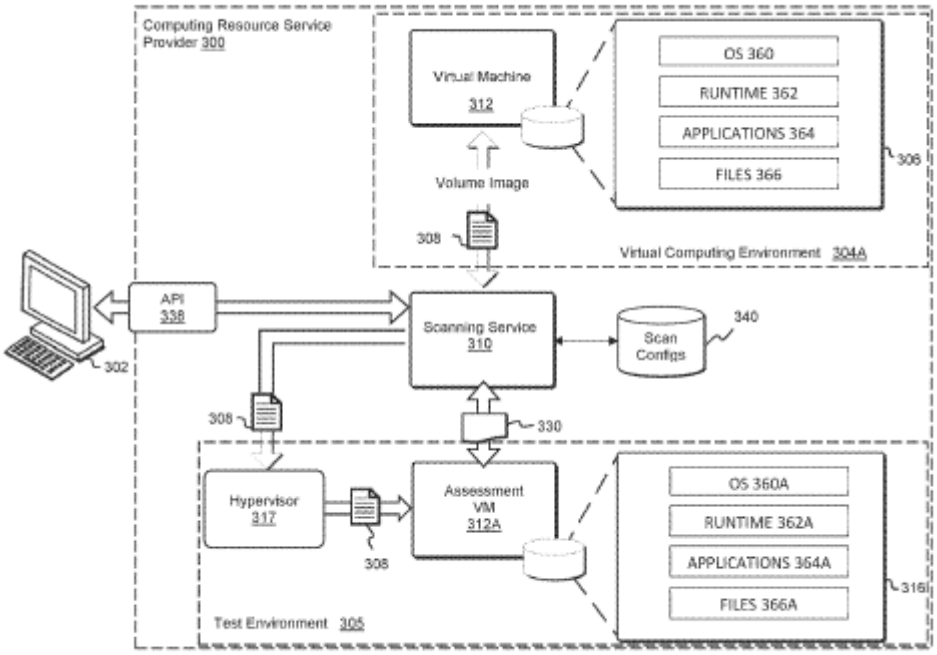

FIG. 3A

55. In Figure 3A, Veselov describes analyzing the duplicate virtual machine, and explains that the scanning service performs the security assessment on the assessment machine via a security agent program installed on the assessment machine. 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. Figure 3B is similar: Veselov analyzes a duplicate virtual machine by installing a security agent program on the assessment machine. *See* Ex.1007 (Veselov), FIG. 3B.

56. 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.

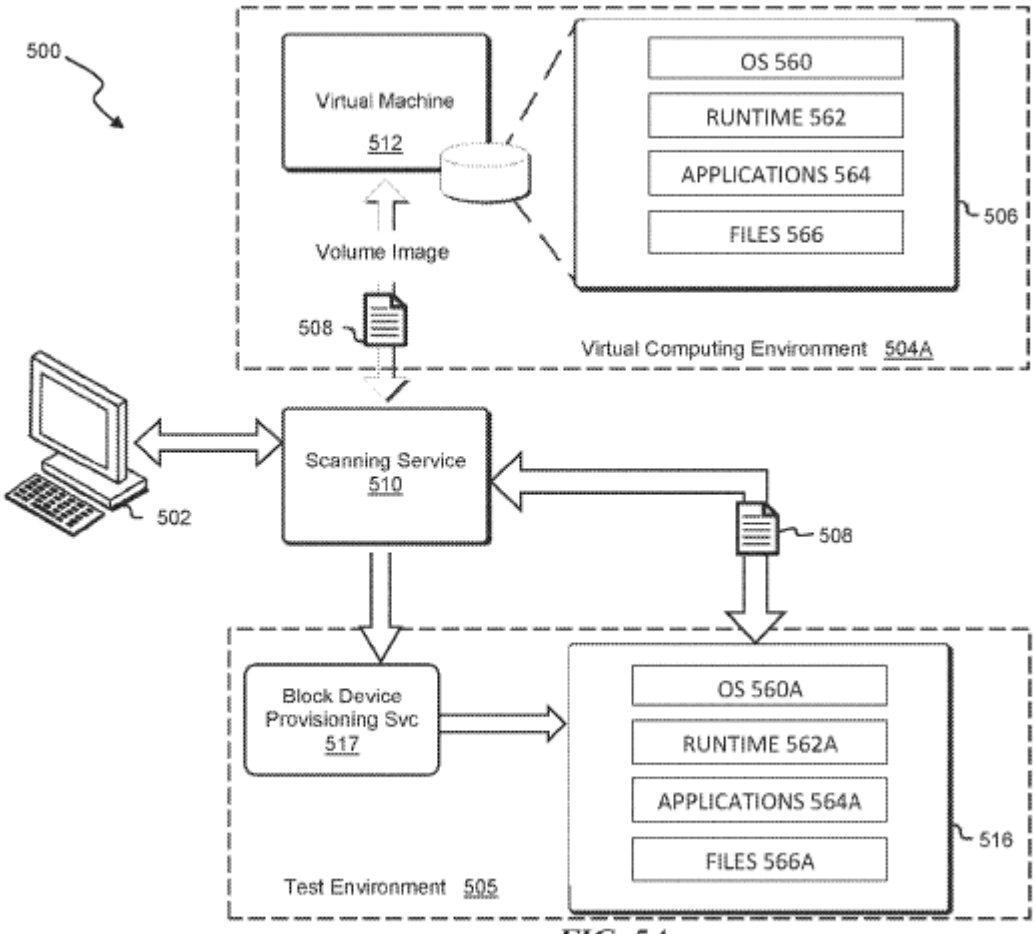

FIG. 5A

PROTECTIVE ORDER MATERIAL
REDACTED

57. 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. Additionally, Veselov's Figures 5A and 5B implementations do not install a security agent program on the snapshot that is to be analyzed.

58. 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.

59. 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

PROTECTIVE ORDER MATERIAL
REDACTED

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 Basavapatna

60. Basavapatna is titled "Calculating Quantitative Asset Risk" and published on July 25, 2013 from an application filed on January 19, 2012. Basavapatna, Cover.

61. Basavapatna relates to "the field of computer security assessment and, more particularly, to calculating risk metrics for assets in a system of computing assets." Ex.1008 (Basavapatna), ¶1. The "asset" described in Basavapatna is a physical "computer or other electronic device" that is in "a home" or "office[]," for example, and may be connected over a network to other assets via the Internet. Ex.1008 (Basavapatna), ¶2. Basavapatna does not discuss snapshots or virtual machines at all. *See, e.g.*, Ex.2172 (2nd Stavrou Tr.) 220:20-221:3 ("Q. … Basavapatna is directed to security assessments of physical computing assets; right? A. Correct."), 221:15-222:12 (agreeing no disclosure of "virtual" assets).

62. Basavapatna explains that each asset in a system of assets may be "at risk" due to "multiple threats at any given time" and that "[e]ach threat can correspond to a potential attack on the asset by a particular virus malware, or other unauthorized entity." Ex.1008 (Basavapatna), ¶2. Additionally, Basavapatna notes current drawbacks in security assessments of computers, noting that "some

PROTECTIVE ORDER MATERIAL
REDACTED

known vulnerabilities may not be associated with known threats." Ex.1008 (Basavapatna), ¶2.

63. Basavapatna's techniques therefore identify a "standardized vulnerability score" for "a particular vulnerability in a plurality of known vulnerabilities" where "the standardized vulnerability score indicat[es] a relative level of risk associated with the particular vulnerability relative other vulnerabilities." Ex.1008 (Basavapatna), Abstract. Additionally, Basavapatna's processes provide for a "countermeasure component score" and a "risk metric" with respect to a "particular asset." Ex.1008 (Basavapatna), Abstract. The countermeasure component score is "identified" and "indicates an estimated probability that a countermeasure will mitigate risk associated with the particular vulnerability on the particular asset." Ex.1008 (Basavapatna), Abstract. The "risk metric" is "for the particular asset" and "the particular vulnerability is determined from the vulnerability composite score and the countermeasure component score." Ex.1008 (Basavapatna), Abstract.

PROTECTIVE ORDER MATERIAL
REDACTED

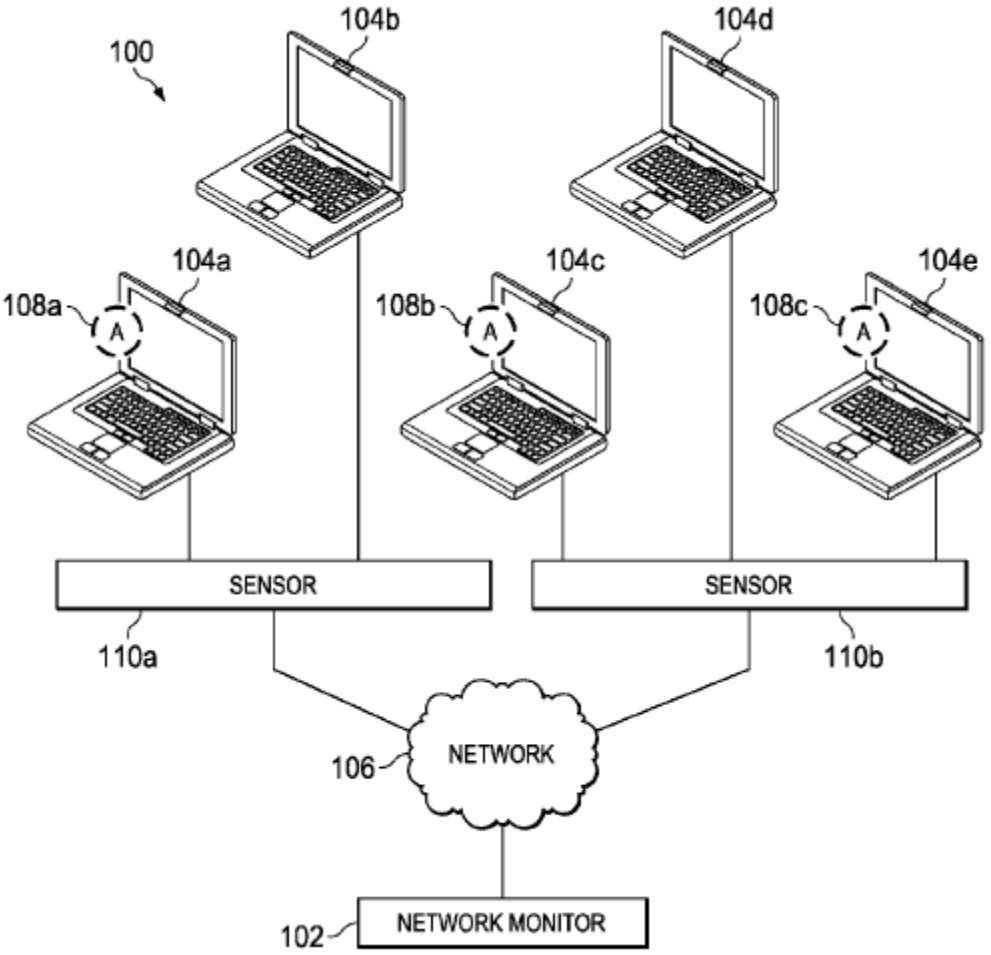

64. Basavapatna's Figure 1 provides an "example asset system" that is "monitored by a network monitor." Ex.1008 (Basavapatna), ¶19. The "asset system" includes assets "connected to each other" and/or to "other systems" via a network. Ex.1008 (Basavapatna), ¶19. Basavapatna discloses "host-based sensors" and "network-based sensors" in its asset system. Ex.1008 (Basavapatna), ¶¶26-28. The "host-based sensors" can include "agent-based or otherwise software-based sensors that are installed on respective assets." Ex.1008 (Basavapatna), ¶26. The "host-based sensors" can "run various analyses" on the

system assets and may "identify vulnerabilities on the assets" or "identify viruses or other malware executing on the assets." Ex.1008 (Basavapatna), ¶26. These "host-based sensors" may additionally "provide one or more passive countermeasures for threats" such as "antivirus and other antimalware software." Ex.1008 (Basavapatna), ¶26.

65. Basavapatna's system also includes "network-based sensors" which are "hardware devices and/or software in a data communication path" "protected by the sensor and the network resources that the sensor is attempting to access." Ex.1008 (Basavapatna), ¶27. Basavapatna explains that when an asset tries to send or receive information through or over the network through a network-based sensor, "the sensor analyzes information about the asset" along with "the information being sent or received and determines whether to allow the communication." Ex.1008 (Basavapatna), ¶28.

66. Additionally, Basavapatna explains that the host-based sensors and network-based sensors "monitor the assets themselves and/or network traffic to and from the assets." Ex.1008 (Basavapatna), ¶23. Sensor, as used in Basavapatna, refers to "various types of monitoring and protection systems including, for example, firewalls, host intrusion prevention systems, network intrusion prevention systems, network access control systems, intrusion detection systems, anti-virus software, and spam filters." Ex.1008 (Basavapatna), ¶23. And,

PROTECTIVE ORDER MATERIAL
REDACTED

the assets in Basavapatna's system may be "vulnerable to many different threats at any given time." Ex.1008 (Basavapatna), ¶30. Accordingly, Basavapatna provides for a "risk metric for each asset and each threat" to address these vulnerabilities. Ex.1008 (Basavapatna), ¶23. The "risk metric is a quantitative measure of the risk that a threat poses to an asset, both in terms of the probability that the threat will affect the asset and the magnitude of the effect that the threat will cause." Ex.1008 (Basavapatna), ¶23.

