

Filed on behalf of: Orca Security Ltd.

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

WIZ, INC.,
Petitioner,

v.

ORCA SECURITY LTD.,
Patent Owner.

Case IPR2024-00865
Patent 11,693,685

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

Table of Contents

I.	Introduction.....	3
II.	Qualifications.....	4
III.	Summary Of Materials Reviewed And Considered	7
IV.	Legal Standards.....	8
V.	Level Of Skill In The Art And Perspective Applied In This Declaration.....	12
VI.	Background.....	15
	A. Overview of the '685 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 Hufsmith	31
	C. Overview of Price.....	42
VIII.	The Claims Of The '685 Patent Would Not Have Been Obvious ...	47
	A. Ground 1: Claims 1-22 Would Not Have Been Obvious Over Veselov and Hufsmith	47
	1. A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov and Hufsmith and Would Not Have Had a Reasonable Expectation of Success in Doing So.....	47
	2. Even If Combined, Veselov and Hufsmith Do Not Disclose or Suggest Multiple Limitations in the Independent Claims.....	73
	B. Ground 2: Claims 1-22 Would Not Have Been Obvious Over Veselov, Hufsmith, and Chari	103
	C. Ground 3: Claims 2-3 and 14-15 Would Not Have Been Obvious Over Veselov, Hufsmith, Chari, Price, and Huseinović	107

1.	A Person of Ordinary Skill in the Art Would Not Have Been Motivated to Combine Veselov, Hufsmith, Chari, Price, and and Huseinović , and Would Not Have Had a Reasonable Expectation of Success in Doing So	108
2.	Even If Combined, Veselov, Hufsmith, Chari, Price, and Huseinović Do Not Disclose or Suggest Multiple Limitations in Claims 2-3 and 14-15.....	109
IX.	Objective Indicia of Non-obviousness.....	111
A.	Orca's SideScanning™ Uses The '685 Patent Inventions ..	111
B.	Long-Felt, Unresolved Need	111
C.	Orca's claimed inventions overcame industry skepticism to become widely accepted--and widely praised—in the industry.	117
D.	Extensive Industry Praise Supports Non-Obviousness of the Claims.....	122
E.	Evidence Shows Wiz Copied Orca's SideScanning™	129
1.	Evidence shows Wiz had access to Orca's invention	131
2.	Evidence shows Wiz's Workload Scanning is substantially similar to Orca's SideScanning™	150
F.	Commercial Success.....	178
X.	CONCLUSION.....	181

I. Introduction

1. I, David Kaeli, submit this declaration on behalf of Orca Security Ltd. (“Patent Owner”) to provide my opinions concerning the validity of claims 1-22 of U.S. Patent No. 11,693,685 (the “’685 patent,” Ex.1001), in support of Patent Owner’s Response to the Petition for *Inter Partes* review of the ’685 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 December 9, 2024, this *Inter Partes* review (IPR) was instituted (Paper 8).

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

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

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

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

II. Qualifications

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

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

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

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

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

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

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

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

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

III. Summary Of Materials Reviewed And Considered

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

15. My opinions are additionally guided by my appreciation of how a person of ordinary skill in the art would have understood the claims of the '685 patent at the time of the invention, which I have been asked to assume is January 28, 2019. I will sometimes refer to this time as the "relevant time."

IV. Legal Standards

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

29. With respect to the '685 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 '685 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 '685 patent.

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

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

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

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

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

Ex.1002 (Stavrou Decl.), ¶22. At the time of the earliest priority date for the ’685 patent (January 28, 2019), I met Dr. Stavrou’s proposed qualifications to be a person of ordinary skill in the art, and my opinions in this Declaration would not change if the Board were to adopt and apply that proposed level of ordinary skill.

VI. Background

A. Overview of the '685 Patent

35. The '685 patent is titled "Virtual Machine Vulnerabilities and Sensitive Data Analysis" and generally relates to innovative systems and methods for securing virtual cloud assets, such as VMs. Ex.1001 ('685 patent), Abstract, cls. 1, 13, 22.

36. The '685 patent explains that, around the time of its priority date (January 2019), public cloud service providers, like 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 ('685 patent), 1:20-29. The '685 patent also provides context for "virtualization" and explains that "virtualization can be achieved by means of virtual machines," each of which can "emulate[] a number of 'computers' or instances, all within a single physical device." Ex.1001 ('685 patent), 1:27-34.

37. The '685 patent discusses problems that existed with trends in the cybersecurity field. For example, the '685 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 ('685 patent), 1:39-43. And, the '685 patent also

explains ways the industry had tried to deal with such issues. For example, the '685 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 ('685 patent), 1:44-2:14. As the '685 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 ('685 patent), 1:49-52. 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 ('685 patent), 1:53-64. 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 ('685 patent), 2:1-6. And "[u]sing agents is a cumbersome solution for a number of reasons, including IT resource management, governance, and performance." Ex.1001 ('685 patent), 2:11-13. "For example, installing agents in

a large data center may take months.” Ex.1001 (’685 patent), 2:13-14. The ’685 patent’s systems and method “overcome the[se] deficiencies.” Ex.1001 (’685 patent), 2:15-16.

38. The ’685 patent aims to solve such problems. Ex.1001 (’685 patent), 3:11-12, 3:32-35. Figure 1B of the ’685 patent, reproduced below, illustrates a diagram of a cloud computing platform to which the ’685 patent’s techniques for “scan[ning] for embedded vulnerabilities in machines.




FIG. 1B

39. The cloud computing platform of the ’685 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 (’685 patent), 3:32-67, Fig. 1B. 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 ('685 patent), Fig. 1B, 3:52-4:21. The storage may be part of the server or, as shown in Figure 1B, separated from the server. Ex.1001 ('685 patent), Fig. 1B, 3:52-60. In the example shown in Figure 1B, a "high-speed connection" such as an "optic fiber" connects the storage with the server. Ex.1001 ('685 patent), Fig. 1B, 3:52-60. 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 ('685 patent), Fig. 1B, 4:16-21. The security system is configured to perform the steps of the '685 patent's invention and may be separate from the client environment, as shown in Figure 1B, or part of the client environment. Ex.1001 ('685 patent), Fig. 1B, 4:1-15. The challenged claims of the '685 patent demonstrate the embodiment where the security system interfaces with the client environment through APIs. Ex.1001 ('685 patent), cls. 1, 13, 22.

40. The security system of the '685 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 ('685 patent), 4:43-54. For example, the security system can detect vulnerable code (Ex.1001

('685 patent), 5:30-32), cyber threats that do not represent vulnerabilities such as unencrypted sensitive data (Ex.1001 ('685 patent), 6:19-30), and changes in sensitive machine areas such as changed application files without installation (Ex.1001 ('685 patent), 6:31-37). 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 ('685 patent), 4:43-54), generate reports on detected vulnerabilities and cyber threats and sensitive data, including, “for example, as alerts” (Ex.1001 ('685 patent), 6:38-53, 6:19-53, 7:50-64), and mitigate detected vulnerabilities and cyber threats. Ex.1001 ('685 patent), 7:65-8:3. Mitigation actions may include, for example, blocking traffic from untrusted networks, quarantining an infected VM, or halting the operation of the VM. Ex.1001 ('685 patent), 8:1-3.

41. In order to accomplish the scanning and detection operations I discuss above, the '685 patent provides a series of steps. First, the security system queries for and receive an address of a virtual disk to be assessed (Ex.1001 ('685 patent), 4:22-31). For example, the security system queries a cloud management console for the address of virtual disks that serve the VMs. Ex.1001 ('685 patent), 4:55-67. After the security system receives the location of the virtual disk, the system may generate a snapshot. Ex.1001 ('685 patent), 5:1-11, cls. 1, 13, 22. 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 ('685 patent), 4:58-59. The security system then analyzes the snapshot to detect vulnerabilities and cyber threats and sensitive data, without installing agents on the server or the VM. Ex.1001 ('685 patent), 6:38-53, 6:19-53, 7:50-64. The snapshot analysis does not involve interaction with or information from the VM and does not require agents to be installed in the system server or VM. Ex.1001 ('685 patent), 6:38-53, 6:19-53, 7:50-64. The '685 patent provides multiple techniques for analyzing snapshots to detect different types of vulnerabilities and cyber threats. Ex.1001 ('685 patent), 6:38-53, 6:19-53, 7:50-64. 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 ('685 patent), 6:38-53, 6:19-53, 7:50-64. The security system then generates reports on detected vulnerabilities and cyber threats and sensitive data, including, "for example, as alerts." Ex.1001 ('685 patent), 6:38-53, 6:19-53, 7:50-64. 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 ('685 patent), 6:38-53, 6:19-53, 7:50-64.

B. Claim Construction

42. In addition to the explicit constructions discussed below, I have applied the plain and ordinary meanings of the claim terms in the '685 patent, as a person of ordinary skill in the art would have understood them in the context of that patent in approximately January 2019.

1. “Analyzing the Snapshot”

43. Claims 1 and 22 recite “analyze the at least one snapshot to detect vulnerabilities and sensitive data, wherein analyzing the at least one snapshot requires no interaction with the virtual machine.” Similarly, claim 13 recites “analyzing the at least one snapshot to detect vulnerabilities and sensitive data, wherein analyzing the at least one snapshot requires no interaction with the virtual machine.” Dr. Stavrou labels these as elements 1.6, 13.6, and 22.6. Ex.1002 (Stavrou Decl.), ¶153.

44. Dr. Stavrou states that he was “asked to assume that ‘[analyze/analyzing] the at least one snapshot’ also encompasses analysis of a VM instantiated from the snapshot.” Ex.1002 (Stavrou Decl.), ¶83; *see also* ¶84 (“Accordingly, my discussion of Ground 1 below applies an understanding of ‘analyze the at least one snapshot’ in claims 1 and 22, as well as ‘analyzing the at least one snapshot’ in claim 13, as encompassing both direct analysis of the snapshot data and analysis of a VM instantiated from snapshot.”). Thus,

Dr. Stavrou treats limitations in the challenged claims that recite “[analyze/analyzing] the at least one 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.), ¶¶81-84.

45. I disagree that this claim element encompasses “[2] analysis of a VM instantiated from [the] snapshot” because that is not how a person of ordinary skill in the art would understand the claim element. Dr. Stavrou does not try to justify this claim element encompassing “analysis of a VM instantiated from a snapshot,” but simply states that he was “asked to assume” this claim construction. Ex.1002 (Stavrou Decl.), ¶83; *see* Ex.2077 (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.

46. A person of ordinary skill in the art would not understand the “[analyze/analyzing] the at least one snapshot” limitations in the ’685 patent to encompass “analysis of a VM instantiated from a snapshot.” As an initial matter, there appears to be no disagreement that an instantiated VM is not a snapshot—otherwise Dr. Stavrou would not have enumerated it as separate from “direct analysis of snapshot data” in his assumed construction. Rather, an instantiated VM is a VM, and analysis of a VM instantiated from a snapshot is an analysis of a VM—not a snapshot. Once a VM is instantiated from a snapshot, e.g., when a duplicate VM is instantiated from a snapshot for subsequent monitoring, the

instantiated VM is no longer a “snapshot,” and monitoring, assessing, or evaluating it is not “analyze/analyzing” a “snapshot” as claimed. This would have been readily understood by a person of ordinary skill in the art reading the '685 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 ('685 patent), 4:58-59. The claims explicitly state that “analyzing the at least one snapshot requires no interaction with the virtual machine.” The specification also distinguishes analyzing a snapshot by a “security system 140 to detect vulnerabilities” (Ex.1001 ('685 patent), 5:19-20), 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.* 5:54-61).