67. Basavapatna sets forth example processes for either a "threat-centric risk metric for an asset and a threat" and a "vulnerability-centric risk metric for an asset and vulnerability." Ex.1008 (Basavapatna), ¶¶74, 100. Basavapatna explains that the system "aggregate[]" risk metrics "on a per-threat or per-vulnerability basis" to calculate an "aggregate risk metric for the particular vulnerability from the risk metrics for the assets for [a] particular vulnerability." Ex.1008 (Basavapatna), ¶¶122-124.

C. Overview of Czarny

68. Czarny is titled "Computer Security Vulnerability Assessment" and describes that an assessment "is performed with product binary data and product vulnerability data that correspond with product identification data." Ex.1084 (Czarny) Abstract. I understand that Czarny is relied upon by Dr. Stavrou and

Petitioner for its disclosures of using "binary hashes of binary level files" of software products. *See, e.g.*, Ex.1084 (Czarny) 4:7-15, 11:16-12:10.

D. Overview of Giakouminakis

69. Giakouminakis describes a "security tool" which "can identify vulnerabilities in a computing system and determine a risk level of the vulnerabilities based on base and optional CVSS vectors and additional factors that represent the evolving nature of vulnerabilities." Ex.1044 (Giakouminakis) Abstract. Giakouminakis states that the "security tool 102 can be utilized to test and analyze the security of the computing systems 104 and/or the VMs 114" and "can be configured to run on one or more of the computing systems 104." Ex.1044 (Giakouminakis) 3:36-40.

VIII. The Claims Of The '326 Patent Would Not Have Been Obvious

70. Petitioner and Dr. Stavrou challenge claims 1-28 of the '326 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 Pub. No. 20130191919 ("Basavapatna," Ex.1008).

71. In my opinion, none of the claims of the '326 patent would have been obvious over Petitioner's and Dr. Stavrou's prior art, for reasons I explain below.

PROTECTIVE ORDER MATERIAL
REDACTED

A. Ground 1: Claims 1-21 and 28 Would Not Have Been Obvious Over Veselov and Basavapatna

72. In Ground 1 of the Petition, Petitioner and Dr. Stavrou challenge claims 1-21 and 28 of the '326 patent based on the combination of Veselov and Basavapatna. *See* Ex.1002, Stavrou Decl. §XII.

73. In my opinion, Petitioner and Dr. Stavrou have failed to show that any of claims 1-21 and 28 of the '326 patent would have been obvious over the combination of Veselov and Basavapatna.

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

74. Dr. Stavrou asserts that a person of ordinary skill in the art would have been motivated to make the combination of Veselov and Basavapatna on which all of Petitioner's grounds depend. *See* Ex.1002 (Stavrou Decl.), ¶¶104-113. I disagree. To a person of ordinary skill in the art, Veselov and Basavapatna 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 Basavapatna with the disclosures that Petitioner and Dr. Stavrou rely on for their challenges to the '326 patent claims. Indeed, in view of these references' distinctions, which I discuss in more detail below, it appears to me that Petitioner and Dr. Stavrou are working backwards from the '326 patent's claims, looking for disclosures in

PROTECTIVE ORDER MATERIAL
REDACTED

disparate references, without meaningfully considering whether the references they cite would, in fact, be combined in the first place. 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**

75. Dr. Stavrou's analysis mixes disclosures from Figures 1, 3A, 3B, 5A, and 5B in Veselov without explaining why a person of ordinary skill in the art would have wanted to combine or modify them. *See, e.g.*, Ex.1002 (Stavrou Decl.), ¶153 (opining about "scanning service" in Figures 1, 3A-3B, and 5A-5B). However, 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 Basavapatna.

76. Specifically, 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.

77. Figure 3A is reproduced below.

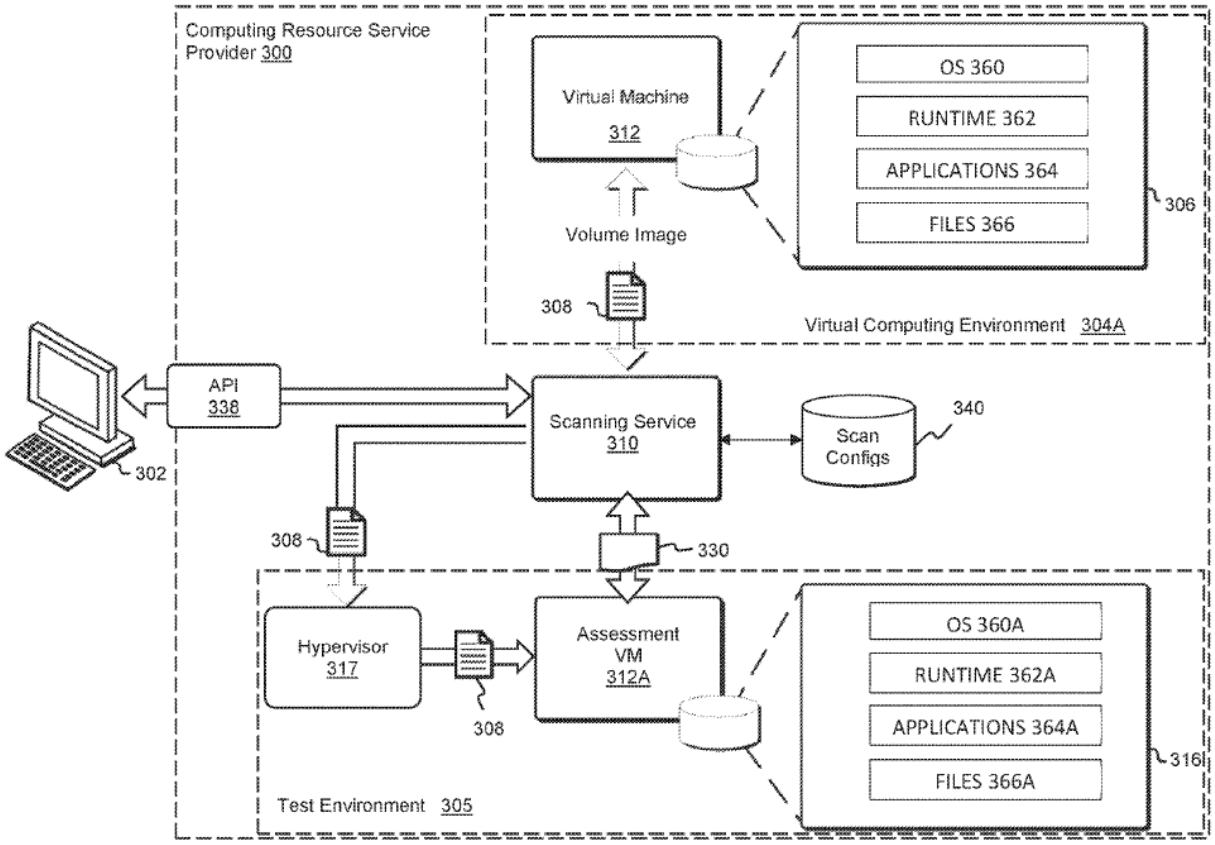

FIG. 3A

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. Figure 3B is similar to Figure 3A.

78. 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.

79. 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.

80. As I explained above in the Claim Construction section, the use of an assessment virtual machine instantiated from a snapshot, as in this example of

PROTECTIVE ORDER MATERIAL
REDACTED

Veselov, does not teach analyzing the snapshot as required by the claims. *See* Section VI.B.1.

81. 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.

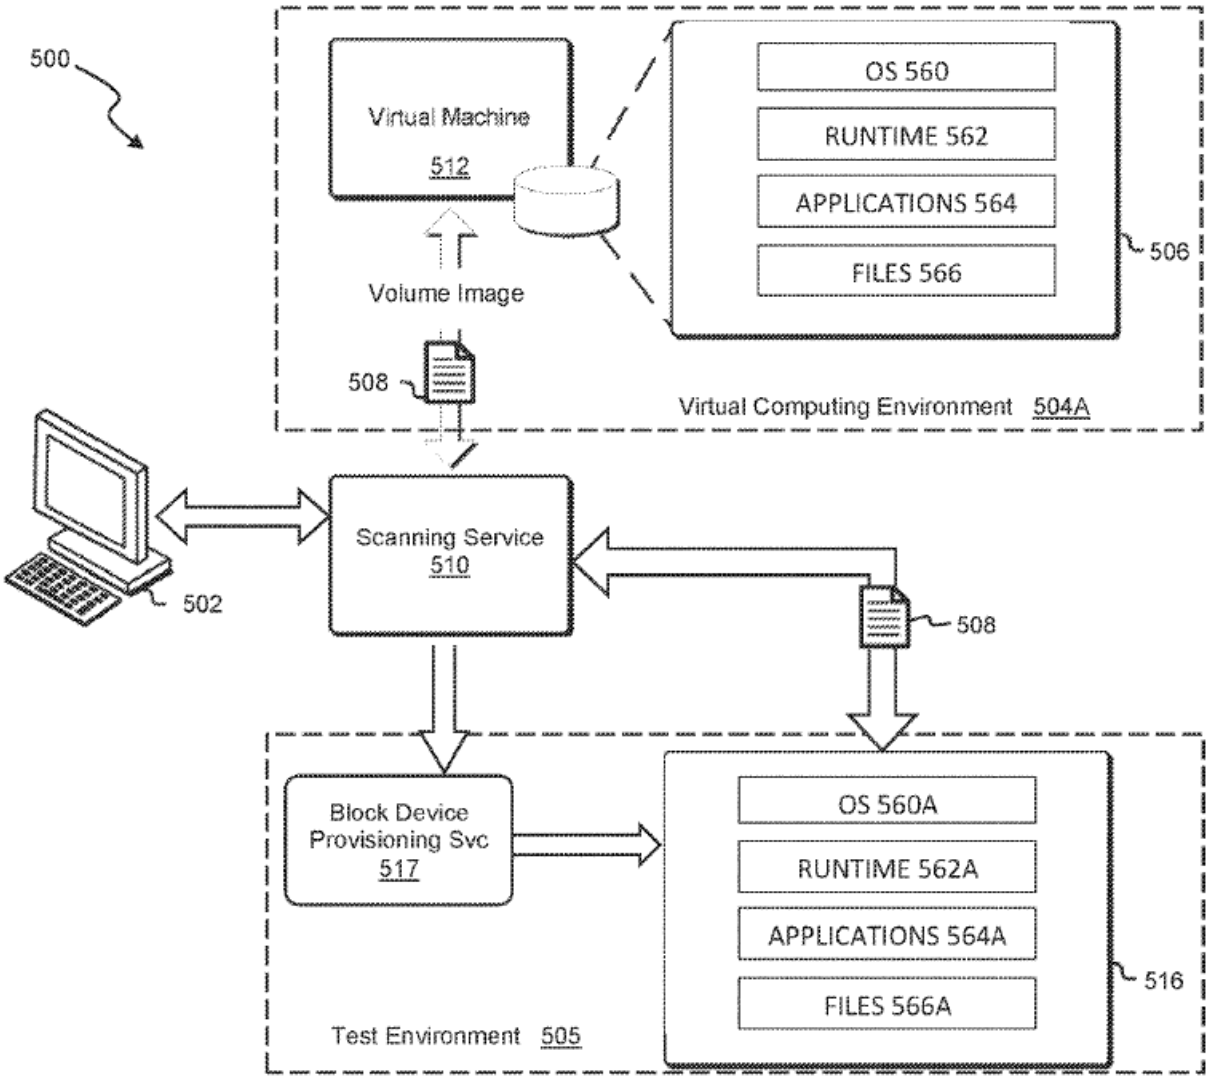

FIG. 5A

PROTECTIVE ORDER MATERIAL
REDACTED

82. 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 (Veselov) 17:24-29. Veselov does not contemplate or describe installing an agent on the volume image 508 or on a snapshot. Veselov 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 (Veselov) 18:38-44.

PROTECTIVE ORDER MATERIAL
REDACTED

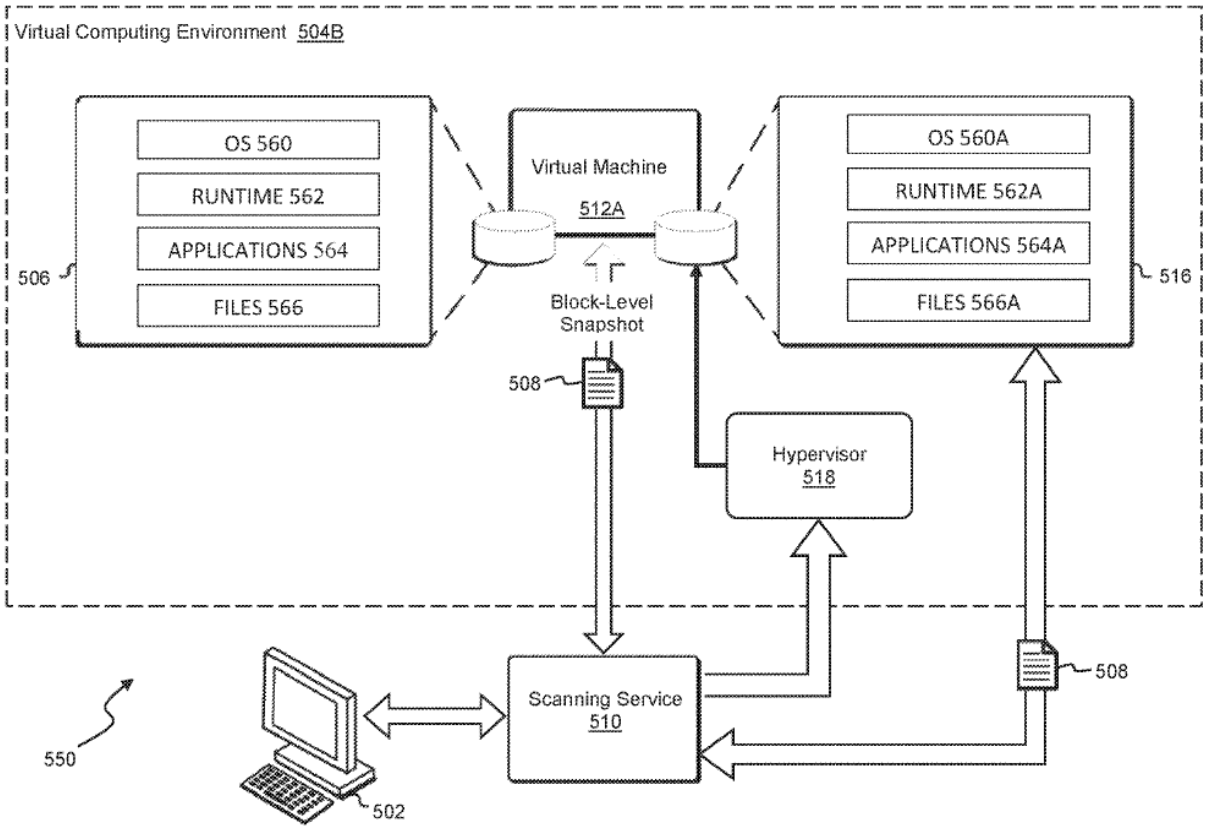

FIG. 5B

Ex.1007 (Veselov) Fig. 5B

83. Veselov makes clear that the implementation of Figures 3A-3B, which analyzes a duplicate virtual machine, and the implementation of Figures 5A-5B, which analyzes a snapshot, are alternative implementations. For example, Veselov states that "**rather than recreating and instantiating a duplicate VM**, the scanning service is configured to analyze the volume image 508 as a data file." Ex.1007 (Veselov) 17:1-3. Thus, Veselov indicates that either a snapshot is analyzed as a data file, or a VM is instantiated from a snapshot.

PROTECTIVE ORDER MATERIAL
REDACTED

84. The difference between these implementations, where Figures 3A-3B instantiate duplicate/assessment virtual machines and install agents on those assessment virtual machines, 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 Basavapatna with a reasonable expectation of success. As detailed below, a person of ordinary skill in the art would not have been motivated to combine the teachings of Basavapatna that Wiz and Dr. Stavrou rely upon, and would not have reasonably expected success in doing so.

b. A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov's Snapshot-Based Analysis with Basavapatna's Agent-Based Analysis as Proposed

85. Considering Dr. Stavrou's reasons to combine Veselov and Basavapatna, persons of ordinary skill in the art still would not find motivation to combine the disclosures. Specifically, a person of ordinary skill in the art considering the disclosure of Veselov would not be motivated to look to the teachings of Basavapatna absent hindsight reasoning guided from the disclosures of the '326 patent.

86. As I discuss above, 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[]

PROTECTIVE ORDER MATERIAL
REDACTED

other useful work." Ex.1007 (Veselov), 12:21-43. Thus, Veselov's security systems act upon a duplicate of a target VM asset instantiated from a snapshot of the target VM asset, or analyze a snapshot, and do not analyze the target VM asset itself. Basavapatna, in contrast, discloses a security process where active target electronic assets are monitored continuously by sensors. Ex.1008 (Basavapatna), ¶¶2, 24-28. Basavapatna explains that it employs "host-based sensors" and "network-based sensors" which can include "agent-based or otherwise software-based sensors that are installed on respective assets" and that "monitor the network traffic passing through the input/output subsystem." Ex.1008 (Basavapatna), ¶¶24-28. A person of ordinary skill in the art would have therefore understood that, although virtual machines are not disclosed in Basavapatna, any virtual machines in a hypothetical combination of Veselov and Basavapatna that were subject to Basavapatna's processes would only be combinable with Veselov's Figures 3A-3B implementations, which do not analyze a snapshot. And, as explained below, because Basavapatna's agent-based solutions are fundamentally different from and inapposite to Veselov's snapshot-based solutions, a person of ordinary skill in the art would have not sought to combine its teachings with the Figures 5A-5B teachings of Veselov, which describe only snapshot-based analyses.

87. Dr. Stavrou asserts that a "POSA would have recognized that Basavapatna's security-assessment techniques would operate similarly in the

PROTECTIVE ORDER MATERIAL
REDACTED

context of Veselov's snapshot-based security assessments of virtual cloud assets taught by Veselov." Ex.1002 (Stavrou Decl.), ¶108. Dr. Stavrou also asserts that "security assessments of physical machines and security assessments of VMs (and snapshots of their virtual disks) typically examined the same types of data (e.g., applications, configurations, files, etc.) and identified the same types of security issues." Ex.1002 (Stavrou Decl.) ¶108. I disagree. Dr. Stavrou's analysis is not grounded in the disclosures of the references he uses for his obviousness combination, and as a result, he oversimplifies and ignores critical aspects of Veselov and Basavapatna. A person of ordinary skill in the art would recognize material differences between Veselov's snapshot-based analysis of VMs and Basavapatna's sensor-based monitoring of active non-virtual assets and would not endeavor to combine the two references.

88. A VM's snapshot is, as the '326 patent states, "a copy of the machine's virtual disk (or disk file) at a given point in time." Ex.1001 ('326 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

PROTECTIVE ORDER MATERIAL
REDACTED

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).

89. But the teachings of Basavapatna that Dr. Stavrou relies upon would not be combinable with this limited scope of data available from Veselov's non-executing snapshots. Dr. Stavrou relies on at least the following teachings from Basavapatna: "determining risks associated with vulnerabilities and assets" and "determining application usage." Ex.1002 (Stavrou Decl.) ¶107. However, to "determin[e] risks associated with vulnerabilities and assets" and "determine application usage," Basavapatna implements host and network sensors to monitor and collect vulnerability data from target assets for detecting and determining vulnerabilities and setting forth remediations.

90. Specifically, to achieve the teachings that Dr. Stavrou relies upon, in Basavapatna's section 2.0, Basavapatna states that its "network monitor 102 receives data from several sources in order to determine a risk metric for a given asset and a given threat and/or vulnerability." Ex.1008 (Basavapatna) ¶33. Basavapatna provides in Figure 2 "an example of the sources of data used by a network monitor 102" which includes "vulnerability detection data 206" and "asset

PROTECTIVE ORDER MATERIAL
REDACTED

configuration data 207." Ex.1008 (Basavapatna) ¶¶33-34. Basavapatna explains that "vulnerability detection data 206 specifies, for each asset and for each threat, whether the asset is vulnerable to the threat, not vulnerable to the threat, or of unknown vulnerability." Ex.1008 (Basavapatna) ¶¶33-34. Further, "configuration data 207 specifies, for each asset, details of the configuration of the asset." Ex.1008 (Basavapatna) ¶¶33-34.

91. In section 2.1.1, Basavapatna explains more detail on Asset Configuration Data, stating that "asset configuration data 207 is received from one or more configuration data source(s) 209." Ex.1008 (Basavapatna) ¶¶35-37. Basavapatna states that in one example "configuration data source(s) 209 are the assets and/or sensors themselves." Ex.1008 (Basavapatna) ¶36.

92. In section 2.1.4, Basavapatna provides more detail on Vulnerability Detection Data. Ex.1008 (Basavapatna) ¶¶57-62. In this section, Basavapatna states that "vulnerability detection data 206 is received from one or more vulnerability data source(s) 212" and that "the vulnerability data source(s) 212" may be "one or more data aggregators." Ex.1008 (Basavapatna) ¶58. Basavapatna further explains that "data aggregators receive vulnerability detection data from individual sensors in the system. The individual sensors can be host-based sensors and/or network-based sensors." Ex.1008 (Basavapatna) ¶58. Basavapatna also explains that "data aggregators can receive vulnerability detection data from the

PROTECTIVE ORDER MATERIAL
REDACTED

assets themselves or from the sensors monitoring the assets." Ex.1008 (Basavapatna) ¶58.

93. In section 3.0, Basavapatna explains that "[r]isk metrics can be generated to illustrate or assess the risk evident in particular assets by virtue of the asset's exposure to a particular vulnerability" and "risk metrics can be calculated using either or both of threat-centric and vulnerability-centric techniques." Ex.1008 (Basavapatna) ¶72. In section 3.1, Basavapatna explains an "example process 300 for generating a threat-centric risk metric for an asset and a threat" which includes a step of "analyzing the vulnerability detection data." Ex.1008 (Basavapatna) ¶¶73-98. In section 3.2, Basavapatna explains an "example process 350 for generating a vulnerability-centric risk metric for an asset and a vulnerability" which includes a step of "analyzing the vulnerability detection data 206." Ex.1008 (Basavapatna) ¶¶99-113. As detailed in the previous paragraph, "vulnerability detection data" is received "from individual sensors in the system" which "can be host-based sensors and/or network-based sensors." Ex.1008 (Basavapatna) ¶58.

94. Thus, it is clear that the teachings that Dr. Stavrou relies upon from Basavapatna, "determining risks associated with vulnerabilities and assets" (Ex.1002 (Stavrou Decl.) ¶107) rely on vulnerability detection data, and that such data is collected by what Basavapatna refers to as "sensors."

PROTECTIVE ORDER MATERIAL
REDACTED

95. Section 1.0 of Basavapatna provides more detail on Basavapatna's sensors. With respect to host-based sensors, Basavapatna explains:

> "host-based sensors 108 can include agent-based or otherwise software-based sensors that are installed on respective assets 104. For example, host-based sensor 108 a is installed on asset 104 a, host-based sensor 108 b is installed on asset 104 c, and host-based sensor 108 c is installed on asset 104 e. The host-based sensors 108 run various analyses on their respective assets 104, for example, to identify vulnerabilities on the assets 104 or to identify viruses or other malware executing on the assets 104. The host-based sensors may also provide one or more passive countermeasures for threats, as described above. Example host-based sensors can include antivirus and other antimalware software."

Ex.1008 (Basavapatna) ¶ 26.

96. With respect to network-based sensors, Basavapatna explains:

> "network-based sensors 110 are hardware devices and/or software in a data communication path between assets 104 protected by the sensor and the network resources that the asset is attempting to access." Basavapatna further explains: "[w]hen an asset 104 tries to send information through the network 106 or receive information over the network 106 through a network-based sensor 110, the sensor analyzes information about the asset 104 and the information being sent or received and determines whether to allow the communication."

Ex.1008 (Basavapatna) ¶27.

PROTECTIVE ORDER MATERIAL
REDACTED

97. Two things are therefore clear from Basavapatna's disclosure: to determine risks associated with vulnerabilities and assets and provide those teachings, Basavapatna: (1) installs agents on target assets (i.e., computing devices); and (2) relies upon execution and activity of its target assets (e.g., sending or receiving information over a network).

98. However, installing agents on target assets is not compatible with Veselov's assessments of snapshots in Figures 5A-5B. Likewise, execution and activity of target assets does not occur in Veselov's assessment of snapshots in Figures 5A-5B. Dr. Stavrou confirmed this is the case, as he stated: "an agent can only be installed on a virtual machine" and "I don't think that an agent can be installed on a snapshot." Ex.2172 (2nd Stavrou Transcript) 53:14-20. Thus, Basavapatna's agents could not be implemented on Veselov's snapshots, and thus Basavapatna's agents could not collect the data needed by its process of generating a risk metric and actually generate those risk metrics.