47. In fact, as I discuss further below, it appears that Dr. Stavrou “assume[d] that ‘[analyze/analyzing] the at least one snapshot’ also encompasses analysis of a VM instantiated from the snapshot,” as he was “asked” to (Ex.1002 (Stavrou Decl.), ¶ 83), simply to try to buttress opinions he offers that depend on disclosures from prior art references, such as Veselov, that discuss evaluating duplicate VMs, not snapshots. *See* Section VIII.A.2.b (discussing limitations 1.6, 13.6, 22.6).

VII. Overview of the Asserted References

A. Overview of Veselov

48. Veselov is titled “Security Assessment of Virtual Computing Environment Using Logical Volume Image” and discloses systems and techniques that perform a security assessment on “virtualized reproductions of computing resources” such as VMs or instances of VMs. Ex.1007 (Veselov), Title, Abstract, 3:20-23. Figure 1 illustrates an exemplary environment and components:




FIG. 1

49. As seen in Figure 1, Veselov discloses embodiments of a computing environment involving a computing resource service provider (shown as no. 104), in which user devices (shown as no. 102) request or manage resource allocation

for server computers (shown as no. 142). Ex.1007 (Veselov), 4:19-29. These server computers are allocated as virtual computing resources. Ex.1007 (Veselov), 4:19-29. Figure 1 also depicts an API (shown as no. 120) that allows a user to connect, through, for example, a web-based application, to the computing resource service provider, and provide input. Ex.1007 (Veselov), 4:29-41.

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

resources “may be costly monetarily” and will require “an extended period of downtime.” Ex.1007 (Veselov), 3:8-11.

51. Veselov’s overarching goal is therefore to avoid downtime of VMs and reduce assessment complexity when performing security assessments. Throughout its disclosure, Veselov aims to preserve the activity of the original virtual machine instance. For example, Veselov provides systems that “include a scanning service that facilitates duplication of all or a portion of the target computing resource, and then perform[] the security assessment on the duplicate computing resource”—“to avoid consuming processing time, processing power, and storage space of the target computing resource.” Ex.1007 (Veselov), Abstract; *see also* Ex.1007 (Veselov), 10:9-11 (generating a scannable volume from a snapshot includes “launching a duplicate virtual machine instance.”). And, Veselov provides that performing a robust security assessment may reduce the target computing device’s capabilities and may not be as responsive should an assessment be performed on the target machine. Ex.1007 (Veselov), 2:56-67. And Veselov cautions against “downtime.” Ex.1007 (Veselov), 3:8-11. Veselov explains that its solution allows the assessment to “be performed without impacting the customer’s virtual machine,” which can continue to “servic[e] requests from its customers or perform[] other useful work.” Ex.1007 (Veselov), 12:21-43. For the system to perform such “other useful work” Veselov must keep its VMs active

during the capture of a snapshot and subsequent vulnerability analysis. Additionally, Veselov 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.

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

53. In Veselov’s first solution, Veselov’s scanning service sets up and launches an active virtual machine instance (a duplicate machine for assessments). Ex.1007 (Veselov), 3:54-59. The duplicate machine has the same configuration as

the original machine instance at the time of the snapshot capture. Ex.1007 (Veselov), 3:54-59. Figures 3A and 3B are informative.





FIG. 3A

54. In Figure 3A, the scanning service performs the security assessment on the assessment machine via a security agent program. Ex.1007 (Veselov), 14:13-27. The scanning service receives assessment results via the agents on the assessment machine. Ex.1007 (Veselov), 14:28-34. The assessment machine can implement processing activity in order to demonstrate real-time dynamic or static data, network connections, and communications to other computing resources. Ex.1007 (Veselov), 15:41-59. Additionally, a user or administrator may 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.

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



56. In exemplary Figure 5A, the scanning service analyzes snapshot data in a “test environment” outside of the virtual computing environment. Ex.1007 (Veselov), 17:9-12. In both Figures 5A and 5B, the user communicates directly with the scanning service outside the target resource computing environment. Ex.1007 (Veselov), Figs. 5A-B. 17:9-67. An API is not depicted in Figures 5A and 5B, and is not used for communication between the user and scanning service. The scanning service interfaces with VM resources through provisioning services such as a hypervisor. Ex.1007 (Veselov), Figs. 5A-B. 17:9-67.

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

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

B. Overview of Hufsmith


59. Hufsmith is titled “Combined Threat Score for Container Images.” Ex.1078 (Hufsmith), Cover. Hufsmith discloses “a process for determining threat scores for container images or distributed applications that consider the results of a multitude of different scanners and other factors.” Ex.1078 (Hufsmith), Abstract. Hufsmith’s “Field” is disclosed as “techniques that combine metrics of heterogenous vulnerability scans of container images.” Ex.1078 (Hufsmith) ¶2. Hufsmith is therefore directed to methods for “determining combined threat score[s] for container images,” not snapshot-based assessments of VMs. *See* Ex.1078 (Hufsmith), Title, ¶2, ¶¶4-6. Hufsmith does not discuss snapshots at all.

60. Hufsmith discusses problems in the field of scanning container images and teaches solutions involving its techniques for analyzing scores for container images to determine a combined threat score for the images. For example, Hufsmith discloses that components of distributed applications, called “services,” may be composed of “one or more executing bodies of code.” Ex.1078 (Hufsmith), ¶5. And, Hufsmith explains that such “bodies of code and other resources” can be difficult to secure because, for example, services arise from “diverse sets of bodies of code,” “can undergo relatively frequent version changes,” and can be “downloaded from third parties ... that may be un-trusted or accorded less trust than code built in-house.” Ex.1078 (Hufsmith), ¶5.

Additionally, Hufsmith provides that in such “potentially highly complex” environments, even where potential vulnerabilities are detected, “a degree to which they pose a threat” in an execution environment with distributed code and other resources can be “difficult to quantify.” Ex.1078 (Hufsmith), ¶6. And, Hufsmith explains that such difficulties are heightened for applications built with containers such as DockerTM containers, Open Container Initiative (OCI) containers, or with KubernetesTM because these containers “generally virtualize at the operating system level, in contrast to virtual machines” which also emulate underlying hardware. Ex.1078 (Hufsmith), ¶26. For example, containers are generally formed from multiple layers and offer OS-level virtualization with various benefits not seen in virtual machines, “including lower computational load, faster spin up, and sharing of resources,” as well as the ability to reuse “intermediate images of read-only bodies of code and other resources” “across multiple container images.” Ex.1078 (Hufsmith), ¶¶26-27. However, such features can cause difficulty with securing distributed application due to “false positives” in the container context along with an inability to obtain holistic views into container exposure (i.e., because of the container’s layered nature). Ex.1078 (Hufsmith), ¶¶28-29. Also, Hufsmith explains that the prior art lacks “multi-sourced vulnerability assessments” and “combined threat scores for container images.” Ex.1078 (Hufsmith), ¶¶28-29.

61. Accordingly, Hufsmith teaches techniques to combine metrics of heterogenous vulnerability scans of container images. Ex.1078 (Hufsmith), ¶2. Hufsmith teaches that its techniques are useful in the context of “software development related to distributed applications” (Ex.1078 (Hufsmith), ¶2) and its techniques assist software developers during development of container images, i.e., before containers are instantiated from a container image. Ex.1078 (Hufsmith), ¶153.

62. Hufsmith provides an architecture to implement a scanning engine to facilitate the vulnerability scans. Ex.1078 (Hufsmith), Fig. 1A.



63. Figure 1A depicts a vulnerability scanning engine (shown as no. 12) that implements a controller (shown as no. 42), a schema translator (shown as no. 44), a scan selector (shown as no. 46), a layer evaluator (shown as no. 50), a scan configurer (shown as no. 48), and a results engine (shown as no. 54) that coordinate

to arbitrate which layers and which portions of layers will be scanned by one of various scanning applications (shown as no. 16). Ex.1078 (Hufsmith), ¶68. The vulnerability scanning engine decides which scanner application will perform the scan among a set of heterogeneous scanner applications that are each configured to apply a different type of scan to a different type of body of code and other resources such as images and files. Ex.1078 (Hufsmith), ¶68.

64. Figure 1B demonstrates the results engine (shown as no. 54) communicatively coupled (via a network shown as no. 21) with the container manager (shown as no. 20) and the container score database (shown as no. 102).




FIG. 1B

65. Hufsmith provides that the results of the vulnerability scanning engine will be translated into a unified schema of scanner results or properties, which may be stored in a database “in association with an identifier of the container image.” Ex.1078 (Hufsmith), ¶101. The database may be a container score database which comprises a number of container score records (shown as nos. 101A to 101N) that are referenced by, for example, a cryptographic hash value of the container image. Ex.1078 (Hufsmith), ¶101. In these scenarios, where the container score records correspond to container images that constitute a distributed application, a given container score record may reference one or more other container score records. Ex.1078 (Hufsmith), ¶104. These container score records “assist a developer in determining which one of a plurality of container images to utilize based on a risk associated with each container image.” Ex.1078 (Hufsmith), ¶103.

66. Figure 4, below, illustrates a process for “determin[ing] one or more score records for a distributed application having one or more container images. Ex.1078 (Hufsmith), Fig. 4, ¶¶18, 199.




FIG. 4

67. First, Hufsmith provides that the system obtains a “distributed application” (depicted by block 420). Ex.1078 (Hufsmith), ¶201. The second step (depicted by block 421) provides that the system will identify container images associated with the distributed application “(which may be a multi-container distributed application).” Ex.1078 (Hufsmith), ¶202.

68. Then, at the third step (depicted by block 422), the system creates a distributed application score record for the distributed application. Ex.1078 (Hufsmith), ¶¶201-202. And, when the system determines that a container image to process exists (depicted by block 422), the system will execute steps 402-408. Ex.1078 (Hufsmith), ¶¶201-202. The system therefore obtains a container image (depicted by block 402). Ex.1078 (Hufsmith), ¶200. Then, scanner properties are determined for the container image (depicted by block 406). Ex.1078 (Hufsmith), ¶200. The system then proceeds to Figure 2 (below). Ex.1078 (Hufsmith), ¶200.




FIG. 2

69. At Figure 2, the system first obtains a container image (depicted by block 202). Ex.1078 (Hufsmith), ¶152. Then, Hufsmith provides that the system determines whether there are additional layers in the container image to be processed (depicted by block 204). Ex.1078 (Hufsmith), ¶152. For example, the system will start with the base or top layers of a multi-layer container, and, if there are additional layers to process, then the system will continue to select a layer (depicted by block 206). Ex.1078 (Hufsmith), ¶152. Then, the system determines whether there is additional scanner criteria that should be applied to a given layer (depicted by block 208). Ex.1078 (Hufsmith), ¶152. If the system determines that there are additional criteria to apply, then an option for “next scanner criteria” may be selected (depicted by block 210). Ex.1078 (Hufsmith), ¶152. The system determines whether a given layer is satisfied by selected criteria (depicted by block 212), and if so, the scanner that corresponds to the chosen criteria will be selected to scan the given layer (depicted by block 214). Ex.1078 (Hufsmith), ¶152. Next, the system executes a command to the scanner to scan, and the results of the scanning are received and translated (depicted by blocks 216-222). Ex.1078 (Hufsmith), ¶152. Then, once there are no further layers to scan, the system will filter potential vulnerabilities and calculate metrics (depicted by blocks 224, 226). Ex.1078 (Hufsmith), ¶152. Subsequently, the system may store results and present results (depicted by blocks 228, 230). Ex.1078 (Hufsmith), ¶152. Then, the system

resumes at Figure 4, where a container score record for the container image may be created and a distributed application score may be updated. Ex.1078 (Hufsmith), ¶¶200, 202. The system will continue to process container images iteratively until it finds no additional container images to process. Ex.1078 (Hufsmith), ¶202 (describing blocks 426, 430 of Figure 4).