99. The only example of Veselov that contemplates installing agents is the assessment of a duplicate VM described with reference to Figures 3A-3B, which does not correspond to "analyzing the snapshot" when that phrase is correctly construed and as used in the '326 patent's claims.[1] In the example of

[1] Dr. Stavrou's assertion in his analysis of limitation 1.2.c that "Veselov's snapshot-based security analysis can be performed in multiple ways" (Ex.1002

PROTECTIVE ORDER MATERIAL
REDACTED

Figures 5A-5B of Veselov, installing agents is not contemplated, nor would a person of ordinary skill in the art have expected to be able to install agents. As Veselov explains, in the Figures 5A-5B examples, the scanning service "can read files and other data such as the OS 560A, runtime environment 562A, software applications 564A, and files/file system 566A." Ex.1007 (Veselov) 17:23-29; 18:34-38 ("string or other pattern comparisons to identify files, software packages, application versions, and the like, for software installed on the target virtual machine 312"). Installing applications is not contemplated. Indeed, Veselov expressly cautions against installing applications, as it suggests that the "scanning service 510 may mount volume images as read only to avoid inadvertently modifying the volume corresponding to a volume image." Ex.1007 (Veselov) 18:1-3, *see also* 18:27-29 ("read-only virtual hard disk"). A person of ordinary skill in the art would have recognized that a read only disk cannot be modified, and in the context of Dr. Stavrou's suggested combination, Veselov's read only disk could not be modified to have an agent installed. Thus, as Veselov's snapshot analysis **cannot** install an agent, it would not have been obvious to incorporate

(Stavrou Decl.) ¶154) is accurate within the scope of Veselov itself, but not as applied to the '326 patent's claims, because the first approach he points to, where "the scanning service uses the snapshot to launch an active instance of the VM (which Veselov refers to as an assessment VM) that is essentially a duplicate of the target VM at the time the snapshot was captured" (Ex.1002 (Stavrou Decl.) ¶155), is **not** analyzing a snapshot, consistent with the correct interpretation of this limitation, as detailed in the Claim Construction section.

PROTECTIVE ORDER MATERIAL
REDACTED

Basavapatna's teachings of determining risk, and there would not have been a reasonable expectation of success in incorporating Basavapatna's teachings of determining risk, because those teachings require installation of Basavapatna's host-based sensors 108.

100. Likewise, incorporating Basavapatna's teachings of determining risk would not have been obvious based on Basavapatna's requirement of using network-based sensors. As Basavapatna explains, network-based sensors "analyze[] information about the asset 104" when "an asset 104 tries to send information through the network 106 or receive information over the network 106." Ex.1008 (Basavapatna) ¶28. But as I stated above, a snapshot is not a VM, and it follows that a snapshot does not send information through a network or receive information over a network. The only mentions of sending or receiving information in Veselov are described with reference to the duplicate VM assessment techniques in Figures 3A-3B, where Veselov states 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. But in Figures 5A-5B, there is no "assessment virtual machine 312A-B" because the Figures 5A-5B example

PROTECTIVE ORDER MATERIAL
REDACTED

analyzes a snapshot as a data file, where no network communication occurs, and no communications to other computing resources, or other information available in an agent-based system like Veselov's Figures 3A-3B, and thus, there is no "real-time dynamic or static data" or "computational activities" to apply tasks of the security assessment. The data available to analyze in Veselov's Figures 5A-5B implementation is therefore relatively limited in comparison to the dynamic and wide-ranging data available in Veselov's agent-based duplicate virtual machine implementation in Figures 3A-3B, i.e., the types of data that Basavapatna's techniques rely upon. In particular, Veselov teaches using "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 installing and using a security agent as in the Figure 3A-3B embodiments. Additionally, as Dr. Stavrou

PROTECTIVE ORDER MATERIAL
REDACTED

confirmed, it "would not be possible to send or communicate packets or other messages to a snapshot" because a snapshot "doesn't have active processing." Ex.2172 (2nd Stavrou Tr.) 256:21-257:3. Thus, as Veselov's snapshot analysis cannot send or receive communication over a network, it would not have been obvious to incorporate Basavapatna's teachings of determining risk, which require capturing and analyzing such communication.

101. Given these differences, Dr. Stavrou's statement that "Basavapatna's security-assessment techniques would operate similarly in the context of Veselov's snapshot-based security assessments of virtual cloud assets taught by Veselov" is not correct. Ex.1002 (Stavrou Decl.) ¶108. I also disagree with Dr. Stavrou's opinion that agent-based and snapshot-based analyses "examine[] the same types of data" and "identif[y] the same types of security issues," Ex.1002 (Stavrou Decl.) ¶108, because it is contradicted by Veselov and other evidence. Basavapatna's "security-assessment techniques" do not contemplate static analysis, and Dr. Stavrou does not consider or address how Basavapatna's host-based sensors and network-based sensors could be combined with Veselov and capture sufficient data from Veselov's snapshots analyzed as a data file. Basavaptana's risk generation techniques require collecting data obtained from installing and using of agents, but Dr. Stavrou does not establish how that could or would be combined with Veselov's snapshots, in light of his acknowledgment that an agent could not be

PROTECTIVE ORDER MATERIAL
REDACTED

installed on a snapshot. Basavapatna's techniques also require communication over a network, but Dr. Stavrou confirmed that a snapshot could not communicate packets or messages and does not perform any processing. Thus, I disagree with Dr. Stavrou's position that "a POSA would have been motivated to combine the teachings of Veselov and Basavapatna to more accurately assess risks associated with particular vulnerabilities and assets as a whole." Ex.1002 (Stavrou Decl.), ¶110. For one, nothing in Veselov itself suggests Veselov's system lacks of accuracy. *See, e.g.*, Ex.1007 (Veselov) 4:5-8 (describing assessment results or findings reviewed by an administrator without suggesting results are inadequate). And Dr. Stavrou's assertion ignores that Basavapatna's suggested increased accuracy would require its system with sensors that continuously monitor data such that the system need not rely on point-in-time captures for a security analysis. *See generally* Ex.1008 (Basavapatna); Ex.2043; Ex.2044. Collecting data in Basavapatna is entirely reliant on the sensors I discuss above. The sensors collect information regarding vulnerabilities, and Basavapatna does not contemplate any other method of data collection. Veselov, however, does not mention or suggest sensors for its snapshot analysis. And adding sensors, as taught by Basavapatna, to Veselov's snapshot analysis (even if technically possible) would add complexity to Veselov, yet Veselov acknowledges that added complexity is a "drawback[] of existing security assessment systems." Ex.1007 (Veselov) 2:42-46. Thus, the data

PROTECTIVE ORDER MATERIAL
REDACTED

required to determine vulnerability and asset risks in Basavapatna would not be available in a system that combined Veselov with Basavapatna, as Dr. Stavrou contemplates. And persons of ordinary skill in the art would have therefore understood that Basavapatna's techniques for running security assessments on non-virtual assets are materially different from Veselov's snapshot-based techniques to analyze VMs and would not have combined the references.

102. Similarly, Dr. Stavrou's statement that "security assessments of physical machines and security assessments of VMs (and snapshots of their virtual disks) typically examined the same types of data (e.g., applications, configurations, files, etc.) and identified the same types of security issues" is contradicted by Basavapatna and Veselov themselves. Ex.1002 (Stavrou Decl.) ¶108. Veselov's read-only analysis of a snapshot, including the operating system, runtime environment, software applications, and files/file system, is not the same type of data as the data collected from Basavapatna's host-based sensors like "antivirus and other antimalware software" or network-based sensors like data communicated through a network, both of which require activity by a computing device, and are not static data available through a snapshot. At best, a person of ordinary skill in the art might have, hypothetically, been led to combine Veselov's duplicate agent-based virtual machine implementation, shown in Figures 3A-3B, with Basavapatna's agent-reliant techniques, which are techniques that are analogous to

PROTECTIVE ORDER MATERIAL
REDACTED

each other, but such a combination would not have met the challenged claims, because such a combination would not analyze a snapshot as expressly required by every independent claim. Only with the benefit of hindsight would a person of ordinary skill in the art have selectively combined the incompatible techniques of Veselov's Figures 5A-5B snapshot analysis with Basavapatna's agent-based analysis as would be required to meet the claims (but one of ordinary skill in the art would not have done so, for the reasons I detail here).

103. Accordingly, I disagree that in a combination of Veselov and Basavapatna's techniques, the references would "operate similarly" as Dr. Stavrou asserts. Ex.1002 (Stavrou Decl.), ¶ 108. Veselov's analysis of snapshots as a data file would not perform their same function when Basavapatna's teachings would be combined, as Mohanty's agents could not be installed on Veselov's snapshots. A POSA would also not have expected Basavapatna's agent-based teachings to have the same data necessary if using Veselov's snapshots. Indeed, 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 (as described in Veselov) from "[a]gent assisted access" (as described in Basavapatna). Ex.1038 (Suneja), 134. As Suneja states, snapshotting does not require guest cooperation, whereas agent assistance

PROTECTIVE ORDER MATERIAL
REDACTED

access requires guest cooperation to facilitate runtime state extraction. Ex.1038 (Suneja), 134. Similarly, Jiang provides distinctions between agent-based (e.g., Basavapatna) and "out-of-the-box" (e.g., Veselov) 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 thus makes clear that a person of ordinary skill would not have found it obvious to combine Basavapatna's agent-based analyses when considering Veselov's snapshot analysis techniques. Nor would Basavapatna's techniques have had known functions and predictable results when applied to a snapshot-based analysis.

104. Additionally, as I detail further below in the section detailing objective indicia of non-obviousness, before the '326 patent, those in the industry commented that only agent-based security analysis implementations, such as those described in Basavapatna, would provide sufficient data to enable detailed monitoring and protection, and the metrics generated in Basavapatna; by contrast, snapshot-based approaches were "lightweight" and provided "rudimentary

PROTECTIVE ORDER MATERIAL
REDACTED

security capabilities." Ex.2106, 1-2. That is why Orca's SideScanning product was considered an advance and was praised, as detailed below, because it innovatively allowed deep visibility into cloud workloads using snapshot-based techniques but "without the need to install agents." Ex.2076 at 11. And 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).

105. Dr. Stavrou's statement that "Basavapatna's teachings are a natural fit with Veselov's since Veselov broadly teaches detecting and reporting security risks for multiple assets, while Basavapatna provides additional details regarding determining risk and prioritizing more important vulnerabilities/assets" is likewise unsupported by the references. Ex.1002 (Stavrou Decl.) ¶112. Veselov's teachings and Basavapatna's teachings are not a "natural fit" but are instead much different.

106. Dr. Stavrou relies on a series of tertiary references to assert his rationale for combining Veselov and Basavapatna. *See* Stavrou Decl. (Ex.1002), ¶¶106-113 (citing 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)).

PROTECTIVE ORDER MATERIAL
REDACTED

But I have reviewed the cited portions of these, and none of them appear to explain why persons of ordinary skill in the art would make the purported combination of Veselov and Basavapatna.

107. In order, in paragraph 107, Dr. Stavrou cites Klemperer, Tang, Ammons, Milner, Khalid, Hufsmith, NIST Incident Handling, Wysopal, Giakouminakis, Crabtree, and Price, which I address below in the following paragraphs in that order.

108. First, Dr. Stavrou relies upon Klemperer for "taking a snapshot of a VM." Ex.1002 (Stavrou Decl.) ¶107. In the portion cited by Dr. Stavrou (pages 24-25), Klemperer describes a technique for "[s]napshotting guest memory." Ex.1014 (Klemperer), 24-25. However, this portion of Klemperer does not support combining Basavapatna's agent-based techniques with Veselov's snapshot-based analysis, because a snapshot (even given Klemperer's teachings) could not be used to support installation of Basavapatna's agents.

109. Next, Dr. Stavrou relies upon page 5643 of Tang for "matching vulnerabilities using signature database." Ex.1002 (Stavrou Decl.) ¶107. Tang describes approaches to "antivirus approaches" which is like Basavapatna's agent-based sensors. Ex.1034 (Tang) 5643. Tang mentions "snapshot & rollback" techniques, but Tang does not contend or support that "antivirus approaches" like Basavapatna's are combinable with an analysis of a snapshot, like Veselov's

PROTECTIVE ORDER MATERIAL
REDACTED

snapshot-based analysis in Figures 5A-5B. Rather, Tang mentions snapshots as a contrast to antivirus approaches, and to describe "rollbacks to a snapshot previously saved." Ex.1034 (Tang) 5643.

110. Dr. Stavrou then turns to Ammons for "image/snapshot-based analysis." Ex.1002 (Stavrou Decl.) ¶107. The cited portions of Ammons describe a "checkpoint operation … that takes a 'snapshot' virtual image" to perform "passive monitoring" of a virtual server. Ex.1043 (Ammons) ¶43. However, Ammons' techniques specifically operate "without utilizing agents executing within the virtual server" because "agents use resources that could otherwise be devoted to the designed function of the computer system" and agents cause impact "on the capacity, function and/or communications of the computer system." Ex.1043 (Ammons) Abstract, ¶7. Thus, Ammons would have led a person of ordinary skill in the art away from combining Veselov's snapshot-based assessments with Basavapatna's agent-based techniques, because it specifically discourages agent-based approaches. Ammons does not support combining Veselov and Basavapatna.

111. Dr. Stavrou also turns to Milner for "matching software packages to database of known vulnerabilities." Ex.1002 (Stavrou Decl.) ¶107. Milner uses a "security tool" for its techniques which is "implemented as an application program that is capable of executing on the computing systems." Ex.1076 (Milner) 5:56-

PROTECTIVE ORDER MATERIAL
REDACTED

59. Milner's security tool is therefore like Basavapatna's agent-based analysis techniques. But Milner does not address snapshots, or support combining the software agent-based analysis techniques of Basavapatna with a snapshot-based analysis like the ones described in Veselov's Figures 5A-5B.

112. The techniques of Khalid, which Dr. Stavrou refers to for "comparing characteristics of suspicious objects to CVE database" (Ex.1002 (Stavrou Decl.) ¶107) use a "mobile device[] … deploying an endpoint application." Ex.1077 (Khalid) 17:63-66. Thus, Khalid's techniques are again like Basavapatna's agent-based techniques. But Khalid does not discuss snapshots, and does not support combining the agent-based analysis techniques of Basavapatna with Veselov's snapshot-based analysis in Figures 5A-5B.

113. Dr. Stavrou then turns to Hufsmith for "analysis of container/VM images includes CVE/CWE/malware matching." Ex.1002 (Stavrou Decl.) ¶107. Hufsmith relates to "software development related to distributed applications" and "container images" (Ex.1078 (Hufsmith) ¶2) and does not discuss or relate to snapshots or any snapshot-based analysis of a virtual machine. Hufsmith also does not discuss software agents like those discussed in Basavapatna, or disclose an executing virtual machine which is analyzed by a software agent. Thus, Hufsmith also does not support combining the teachings of Veselov and Basavapatna.

PROTECTIVE ORDER MATERIAL
REDACTED

114. The NIST Incident Handling document referred by Dr. Stavrou as describing "prioritizing incidents based on impact of systems and information" (Ex.1002 (Stavrou Decl.) ¶107) also does not support combining an agent-based analysis like that in Basavapatna with Veselov's snapshot scanning technology. The NIST Incident Handling document discusses obtaining "snapshots through full forensic disk images" for "investigatory and evidentiary purposes" but does not discuss any agent-based analysis, much less combining agent-based analysis with snapshot-based analysis. Ex.1025 (NIST Incident Handling), 44.

115. Dr. Stavrou then turns to Wyspoal for "context-sensitive risk levels." Ex.1002 (Stavrou Decl.) ¶107. Wysopal uses an "analysis engine" that "receives virtual machine files from a user" with a "static analyzer" that "emulates the virtual machine by loading the image into memory and extracting all (or in some cases a predefined subset) of the files within the system." Ex.1036 (Wysopal), 6:40-51. In some respects, Wysopal's "static analyzer" technique is similar to Veselov's snapshot-based analysis that analyzes a snapshot as a file system. But in another, different example, Wysopal discusses a "dynamic analysis" in which an "executing virtual machine can then be analyzed using various network and host vulnerability scanners." Ex.1036 (Wysopal), 7:40-47. This "dynamic analysis" technique is more akin to Basavapatna's agent-based analysis, which is, as Wysopal acknowledges, a different technique than a snapshot-based analysis (as in

PROTECTIVE ORDER MATERIAL
REDACTED

Veselov's Figures 5A-5B). Thus, Wysopal does not support combining the teachings of Veselov and Basavapatna as Dr. Stavrou attempts to; rather, Wysopal supports that there is a distinction between analysis of a snapshot as a data file (the "static analyzer" technique) and analysis of executing a virtual machine using agents (the "dynamic analyzer" technique) and that a person of ordinary skill in the art would not have sought to combine the two approaches.

116. Next, Dr. Stavrou cites Giakouminakis for "prioritizing based on factors such as connection to other assets in the network" and "prioritizing vulnerabilities and using thresholds." Ex.1002 (Stavrou Decl.) ¶107. But Giakouminakis' "security tool" is "configured to run on one or more of the computing systems 104," i.e., it is an agent, like Basavapatna's agent-based analysis. Ex.1044 (Giakouminakis) 3:36-40. Giakouminakis does not discuss snapshots, and similarly does not support combining Veselov's snapshot-based analysis with Basavapatna's teachings.

117. Dr. Stavrou also turns to Crabtree for "prioritizing software vulnerabilities." Ex.1002 (Stavrou Decl.) ¶107. Dr. Stavrou specifically cites 11:52-12:14 of Crabtree, which discusses analyzing "vulnerabilities and exploits found in software executed in the sandbox environment." Ex.1046 (Crabtree) 11:54-60. Thus, Crabtree relies upon a computing device's execution to determine vulnerabilities, like Basavapatna's agent-based techniques, and unlike Veselov's

PROTECTIVE ORDER MATERIAL
REDACTED

snapshot-based analysis, which cannot support execution of software. Thus, Crabtree does not support a motivation to combine Veselov's snapshot-based analysis with Basavapatna's teachings.

118. Finally, in paragraph 107, Dr. Stavrou relies on Price's paragraph 43. Ex.1002 (Stavrou Decl.) ¶107. But Price refers to VM images, not snapshots. *See* Ex.1048 (Price) Abstract. And Price analyzes VM images only for VMs identified as offline, but for VMs identified as online, Price installs an agent on those VMs, just like Basavapatna installs agents on its host devices. Ex.1048 (Price) Abstract. Price therefore does not support combining Basavapatna's agent-based techniques with Veselov's snapshots, and supports that a person having ordinary skill in the art would not have combined those teachings.

119. In paragraph 111, Dr. Stavrou relies upon Canzanese, Crabtree, Hibbert, and Hufsmith. I have addressed Crabtree and Hufsmith above. With respect to Canzanese, Dr. Stavrou cites paragraphs 7-8. Ex.1002 (Stavrou Decl.) ¶111. These paragraphs relate to "log data generated by security and operations systems" and identifying "groups of entities in an enterprise network that are impacted by one or more alerts in the logs of data generated by security systems." Ex.1045 (Canzanese) ¶¶7-8. Canzanese does not mention snapshots, and does not mention agent-based analysis of computing devices, and therefore does not support combining the teachings of Veselov and Basavapatna.

PROTECTIVE ORDER MATERIAL
REDACTED

120. Dr. Stavrou relies on Hibbert for "alerts may 'indicate the degree of danger'." Ex.1002 (Stavrou Decl.) ¶111. But Hibbert, like Basavapatna, is directed to sensor-based collection of vulnerability data, as it teaches "an agent on a digital device" with a "monitor module" and "vulnerability module." Ex.1068 (Hibbert) Abstract. Hibbert does not discuss snapshots, and thus does not support a motivation to combine Veselov and Basavapatna.

121. Finally, in paragraph 112, Dr. Stavrou relies upon NIST Incident Handling, Giakouminakis, and Crabtree. I have addressed these references above.

122. 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 Basavapatna, 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 Basavapatna, or simply further support that there are differences between the applied teachings of Veselov and Basavapatna, such that a person of ordinary skill in the art would have not combined their teachings. Dr. Stavrou does not acknowledge these failings. Thus, Dr. Stavrou has not shown that a person of ordinary skill in the art would have been motivated to combine Veselov and Basavapatna.

PROTECTIVE ORDER MATERIAL
REDACTED

123. 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.), ¶¶112-113) in making Petitioner's and Dr. Stavrou's proposed Veselov-Basavapatna combination is based on mischaracterizations of Veselov and Basavapatna. As detailed above, there are significant gaps between the data available in Veselov's snapshot analysis in Figures 5A-5B and the data needed by Basavapatna to determine vulnerability and asset specific risk metrics. Nothing in Dr. Stavrou's declaration or tertiary references attempts to bridge these gaps or demonstrate how the combination of Veselov and Basavapatna would even theoretically operate or function. And, as Dr. Stavrou conceded, it is not possible to install an agent on a snapshot, and thus not reasonable to expect success in installing Basavapatna's sensors on Veselov's snapshots. Even if such a combination of Basavapatna's techniques with Veselov's snapshot analyses were theoretically possible (and to be clear, I do not believe such a possibility exists), adding Basavapatna's techniques to Veselov would add complexity (e.g., agent installation, additional hardware (network sensors) installation and configuration, etc.), which is something that Veselov seeks to avoid, and thus, a person of ordinary skill in the art would not have added Basavapatna's techniques to Veselov. *See* Ex.1007 (Veselov), 2:42-46 (noting "drawbacks of existing security assessment systems" including "complexity of installation"). This lack of analysis

PROTECTIVE ORDER MATERIAL
REDACTED

demonstrates that a person of ordinary skill would not have had a reasonable expectation of success in combining Veselov's snapshot-based techniques with Basavapatna's sensor-based techniques. Contrary to Dr. Stavrou's conclusion, "the techniques applied in the combination of Veselov and Basavapatna—including … analyzing snapshots to detect vulnerabilities (e.g., based on matching against known lists), determining risks associated with vulnerabilities and assets … [and] prioritizing vulnerabilities/assets" (Ex.1002 (Stavrou Decl.) ¶113) would, in fact, have "presented meaningful technical challenges" and "produced unexpected results." Dr. Stavrou has not explained how to overcome, for example, the meaningful technical challenge of installing Basavapatna's agents on Veselov's read-only snapshots, or the meaningful technical challenge of Basavapatna's determination of risks associated with vulnerabilities and assets without agent-based data based on processing that Basavapatna requires when considering Veselov's snapshot-based system that does not have processing occurring. And a person of ordinary skill would recognize that combining Veselov's snapshot technology with Basavapatna's sensor-based technology would generate unpredictable results and challenges, at least because each reference purports to – and indeed, is designed to – mitigate the need for the technical features of the other.

PROTECTIVE ORDER MATERIAL
REDACTED

124. For the foregoing reasons, it is my opinion that Petitioner 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-Basavapatna 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-Basavapatna combination.

125. For at least the foregoing reasons, as well as the objective indicia I discuss in Section IX below, it is my opinion that Petitioner and Dr. Stavrou have not shown that any of the independent claims of the '326 patent—claims 1, 15, 18—would have been obvious. And it follows that Petitioner and Dr. Stavrou have not shown that any of the claims depending on those independent claims, including but not limited to the dependent claims challenged in Petitioner's Ground 1 (claims 2-20 and 28), would have been obvious. Thus, Petitioner and Dr. Stavrou have not shown that any of claims 1-21 and 28 of the '326 patent would have been obvious.

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

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

PROTECTIVE ORDER MATERIAL
REDACTED

a. **Dependent Claim 6**

127. Dependent claim 6 depends from claim 3, which further depends from independent claim 1 and recites: "determining the risk associated with one of the plurality of cyber vulnerabilities further comprises determining whether the one of the potential cyber vulnerabilities corresponds to an application that is in use by the respective protected virtual cloud asset."

128. Veselov and Basavapatna do not render obvious this additional limitation for the same reasons I discuss above with respect to claim 1, and further because Veselov and Basavapatna fail to disclose or render obvious these claim features.

129. Dr. Stavrou relies only on Basavapatna as allegedly rendering obvious this claim feature, first stating the "determination of the vulnerability-centric risk metric includes determining whether the associated vulnerability corresponds to an application that is in use by the respective asset." Ex.1002 (Stavrou Decl.) ¶205. Dr. Stavrou asserts "Basavapatna then explains that a 'vulnerability is applicable to or possessed by an asset when **the asset is running software** or utilizing hardware that is of a type that is associated with or otherwise evidences the vulnerability. Ex.1002 (Stavrou Decl.) ¶205 (emphasis original) quoting Ex.1008 (Basavapatna) ¶104. Dr. Stavrou also asserts "Basavapatna explains that generating the vulnerability detection data (upon which the vulnerability detection

score is based) involves determining application usage, noting an example in which "identifying that a particular asset has an installed **or running version of a particular software product** can be interpreted to mean that the asset is vulnerable to particular vulnerabilities known to exist on or that are attributable to the particular version of the product."). Ex.1002 (Stavrou Decl.) ¶205 (emphasis original) quoting Ex.1008 (Basavapatna) ¶60. I disagree with Dr. Stavrou.

130. Dr. Stavrou does not acknowledge that determining that an asset is "running software" or "running [a] version of a particular software product" requires that the asset (e.g., a computing device) is executing, which would not be the case when considering Veselov's snapshot-based analysis described in Figures 5A-5B. In the examples of Veselov that analyze a snapshot, the snapshot is analyzed as a "data file" and not as an executing virtual machine, as Veselov describes with reference to other examples (e.g., in Veselov's Figures 3A-3B). As Dr. Stavrou admitted, a snapshot does not have active processing that could be analyzed. Ex.2172 (2nd Stavrou Tr.) 256:21-257:3. Thus, Basavapatna's teachings of running software cannot be reconciled with the fact that Veselov's snapshots are not a running virtual machine that would run software, such that Basavapatna's teachings could apply to determine "whether the one of the potential cyber vulnerabilities corresponds to an application that is in use by the respective protected virtual cloud asset" as recited in claim 6.