70. Figure 5 illustrates a process for “determin[ing] a combined threat score for a distributed application having one or more container images or a container image.” Ex.1078 (Hufsmith), ¶19, 204, Fig. 5.





FIG. 5

71. Hufsmith's system of Figure 5 begins by receiving a request for a combined threat score (depicted by block 502). Ex.1078 (Hufsmith), ¶205. Then, the system determines whether a "score record for a distributed application or container image" exists, "such as a score record created by [the] process of FIG. 4." Ex.1078 (Hufsmith), ¶205. If no score record exists, Hufsmith provides that the "process 400 of FIG. 4 may be called to create" one. Ex.1078 (Hufsmith), ¶206. Then, after the system has a score record, the record "may be examined for scanner property weights" (depicted by block 506), which may be determined and stored to the score record if none already exist (depicted by block 508). Ex.1078 (Hufsmith), ¶206. Also, a score record "may be examined for context property weights" (depicted by block 510). Ex.1078 (Hufsmith), ¶207. And, if no context property weights exist, then the system includes an additional step of "requesting context property information, such as environment information," from which "weights are determined for the context properties at block 514." Ex.1078 (Hufsmith), ¶207. Then, weights associated with different scanner properties are modified based on weights of corresponding context properties (depicted by block 516). Ex.1078 (Hufsmith), ¶208. And, after "each of the context property weights are applied, a combined threat score may be generated based on the modified scanner property weights" (depicted by block 518) and reported (depicted by block 520). Ex.1078 (Hufsmith), ¶¶208-209.

72. Accordingly, Hufsmith “determin[es] threat scores for container images or distributed applications that consider the results of a multitude of different scanners and other factors such as context information” for “container images.” Ex.1078 (Hufsmith), Abstract.

C. Overview of Price

73. Price is titled “Security Assessment of Virtual Machine Environments” and is generally directed to performing security assessments for offline and online VMs. Ex.1048 (Price), Cover, Abstract. More specifically, Price involves a set of virtual machines managed by a virtual machine manager, and for each virtual machine, Price determines whether the machine is offline. Ex.1048 (Price), Abstract. For the machines that are determined to be offline, Price’s system implements techniques to collect an image from the offline machine. Ex.1048 (Price), Abstract. In contrast, for the machines that are determined to be online, Price implements agents that are loaded onto each machine and can assess security of the online machines. Ex.1048 (Price), Abstract.



74. Figure 1 depicts the environment in which Price's techniques are implemented. Ex.1048 (Price), Fig. 1. Price provides a computing system (not shown, no. 100 in Fig. 1) with security tools (shown as no. 105) that perform scans on VMs (shown as no. 120) in connection with policy compliance, vulnerability assessment, malware protection, among other security services. Ex.1048 (Price), ¶16. Price also discloses one or more client devices (shown as no. 135) for communicating over a network (shown as no. 110) with other system devices. Ex.1048 (Price), ¶16. The VM manager (shown as no. 130) manages a set of VMs and provides administrators and customers with interfaces for deploying and

maintaining VMs hosted within the server pool (shown as no. 125) in the computing system. Ex.1048 (Price), ¶16.

75. Price provides a series of steps for implementing its scanning techniques. Ex.1048 (Price), Fig. 6. First, the system identifies a VM manager that manages a set of VMs (or VMs outside of the set). Ex.1048 (Price), ¶44. The system may also identify each VM within the set with aspects and conditions of each respective VM. Ex.1048 (Price), ¶44. With this information, the system can determine whether a given VM is online or offline. Ex.1048 (Price), ¶44.






FIG. 6

76. For online machines, the system loads agents onto the target VM via an interface of the VM manager. Figures 3B and 3C are informative:




77. Figure 3B demonstrates a scenario where the security server obtains the identify and locations of VMs and loads agents onto the VMs (or loads agents based on the guest operating systems of the VMs). Ex.1048 (Price), ¶37.



78. Figure 3C of Price demonstrates a scenario where agents are scanning resources of online VMs. Ex.1048 (Price), ¶38. Through the scanning procedures, the agents can generate results data to return to the security server. Ex.1048 (Price), ¶38.

79. For offline machines, the system collects images through the VM manager interfaces or APIs. Ex.1048 (Price), ¶41. Figure 4 is informative:



80. Figure 4 demonstrates a scenario where the system determines that given VMs are offline and collects VM images. Ex.1048 (Price), ¶41. The system can then use the image data to complete security assessments. Ex.1048 (Price), ¶41.

81. In both offline and online scenarios, the methods conclude with a determination of results from the security assessments. Ex.1048 (Price), Fig. 6, ¶¶41, 43, 45. Neither scenario utilizes or analyzes a snapshot of a VM.

VIII. The Claims Of The '685 Patent Would Not Have Been Obvious

82. Petitioner and Dr. Stavrou challenge claims 1-22 of the '685 patent based on two grounds of purported obviousness, each of which depends on the combination of U.S. Pat. No. 11,216,563 ("Veselov," Ex.1007) and U.S. Patent App. Pub. No. 2020/0097662 ("Hufsmith," Ex.1078).

83. In my opinion, none of the claims of the '685 patent would have been obvious over Petitioner's and Dr. Stavrou's prior art, for reasons I explain below.

A. Ground 1: Claims 1-22 Would Not Have Been Obvious Over Veselov and Hufsmith

84. In Ground 1 of the Petition, Petitioner and Dr. Stavrou challenge claims 1-22 of the '685 patent based on the combination of Veselov and Hufsmith. *See* Ex.1002, Stavrou Decl. §XII.

85. In my opinion, Petitioner and Dr. Stavrou fail to show that any of claims 1-22 of the '685 patent would have been obvious over the combination of Veselov and Hufsmith.

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

a. Dr. Stavrou's analysis mixes different disclosures from Veselov without motivation to combine those disclosures.

86. 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* Ex.1002 (Stavrou Decl.), ¶122 (opining about a “communication interface”). Veselov discusses APIs only in relation to the examples in Figures 1 and 3A-3B, where computing environments are implemented by a computing resource service provider.¹ *See, e.g.,* Ex.1007 (Veselov) 5:1-8 (In relation to Figure 1, “[t]he scanning service 110 may access the computing resource service provider 104, or any associated virtual computing environments and/or computing resources, via the API 120 or another API. In some embodiments, the scanning service 110 may use the API 120 to provide user interfaces to the user device 102, enabling the user to configure settings of the scanning service 110 as described further below.”); *see also* 10:37-16:14 (discussing Figures 3A-3B).

87. The Figure 1 example, reproduced below, includes snapshot storage 108, also described as “snapshot storage service 108” that “may be a group of computer systems configured to store snapshot data 146 and make the snapshot data 146 accessible to one or more other computer systems, such as the scanning

¹ I note that “computing resource service provider” is not typically used in the art to refer to a cloud computing platform. Additionally, Veselov does not expressly describe the APIs as cloud computing platform APIs or the computing resource service provider as a cloud computing platform.

service 110.” Ex.1007 (Veselov) 7:34-37; *see also* 7:37-8:15 (describing snapshot storage service 108).




FIG. 1

Ex.1007 (Veselov) Fig. 1

88. In the context of systems that use snapshot storage, Veselov only discusses APIs in relation to a user device transmitting a request to the scanning service for a security assessment, where the “request may be an API call including information corresponding to the requestor, the user associated with the virtual machine 112, the on-demand storage service 108, and/or the particular virtual machine 112.” Ex.1007 (Veselov) 6:32-38. Veselov teaches that the API may

enable the user “to provide settings, commands, software packages, and other user input, to the computing resource service provider 104” and a user “may be required to provide credentials and be authenticated and authorized to modify its virtual computing environments and virtual resource allocations via the API 120” Ex.1007 (Veselov) 4:35-48. Veselov discloses that the scanning service “may determine the corresponding virtualization layer 144 for the virtual machine 112 identified by the request and transmit a comment to the virtualization layer 144 to provide snapshot data 146 (or to provide access to snapshot data 146) corresponding to the virtual machine 112.” Ex.1007 (Veselov) 6:38-43. The virtualization layer itself “may be configured to obtain the snapshot data 146 and send it to the scanning service 110” or store it in the storage service 108. Ex.1007 (Veselov) 6:43-46. But these disclosures do not teach or suggest identifying, querying, or returning the location of a virtual disk, let alone using APIs to do so. To the contrary, the scanning service is interacting with the virtualization layer, which provides snapshot data without the scanning service needing to identify a virtual disk let alone locate it. For further example, when snapshots are stored in the storage service, e.g., in a queryable database, Veselov discloses only that it is queryable based on the snapshot’s own identifier, or by “the identifier of the virtual machine instance from which the snapshot data 146 was derived.” Ex.1007 (Veselov) 7:50-53. Veselov’s scanning service does not have any need or benefit,

absent hindsight, to identify or locate a virtual disk. In summary, the Figure 1 embodiment of Veselov, on its own, cannot meet, e.g., the virtual disk limitations of claim 1, and it is thus unsurprising that Dr. Stavrou mixes disclosure of different examples to try to meet those limitations.

89. The Figure 3A/3B examples are different from Figure 1, including in that they do not include a snapshot storage service. 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 this embodiment, 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. There is no disclosure of using—and no suggestion or need to use—the API to identify a virtual disk or query a location of a virtual disk in this Figure 3A/3B example. Further, the use of an assessment virtual machine does not teach analyzing the snapshot as required by the claims. *See* Section VI.B.1. Thus, the Figure 3A/3B embodiment on its own also cannot meet the claims, including, e.g., the analyzing limitation and the virtual disk limitations.

90. 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. In the Figs. 5A-5B example, the resources are not implemented within a Computing Resource Service Provider (i.e., what Dr. Stavrou appears to identify as a cloud computing platform), which is the entity that provides the API

in Figs. 3A-3B. *See* Ex.1002 (Stavrou Decl.), ¶131 (referring to the user utilizing “the computing resource service provider’s API”). As such, it is not surprising that no API is illustrated in Figs. 5A-5B, let alone cloud computing platform APIs, or the use of such APIs as claimed in the ’685 patent. *See, e.g.,* Ex.1007 (Veselov), Figs. 5A-5B, 17:9-66.




FIG. 5A

91. As shown in Fig. 5B, below, the scanning service 510 communicates with the virtual computing environment through “a provisioning service” such as “a hypervisor 518.” Ex.1007 (Veselov) 17:30:35. Hypervisors and APIs are not the same, and a person of ordinary skill in the art would have understood that they serve different purposes. *See, e.g.,* Ex.2021 [Amazon Web Service (AWS), *What is a Hypervisor?*] (“The hypervisor relays requests for processing power, memory, storage, and other resources to the host machine in several ways, including API calls. An API is a software communication method that allows different applications to exchange data.”). Ex.2022 [Perez-Botero, *Characterizing Hypervisor Vulnerabilities*], §1 (“Virtualized environments are usually implemented with the use of a Hypervisor, which is a software layer that lies between the Virtual Machines (VMs) and the physical hardware. The Hypervisor allocates resources to the VMs, such as main memory and peripherals.”); Ex.2018 [Abraham Silberschatz et al., *Operating System Concepts* 716 (9th ed. 2013)], 63 (“Typically, application developers design programs according to an application programming interface (API). The API specifies a set of functions that are available to an application programmer, including the parameters that are passed to each function and the return values the programmer can expect”).