PROTECTIVE ORDER MATERIAL
REDACTED

131. Dr. Stavrou then contends "[a]t a minimum, this limitation would have been obvious in view of Basavapatna's teachings since it was well known that whether a vulnerable application is in use affects the associated risk." Ex.1002 (Stavrou Decl.) ¶205. To support his argument, Dr. Stavrou relies upon Hibbert and Hufsmith. Neither reference supports his obviousness argument. Hibbert, as detailed above, relies on agents, similar to Basavapatna, meaning it analyzes an operational computing device, not a snapshot as in Veselov, and thus Hibbert does not support analyzing a snapshot to determine that a vulnerability corresponds to an application that is in use. Hufsmith, as detailed above, focuses on pre-deployment container image analysis, which is a different context than snapshot-based analysis, and thus also cannot support obviousness of this limitation.

132. Next Dr. Stavrou asserts the "determination of the threat-centric risk metric" in Basavapatna also teaches this claim feature. Ex.1002 (Stavrou Decl.) ¶206. But Dr. Stavrou's analysis here relies on "the same reasons" and similar disclosure of an asset "running software" that he discusses in the previous paragraph, and thus his analysis fails for the same reasons.

133. Accordingly, Dr. Stavrou has not shown that dependent claim 6 would have been obvious over the teachings of Basavapatna (whether combined with Veselov or not).

PROTECTIVE ORDER MATERIAL
REDACTED

b. **Dependent Claim 8**

134. Dependent claim 8 depends from claim 7, which depends from claim 3, which in turn depends from independent claim 1, and recites: "wherein determining whether the matching installed applications are used by the respective protected virtual cloud asset includes checking configuration files of the matching installed applications to determine whether at least one of the matching installed applications is not in use, and wherein prioritizing reduces priority of the at least one matching installed application not in use."

135. Veselov and Basvapatna do not render obvious this additional limitation for the same reasons I discuss above with respect to claim 1, and further because Veselov and Basavapatna fail to disclose or render obvious these claim features.

136. Dr. Stavrou first relies on his arguments with respect to claim 4, asserting that "Veselov and Basavapatna each teach identifying vulnerabilities on an asset by matching the asset's installed applications against a list of applications known to be vulnerable." Ex.1002 (Stavrou Decl.) ¶214. But that analysis does not address at least the language of claim 8 that requires determining "whether at least one of the matching applications **is not in use**" or "checking configuration files" to do so and thus his previous analysis does not show obviousness of this claim limitation.

PROTECTIVE ORDER MATERIAL
REDACTED

137. Dr. Stavrou also asserts that "Basavapatna's application-usage determinations apply to these matching installed applications since the application-usage determinations are performed as part of risk assessments for the detected vulnerabilities, and the detected vulnerabilities include the matching installed applications." Ex.1002 (Stavrou Decl.) ¶214. But this fails to analyze the claim language, which requires determining whether a matching application "is not in use" by "checking configuration files" and does not show obviousness of this claim limitation.

138. Next, Dr. Stavrou asserts the combination teaches the claim limitation. He first relies on Veselov as allegedly teaching "the snapshot includes configuration files of installed applications and that the security assessment considers these configuration files." Ex.1002 (Stavrou Decl.) ¶214. The quotations he relies on from Veselov do not specifically mention configuration files, but in any event, simply considering configuration files does not teach using configuration files to determine whether at least one of the matching installed applications is not in use, as required by the claim. Rather, the configuration file would merely indicate whether an application exists, or existed at some time, on the virtual machine corresponding to the snapshot. In fact, an application can be uninstalled, but configuration files may remain, meaning that the configuration file would not tell whether an application was in use or not.

PROTECTIVE ORDER MATERIAL
REDACTED

139. Next, Dr. Stavrou asserts "Basavapatna further teaches that the application-usage determination … involves checking application configuration files" because "vulnerability-centric and threat-centric risk metrics are each determined by comparing vulnerability definition data with 'asset configuration data' to determine whether or not an application is in use." Ex.1002 (Stavrou Decl.) ¶215. But even if Basavapatna's vulnerability definition data included "applicability data such as 'software application versions' and 'application settings'" and the other information that Dr. Stavrou relies upon, none of Basavapatna's cited disclosures teach "checking configuration files of the matching installed applications to determine whether at least one of the matching installed applications is not in use" as required. None of the quotations describe "configuration files" of applications. And, in any event, examination of the configuration of the software does not, itself reveal whether an application is in use or not – it just shows whether the application was at one time installed.

140. Finally, Dr. Stavrou falls back on an obviousness rationale, asserting "[a]t a minimum, given Basavapatna's teachings that asset configuration data includes application configurations and that application usage is determined based on the asset configuration data, it would have been obvious to determine application usage by checking application configuration files." Ex.1002 (Stavrou Decl.) ¶216. Dr. Stavrou relies on Carroll (Ex.1079) to support his theories. The

PROTECTIVE ORDER MATERIAL
REDACTED

cited portions of Carroll discuss the Windows registry, which it characterizes as "several databases that track system and application configuration information, as well as user activity" with specific reference to network connections. Ex.1079 (Carroll), 31. Carroll also describes the "System Resource Usage Monitor (SRUM)" as a "lesser known digital artifact that tracks networks to which a computer is connected." Ex.1079 (Carroll), 32. Carroll explains that "An examination of SRUM data would allow you to determine within 59 minutes which applications were running." Ex.1079 (Carroll), 32. But this description in Carroll does not explain how one would be able to use the SRUM data to determine that an application is "not in use" via checking "checking configuration files." Other citations relied upon by Dr. Stavrou relate to opening Office files (e.g., Microsoft Word documents) and not a determination that an application is not in use as required.

141. Accordingly, Dr. Stavrou has not shown that dependent claim 8 would have been obvious over Veselov and Basavapatna.

c. Dependent Claim 14

142. Dependent claim 14 recites: "wherein determining the location of the snapshot of at least one virtual disk further includes querying a cloud management console of the cloud computing environment for the location of the snapshot and

PROTECTIVE ORDER MATERIAL
REDACTED

the location of a specific virtual disk of the respective protected virtual cloud asset."

143. 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.) ¶¶251-252. 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.

PROTECTIVE ORDER MATERIAL
REDACTED

144. 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.) ¶255. For example, Dr. Stavrou asserts that the snapshot of a virtual disk and the underlying virtual disk can have the "same general location (e.g., being located in the same computing environment)." Ex.1002 (Stavrou Decl.) ¶255. 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-Basavapatna combination would *benefit* from such information.

PROTECTIVE ORDER MATERIAL
REDACTED

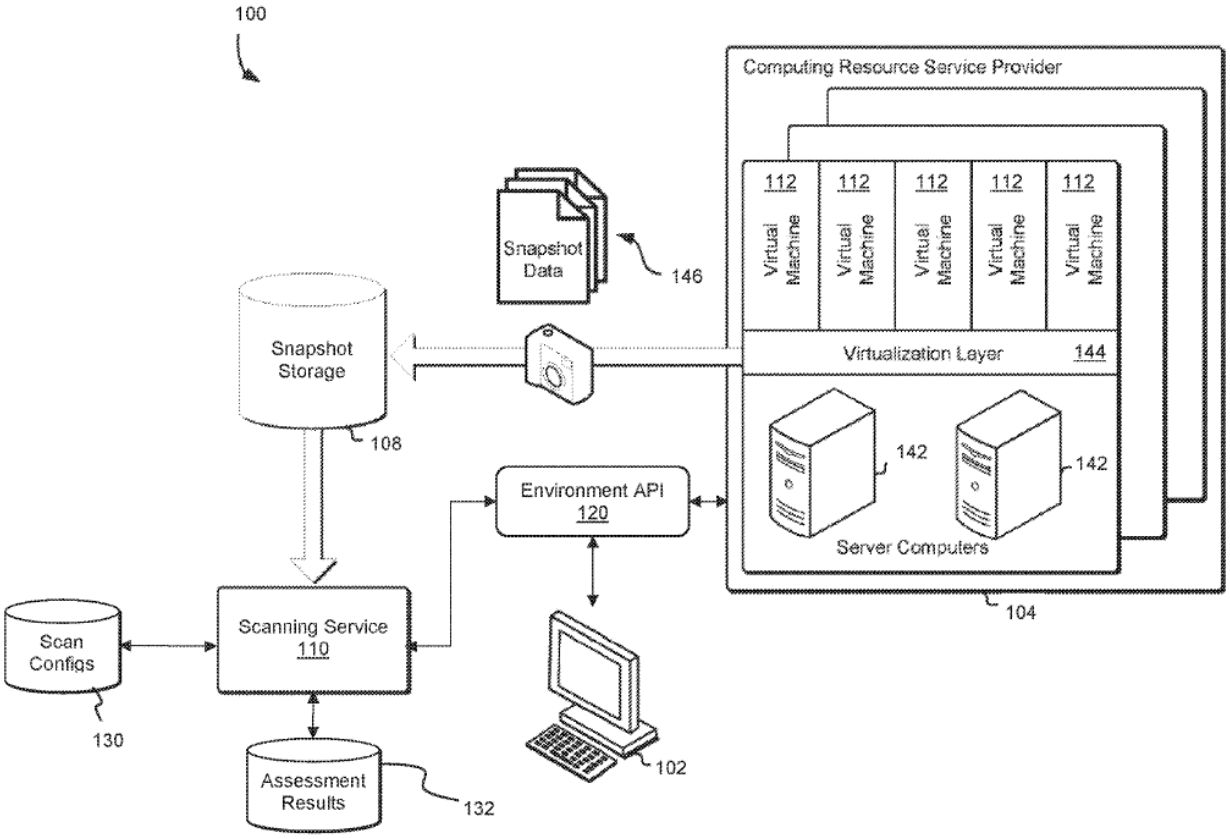

FIG. 1

Ex.1007 (Veselov) Fig. 1

145. 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 determined solely 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

PROTECTIVE ORDER MATERIAL
REDACTED

determination of claim 14. 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.) ¶257.

146. As the '326 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 ('326 patent), Fig. 1B, 3:55-60 ("[S]torage 117 emulat[ing] virtual discs for the VMs" "may be part of the server" hosting VMs, or, as depicted in Figure 1B, separate from but "connected to" that server). And, as Dr. Stavrou admitted at his deposition, "Virtual disks are different from virtual machines"—"[t]hey are definitely not synonymous." Ex.2077, (1st Stavrou Tr.) 97:13-98:10. Thus, the identification of a VM does not necessarily or obviously determine the location of a virtual disk of that VM. Further, identifying a VM does not require or render obvious using an API to "query" such location at least because an identification means the user does not need to query for a location.

147. 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

PROTECTIVE ORDER MATERIAL
REDACTED

apply to Veselov. Ex.1002 (Stavrou Decl.) ¶257. 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.

148. Additionally, in his discussion of claim 14, 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.), ¶¶247-248. 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

PROTECTIVE ORDER MATERIAL
REDACTED

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 '326 patent's claims.[2]

[2] 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. ¶222, 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

PROTECTIVE ORDER MATERIAL
REDACTED

149. Accordingly, Dr. Stavrou has not shown that dependent claims 14 would have been obvious over Veselov and Basavapatna.

B. Ground 2: Claims 4-5 and 17 Would Not Have Been Obvious Over Veselov, Basavapatna, and Czarny

150. Dr. Stavrou sets forth Ground 2 to build upon Ground 1 and Czarny is applied only for "Czarny's well-known methods of using binary and hash comparison for purposes of matching installed application against a list of known vulnerable applications to detect vulnerabilities." Ex.1002 (Stavrou Decl.), ¶¶269-285. But Dr. Stavrou challenges claims 1 and 15, from which claims 4-5 and 17 depend, based on only Veselov and Basavapatna. Therefore, for at least the same reasons Dr. Stavrou fails to show obviousness of the claims based on Veselov and Basavapatna, it is my opinion that Dr. Stavrou fails to show obviousness of these claims based on Veselov, Basavapatna, and Czarny. For example, because Dr.

---

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 ¶222 (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.

Stavrou's asserted reasons to combine Veselov and Basavapatna do not consider the context of these references and their significant differences, a person of ordinary skill in the art would not have combined Veselov and Basavapatna, and further would not have added Czarny either.

151. Additionally, even assuming that a person of ordinary skill in the art would have combined Veselov and Basavapatna as Dr. Stavrou contends, a person of ordinary skill in the art would not have been motivated to modify Dr. Stavrou's Veselov-Basavapatna combination with teachings from Czarny to achieve what is claimed. Czarny describes a "computer security vulnerability assessment" which "determines a correspondence between [] product binary data and [] product vulnerability data." Ex.1084 (Czarny) 1:39-56. However, Czarny does not disclose or mention virtual machines (Ex.2172 (2nd Stavrou Tr.) 246:6-21) or snapshots, and given that Czarny does not contemplate analyzing snapshots, Dr. Stavrou has not shown how the proposed combination of Veselov, Basavapatna, and Czarny would work. Accordingly, for at least these reasons, it is my opinion that claims 4-5 and 17 would not have been obvious over Veselov, Basavapatna, and Czarny.