FIG. 5B

Ex.1007 (Veselov) Fig. 5B

92. Dr. Stavrou fails to explain why a person of ordinary skill in the art would have drawn from separate disclosures of these different examples, including why an API would be used in Veselov's Figure 5A-5B example—i.e., the only example Veselov discloses that is not directed to assessments of duplicate VMs. Nor does he explain why a person of ordinary skill in the art would have picked that one feature from the Figs. 3A-3B examples and not other features (e.g., assessing duplicate VMs) that do not meet the claims. The absence of any API in the Figures 5A-5B example seems to be what motivated Dr. Stavrou to just

“assume,” based on the direction of Petitioner’s counsel and without any explanation or support, that “analyzing [a] snapshot” encompasses analyzing duplicate VMs. *See* Section VI.B.1 (explaining that assessments of duplicate VMs are not analyses of snapshots and that Dr. Stavrou’s contrary assumption is incorrect). To put it simply, interactions between the computing device 502 and other components in Figures 5A-5B of Veselov, the only examples of Veselov that analyze a snapshot, **do not** use APIs, contrary to the express requirements of the claims, and there is no motivation or assertion in Dr. Stavrou’s declaration that Veselov would be modified to take components of other examples and import them into Figures 5A-5B, where such components do not exist.

93. Dr. Stavrou acknowledges that “Veselov describes several different embodiments and figures” but dismisses any differences between them by asserting that “a POSA would have understood that the techniques and components described in the context of a particular figure can also applied [*sic*] in the context of the other embodiments and figures.” Ex.1002 (Stavrou Decl.), ¶91. But even if a technique or component “can” apply to another embodiment or figure, that does not mean the specific techniques or components that Dr. Stavrou mixes between different examples would have been understood to apply by a person of ordinary skill in the art. Dr. Stavrou’s citations of Veselov support this fact. For example, although Dr. Stavrou asserts that Veselov “explains that the techniques discussed

in the context of Fig. 2 can be performed by the various systems described in other embodiments,” he only quotes Veselov as teaching that “[t]he process 200 may be performed by any *suitable* system such as a scanning service or component thereof as described above in connection with FIG. 1.” Ex.1002 (Stavrou Decl.), ¶91 (quoting Veselov 8:52-62). But Dr. Stavrou makes no attempt to show that, e.g., the Figs. 5A-5B systems are “suitable” for the modifications he applies from the Figs. 3A-3B examples. Nor could he, because as described above, the Figs. 5A-5B examples do not employ a computing resource service provider or APIs. *See* Ex.1007 (Veselov) 10:37-39 (Figures 3A-3B use “computing environments implemented by a computing resource service provider”), 16:56-57 (Figs. 5A-B illustrate “additional embodiments” without computing resource service providers or APIs).

94. I have been informed and understand that in an obviousness analysis, it is not appropriate to mix disclosures from multiple references, or multiple examples in one reference, without adequate explanation for why a person of ordinary skill in the art would have done the same. Dr. Stavrou has not provided such explanation.

b. Dr. Stavrou’s proposed motivations to combine Veselov and Hufsmith are not based on what a person

of ordinary skill in the art would have understood at the relevant time

95. Dr. Stavrou asserts that a person of ordinary skill in the art would have been motivated to make the combination of Veselov and Hufsmith on which both of Petitioner's grounds depend. *See* Ex.1002 (Stavrou Decl.), ¶104. I disagree.

96. To a person of ordinary skill in the art, Veselov and Hufsmith 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 the teachings of Hufsmith as Petitioner and Dr. Stavrou assert for their challenges to the '685 patent claims. 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 '685 patent's claims, looking for disclosures in disparate references, without meaningfully considering whether the references they cite would, in fact, be combined in the first place. As noted above, I understand that obviousness arguments based on this type of hindsight are improper.

97. First, Dr. Stavrou asserts that the combination of Hufsmith and Veselov "would provide a more complete assessment of potential security concerns." Ex.1002 (Stavrou Decl.), ¶108. Dr. Stavrou further asserts that a person of ordinary skill in the art "would have also been motivated to combine the teachings of Veselov and Hufsmith to improve security by facilitating a more

efficient, targeted review of critical security issues.” Ex.1002 (Stavrou Decl.), ¶110. I disagree.

98. As I explain below, Veselov and Hufsmith analyze entirely different types of data relating to different entities. As a result, a person of ordinary skill in the art would not have been motivated to apply Hufsmith’s assessments to Veselov, or vice versa. Nor would a person of ordinary skill in the art have had a reasonable expectation of success in combining them—including because Hufsmith’s analysis relates to scanning pre-deployment source code that is not available in Veselov.

99. Veselov discloses analysis of a “logical volume image” including, e.g., with snapshots. As I discussed above, Hufsmith does not discuss snapshots. *See* Section VII.B. Nor does Hufsmith discuss any type of analysis specific to a particular “volume.” Hufsmith relates to a different solution than any problem in Veselov. Specifically, Hufsmith is directed to “techniques [for] combin[ing] metrics of heterogeneous vulnerability scans of container images,” by, e.g., scanning the source code of container images, to provide combined threat scores for container images and multi-container distributed applications. Ex.1078 (Hufsmith) ¶¶2, 184. Hufsmith applies weights to metrics, and annotations to source code, neither of which naturally fit with Veselov. Veselov does not disclose metrics, and does not have access to such image source code for annotation.