C. Ground 3: Claims 22-27 Would Not Have Been Obvious Over Veselov, Basavapatna, and Giakouminakis

152. Dr. Stavrou sets forth Ground 3 to build upon Ground 1 and further relies on the teachings of Giakouminakis. Giakouminakis is applied only for its

PROTECTIVE ORDER MATERIAL
REDACTED

disclosure of "determining risk levels for a particular asset based on analysis of connected assets and additional bases for assessing takeover risk" along with additional disclosures from Basavapatna regarding "how risk levels are determined." Ex.1002 (Stavrou Decl.), ¶¶286-326. But Dr. Stavrou challenges claim 1, from which claims 22-27 depend, based on only Veselov and Basavapatna. Therefore, for at least the same reasons Dr. Stavrou fails to show obviousness of the claims over Veselov and Basavapatna, it is my opinion that Dr. Stavrou fails to show obviousness of claims 22-27 based on Veselov, Basavapatna, and Giakouminakis.

153. Additionally, even assuming that a person of ordinary skill in the art would have combined Veselov and Basavapatna as Dr. Stavrou contends, a person of ordinary skill in the art would not have been motivated to modify Wiz's Veselov-Basavapatna combination with teachings from Giakouminakis to achieve what is claimed.

154. As summarized above, Giakouminakis describes a "security tool" to "identify vulnerabilities in a computing system and determine a risk level of the vulnerabilities." Ex.1044 (Giakouminakis) Abstract. Specifically, similar to Basavapatna, and as detailed above, Giakouminakis describes a "security tool" which "can identify vulnerabilities in a computing system and determine a risk level of the vulnerabilities based on base and optional CVSS vectors and additional

PROTECTIVE ORDER MATERIAL
REDACTED

factors that represent the evolving nature of vulnerabilities." Ex.1044 (Giakouminakis) Abstract. Giakouminakis states that the "security tool 102 can be utilized to test and analyze the security of the computing systems 104 and/or the VMs 114" and "can be configured to run on one or more of the computing systems 104." Ex.1044 (Giakouminakis) 3:36-40. In other words, Giakouminakis utilizes agents on executing computing devices to identify vulnerabilities. Ex.1044 (Giakouminakis) 12:31-36. Giakouminakis does not contemplate snapshots or snapshot-based analyses.

155. Accordingly, a person of ordinary skill in the art would not have combined Giakouminakis' teachings with Veselov and Basavapatna, at least for the same reasons that Basavapatna would not have been combined with Veselov (i.e., because Giakouminakis' teachings require an executing virtual machine, Veselov's snapshot-based analysis in Figures 5A-5B is not an executing virtual machine, and adding Giakouminakis to Veselov would add complexity contrary to Veselov's goals). Thus, just like Dr. Stavrou has failed to explain how the Veselov-Basavapatna combination would work, with a reasonable expectation of success, Dr. Stavrou has failed to show how its purported Veselov-Basavapatna-Giakouminakis combination would work, and that a person of ordinary skill in the art would have had a reasonable expectation of success in combining the teachings. This is particularly true where Veselov's snapshot-based analyses do not provide

PROTECTIVE ORDER MATERIAL
REDACTED

the types of data required by the agent-based analyses of Basavapatna and Giakouminakis, contrary to Dr. Stavrou's opinion that the two different types of analyses involve "comparable data." Ex.1002 (Stavrou Decl.) ¶291.

IX. Objective Indicia of Non-obviousness

156. I have considered objective evidence that supports that a person of ordinary skill in the art would not have considered the claims of the '326 patent to be obvious, as discussed below. It is my opinion that, even assuming that Petitioner and Dr. Stavrou made a threshold showing of obviousness (they have not for the reasons explained above), there is objective evidence of non-obviousness that shows that the claimed combination of features would not have been obvious to a person of ordinary skill in the art.

A. Orca's SideScanning™ Uses The '326 Patent Inventions

157. 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 '326 patent. In that section of his declaration, Mr. Shua maps Orca's SideScanning™ to the claims of the '326 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

PROTECTIVE ORDER MATERIAL
REDACTED

Mr. Shua's mapping, it is my opinion that Orca's SideScanning™ practices the claims of the '326 patent.

B. Long-Felt, Unresolved Need

158. 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 '326 patent describes, "[o]rganizations [] increasingly adapted their applications to be run from multiple cloud computing platforms." Ex.1001 ('326 patent), 1:21-25. Virtualization technologies such as virtual machines, software containers, and serverless functions helped enable this shift to commercial cloud computing. *See* Ex.1001 ('326 patent), 1:39-42.

159. 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

PROTECTIVE ORDER MATERIAL
REDACTED

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 '326 patent describes, the cloud security industry developed a number of methods for protecting cloud computing infrastructures, including: (1) "traffic inspection", (2) using "network scanner[s] deployed out of path," (3) "traffic monitor[ing]," and (4) using "agents installed in each server," but these methods suffered from multiple disadvantages. Ex.1001 ('326 patent), 1:49-2:17.

160. 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. Ex.1001 ('326 patent), 1:52-2:9. The problem of limited coverage, which was especially true with agent-based systems, includes difficulties associated with installing agents on every VM. Thus, while an agent may have better visibility into a VM's security posture, the difficulty in installing agents throughout a platform left big gaps in coverage. See Ex.2076, 11 ("Traditional agents struggle with complex deployment and maintenance."); Ex.1001 ('326 patent), 2:16-17 ("[I]nstalling agents in a large data center may take

PROTECTIVE ORDER MATERIAL
REDACTED

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.").

161. 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.").

162. Accordingly, there was a long-felt, unmet need for a cloud security solution that could provide both visibility and coverage to adequately address "undetected vulnerabilities." Ex.1001 ('326 patent), 1:52-55, 1:60-66, 2:4-9, 2:14-17.

163. 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

PROTECTIVE ORDER MATERIAL
REDACTED

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.

164. "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.[3]

Ex.2076 at 31.

165. REDACTED

[3] *See* '326 patent, cls. 1, 15, 18 (limitations 1.2.c, 15.2.c, 18.2.c.

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

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

166. 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.

167. 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 (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.

168. 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

[4] *See* '326 patent, cls. 1, 15, 18 (limitations 1.4, 15.4, 18.4).

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").

169. REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

171. 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

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

- CRN's "20 Coolest Tech Startups of 2020." Ex.2119, *The 20 Coolest Tech Startups of 2020*, CRN (Nov. 19, 2020) at 9-10. CRN touted "Orca's SideScanning technology" which "automatically assesses the security state of every discovered asset" to provide "complete visibility." 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) at 6-7. Forbes found Orca "noteworthy for its innovative approach to providing instant-on, workload-deep security" for cloud computing systems, emphasizing Orca's "patent-pending SideScanning™ 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")

173. REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

175. 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. 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

- Exhibit 2118, which is titled "Vendor Landscape and TCO Analysis for Achieving Cloud-Wide, Workload-Deep

> Security and Compliance for Public Cloud Environments" and which likewise includes praise of Orca's invention. Ex.2118 (Vendor Landscape), 1. This document describes "Orca uses a patent-pending technique called SideScanning™, … which unobtrusively scans the run-time block storage layer—out of band—to extract valuable data about the environment and the assets communicating in it." Ex.2118 (Vendor Landscape), 5. The report states that a Senior Cloud Security Engineer from a company called Sisense described Orca's solution as "unique in that it locates vulnerabilities with precision and delivers tangible, actionable results—without having to sift through all the noise." Ex.2118 (Vendor Landscape), 4.

176. In summary, Orca's extensive praise for the claimed combination of features in the '326 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™

177. 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 '326 patent's claims obvious.

178. REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

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

PROTECTIVE ORDER MATERIAL
REDACTED

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

182. 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.

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

201. 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™

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

REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

205. 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

PROTECTIVE ORDER MATERIAL
REDACTED

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 determined 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.

PROTECTIVE ORDER MATERIAL
REDACTED

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

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

214. In the discussion that follows, I discuss additional exhibits I have reviewed that further support the substantial similarity between Wiz and Orca's products as it relates to specific limitations of the claims of the '326 patent. REDACTEI

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

a. **Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.1, 1.2, 1.2.a, 1.2.b, 15.1, 15.2, 15.2.a, 15.2.b, 18.1, 18.2, 18.2.a, 18.2.b.**

215. Limitations 1.1, 1.2, and 1.2.a recite "*receiving a request to scan a plurality of protected virtual cloud assets in the cloud computing environment*," *"for each of the requested plurality of protected virtual cloud assets in the cloud*

computing environment," "*determining, using an API or service provided by the cloud computing environment, a location of a snapshot of at least one virtual disk of a respective protected virtual cloud asset*" and "*accessing, based on the determined location and 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 SideScanning™ in relation to limitations 1.1, 1.2, 1.2.a, 1.2.b, 15.1, 15.2, 15.2.a, 15.2.b, 18.1, 18.2, 18.2.a, and 18.2.b of claims 1, 15, and 18 of the '326 Patent.

216. For example, similar to Orca's SideScanning™, Wiz's Workload Scanner is meant to "cover everything" and an "entire environment" (i.e., multiple assets within a cloud computing environment). Ex.2115 at 2, 3.

217. Further, 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.

PROTECTIVE ORDER MATERIAL
REDACTED

> 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.

218. 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.

219. REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

220. 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

221. REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

PROTECTIVE ORDER MATERIAL
REDACTED

REDACTED

b. **Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.2.c-1.2.d, 15.2.c-15.2.d, 18.2.c-18.2.d.**

222. Limitations 1.2.c-1.4 recite "*analyzing the snapshot of the at least one virtual disk to determine an existence of potential cyber vulnerabilities*" and "*determining a risk associated with each of the determined potential cyber vulnerabilities*." Wiz's Workload Scanner is substantially similar to Orca's SideScanning™ in relation to limitations 1.2.c-1.2.d, 15.2.c-15.2.d, and 18.2.c-18.2.d of claims 1, 15, and 18 of the '326 Patent (reproduced in full above).

PROTECTIVE ORDER MATERIAL
REDACTED

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

Stage 6: Vulnerability analysis

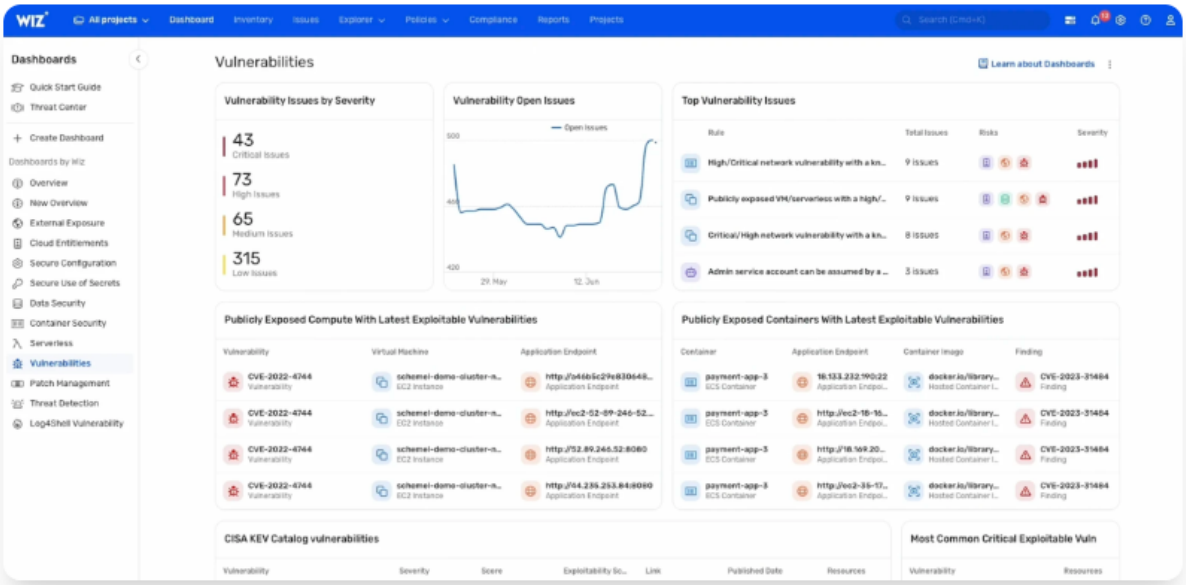

Example vulnerability dashboard that prioritizes issues by contextual severity

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

Ex.2165 at 8

224. Additionally, Wiz detects and prioritizes a catalog of "known exploited vulnerabilities." *See* Ex.2147.

PROTECTIVE ORDER MATERIAL
REDACTED

What is the CISA KEV Catalog?

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

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

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

Ex.2147 at 1-2

c. **Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.3-1.4, 15.3-15.4, 18.3-18.4**

225. Limitations 1.3-1.4 recite "*for each of the requested plurality of protected virtual cloud assets with the determined potential cyber vulnerabilities, determining a risk level to the cloud computing environment*," and "*reporting, for each of the requested plurality of protected virtual cloud assets with the determined potential cyber vulnerabilities, the existence of the potential cyber vulnerabilities, such that the plurality of protected virtual cloud assets with the determined*

PROTECTIVE ORDER MATERIAL
REDACTED

potential cyber vulnerabilities are prioritized based on associated risk levels." Wiz's Workload Scanner uses the features recited in limitations 1.3-1.4, 15.3-15.4, and 18.3-18.4 of claims 1, 15, and 18 of the '326 Patent (reproduced in full above).

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

3. Risk-based prioritization

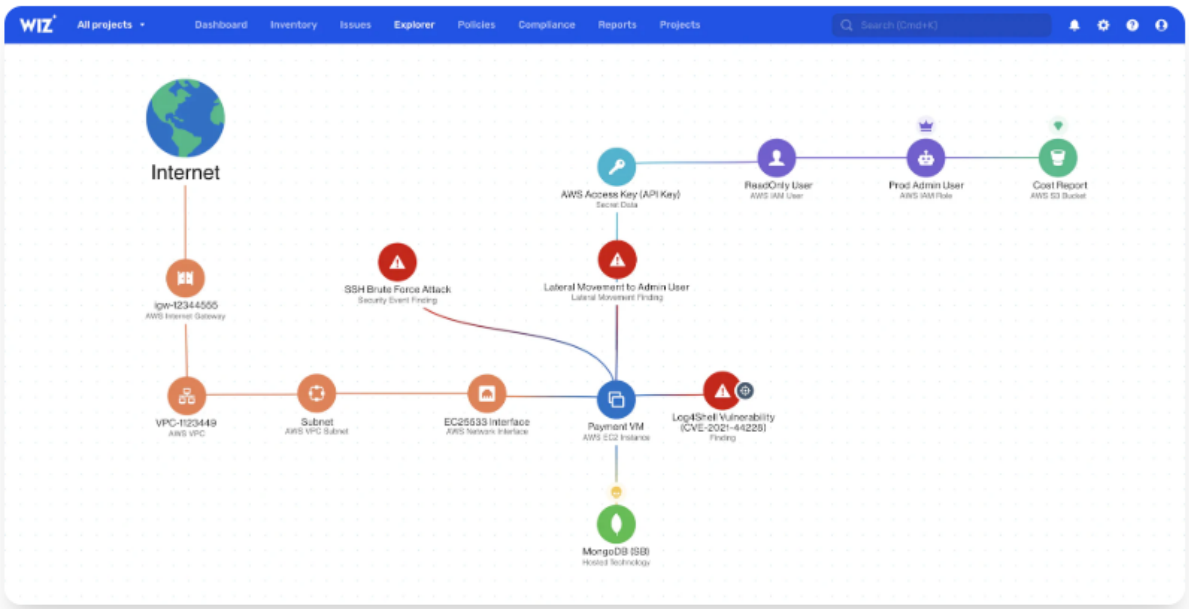

Example visualization of the risk-based context that a vulnerability scanner should provide

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

Ex.2165 at 14-15

227. Wiz's Workload Scanner reports vulnerabilities and orders them based on severity.

Stage 6: Vulnerability analysis

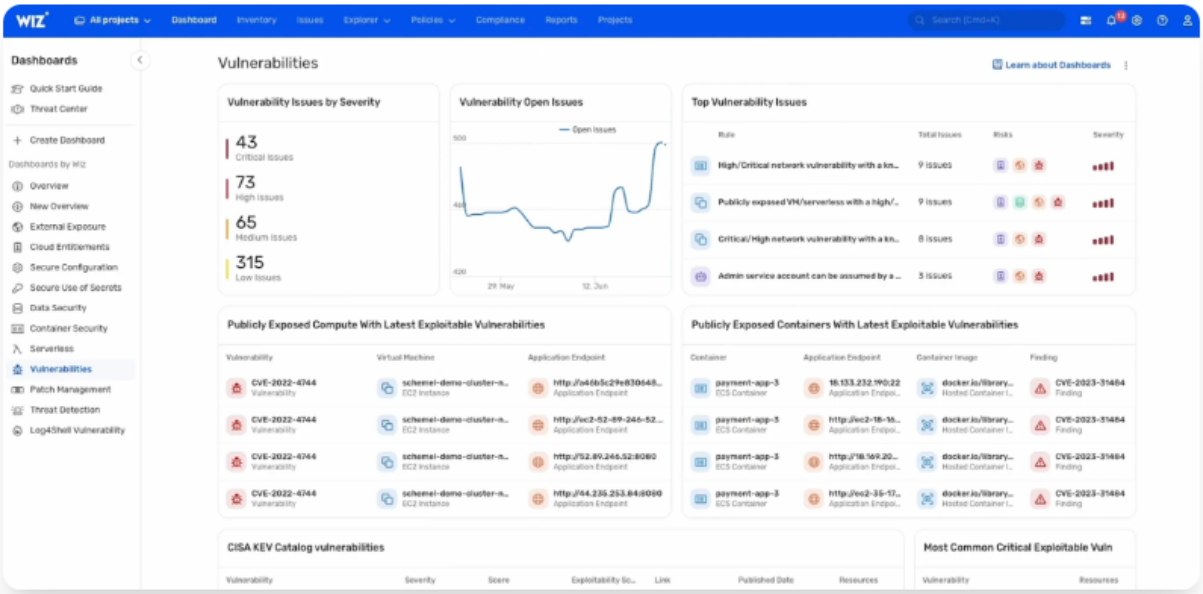

Example vulnerability dashboard that prioritizes issues by contextual severity

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

Ex.2165 at 8

228. Wiz also states that it "analyzes your entire cloud environment to identify which resources are vulnerable and then ranks their severity and exploitability":

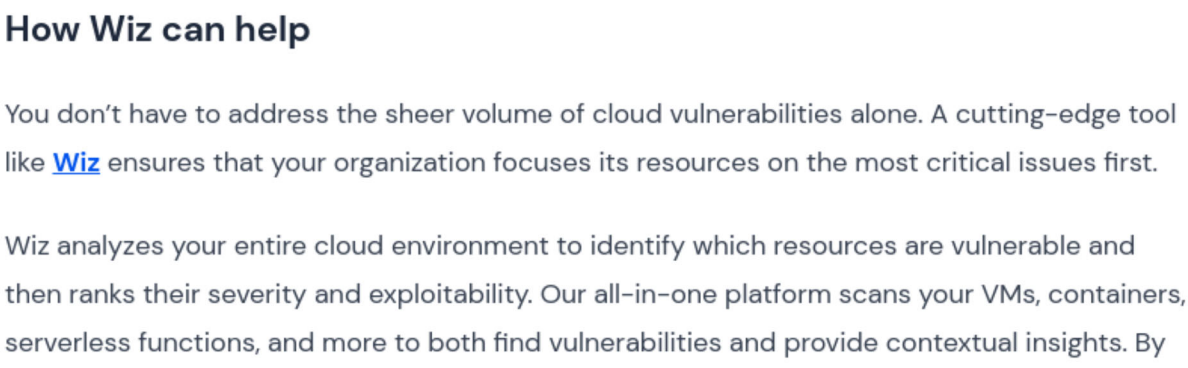

How Wiz can help

You don't have to address the sheer volume of cloud vulnerabilities alone. A cutting-edge tool like Wiz ensures that your organization focuses its resources on the most critical issues first.

Wiz analyzes your entire cloud environment to identify which resources are vulnerable and then ranks their severity and exploitability. Our all-in-one platform scans your VMs, containers, serverless functions, and more to both find vulnerabilities and provide contextual insights. By

Ex.2170 at 11

229. Wiz's Workload Scanner also notes "Top Vulnerability Issues" and provides a "breakdown of each vulnerability's risk impact":

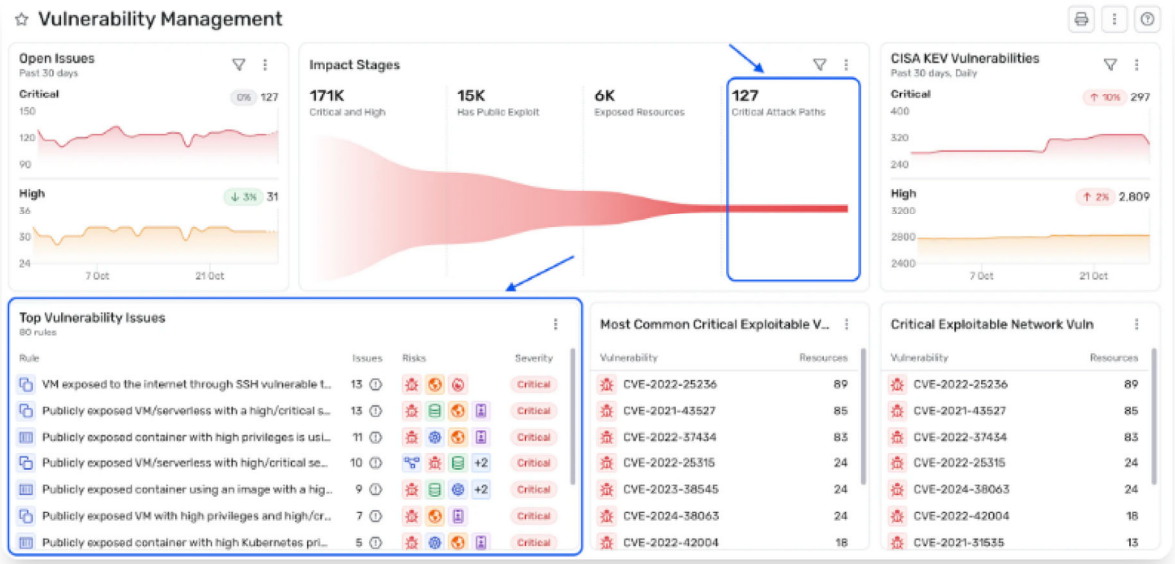

Reduce vulnerability impact

Next up, we want to address vulnerabilities based on vulnerability-specific factors and impact stages. For example, we want to remediate the vulnerabilities based on their severity given by the vendor of the vulnerable component, starting with the critical ones. We also want to address them based on their likelihood of exploitability and focus on the vulnerabilities that have public exploits, as well as vulnerabilities that are on exposed resources.

How we do it – Impact Stages

In Wiz, you can easily find the breakdown of each vulnerability's risk impact in the vulnerability management dashboard, and you can also track progress using the funnel to show you where you are in terms of each impact stage. With this capability, Wiz allows security practitioners to rapidly reduce high-risk and critical vulnerabilities, remove the ones with public exploit, and address vulnerabilities that are exposed to the internet.

Ex.2169

230. Wiz also states that it "correlates vulnerabilities to risks":

How we do it – Wiz Issues

Wiz Issues take contextual analysis to the next level, by identifying critical attack paths and toxic combinations within your environment to provide accurate context-based risk prioritization. Wiz Issues correlates vulnerabilities to risks across misconfigurations, public exposure, identities, secrets, malware, and data so they can focus on the vulnerabilities that truly have an impact to the business.

Ex.2169 at 3

231. Wiz's Workload Scanner also describes "Prioritize vulnerabilities" to "surface the vulnerabilities that should be prioritized."

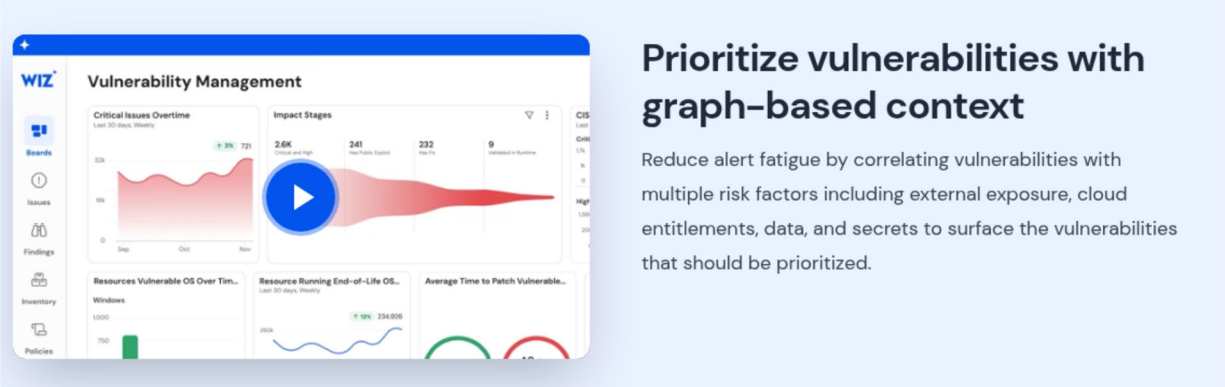

Ex.2168 at 3

232. Wiz also states that prioritization "Filters out irrelevant vulnerabilities and reports on those with the biggest blast radius":

Automated prioritization	• Filters out irrelevant vulnerabilities and reports on those with the biggest blast radius • Helps reduce alert fatigue

Ex.2170 at 9

F. Commercial Success

233. 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

234. 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).

235. It is my opinion that the commercial success of both Orca's and Wiz's products has nexus to the claims of the '326 patent. For example, both systems practice the claims of the '326 patent as discussed above. REDACTED

X. Conclusion

236. For the foregoing reasons, it is my opinion that claims 1-28 of the '326 patent would not have been obvious over the prior art in Petitioner'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.

Dr. David Kaeli

PROTECTIVE ORDER MATERIAL
REDACTED