100. Although Veselov refers to its snapshots as “***volume*** images,” these snapshots are *not* the same as the “container images” and “virtual machine images” used in Hufsmith. Dr. Stavrou clarified at his deposition, “snapshots and images are different.” Ex.2077 (Stavrou Tr.) 272:7-13. I agree. A VM’s snapshot is, as the ’685 patent states, “a copy of the machine’s virtual disk (or disk file) at a given point in time.” Ex.1001 (’685 patent), 4:58-59. A VM image, in contrast, is a data file that “include[s] an installed and configured guest operating system and applications” and is used for running a virtual machine. Ex.2018 [Abraham Silberschatz et al., *Operating System Concepts* (9th ed. 2013)], p.716; *see also*, e.g., Ex.2005, [US Pat. App. Pub. No. 2008/0320583], ¶¶160-63 (distinguishing a “disk image file,” which is “a virtual disk file that stores the contents of the virtual machine’s hard disk drive,” and a “snapshot file,” which is “a file that stores the running state of the VM at a particular moment in time, usually in terms of deltas with respect to a baseline state.”); Ex.2020 [Huawei, *Snapshots vs. Images*] (explaining that images “[c]an’t be created for running instances” whereas snapshots “record[] the state and data of the virtual machine at the moment the snapshot was taken”); *Compare* Ex.2006 [Red Hat Virtualization 4.3 – *Virtual Machine Snapshots*] (“Snapshot are a storage function that allows an administrator to create a restore point of a virtual machine’s operating system, applications, and data at a certain point in time.”) and Ex.2007 [Birk, *Technical Challenges*], 4

(“Snapshots, also referred to as forensic image, of virtual machines provide a powerful tool with which a virtual machine can be cloned by one click including also the running system's memory. This leads to the main benefit that systems hosting crucial business processes do not have to be shutdown for performing a forensic analysis.”) *with* Ex.2025 [Red Hat Open Stack], p. 5 (“A virtual machine image is a file that contains a virtual disk which has a bootable operating system installed on it.”). A container image, as referred to in Hufsmith, is even farther afield from Veselov’s “volume images” or snapshots, because of the additional differences between containers and VMs as discussed below. *See, e.g.*, Ex.1078 (Hufsmith), ¶27; Ex.2075, Wiz Academy, *Containers vs. VMs: What’s the difference?* (April 8, 2024) (“In a nutshell, containers and virtual machines (VMs) are two inherently different approaches to packaging and deploying applications/services in isolated environments.”).

101. Additionally, a person of ordinary skill in the art would have known that snapshots and images are used to address different classes of vulnerabilities. For example, a point-in-time snapshot of a computer system disk is known to allow capturing and cleaning of files infected with malware at the time when the snapshot was taken. Ex.2023 [U.S. Pat. No. 8,220,053 B1 (“Sun”)], 5:49-58. In contrast, because VM images are used to instantiate VMs, scanning, detecting, and removing of malware from offline VM images would prevent the spread of

malwares. Ex.2024 [Wei, *Managing Security of Virtual Machine Images*], Sec. 2.3 (“Anecdotal evidence shows that, if dormant VM images are not managed (e.g., scanned for worms), a virtual environment may never converge to a steady state, because worm-carrying VM images can sporadically run, infect other machines, and disappear before they can be detected”).

102. Dr. Stavrou also recognizes that snapshots and images are not the same thing. *See, e.g.*, Ex.1002 (Stavrou Decl.) ¶37 (“Snapshots *can be more risky than images* because snapshots contain the contents of RAM memory at the time that the snapshot was taken, and this might include sensitive information that was not even stored on the drive itself.”) (emphasis added) (quoting Ex.1018 at 2-6).

103. A person of ordinary skill in the art thus would have recognized that Hufsmith’s teachings were not compatible with those of Veselov, and would not have sought to combine the teachings of the references as Dr. Stavrou asserts.

104. Hufsmith’s security assessments of container images are different than Veselov’s snapshot-based security assessments of virtual machines and a person of ordinary skill in the art would not have combined them. For example, as Wiz acknowledges on a page on Wiz’s website, “[i]n a nutshell, containers and virtual machines (VMs) are two inherently different approaches to packaging and deploying applications/services in isolated environments.” Ex.2075, Wiz Academy, *Containers vs. VMs: What’s the difference?* (April 8, 2024). As a

result, a person of ordinary skill in the art would not have reasonably expected to succeed in applying, e.g., Hufsmith's threat scores to Veselov because, for example, Dr. Stavrou does not explain how Veselov's snapshots could provide sufficient information in order to determine such scores, metrics, or weights, and further to filter or prioritize alerts on that basis. As discussed above, Hufsmith's system has access to code from the container images that Veselov's snapshots would not provide.

105. Dr. Stavrou characterizes Hufsmith as providing a "more detailed discussion of detecting and reporting particular types of security issues such as malware, CVEs, and sensitive data" (Ex.1002 (Stavrou Decl.) ¶108), but nothing that Dr. Stavrou cites from Hufsmith discusses detecting "sensitive data" or security issues relating to such data. *See* Ex.1078 (Hufsmith), ¶¶8, 28, 84, 104, 113-22. Paragraph 8 of Hufsmith discusses a "plurality of scanner properties pertaining to a container" and "determining a combined threat score for the container based on the at least some of the scanner properties" but does not discuss detecting "sensitive data" or security issues relating to such data. Ex.1078 (Hufsmith) ¶8. Similarly, paragraph 28 discusses various drawbacks of "[t]ypical scanning techniques" but is silent about "sensitive data." Ex.1078 (Hufsmith) ¶28. Paragraph 84 discloses scanning "identified resources" to identify vulnerabilities including "those identified in public repositories," but is also silent about "sensitive

data.” Ex.1078 (Hufsmith) ¶84. Paragraph 104 relates to container score records but is also silent about “sensitive data.” Ex.1078 (Hufsmith) ¶104. Paragraphs 113-122 (ten contiguous paragraphs that Dr. Stavrou cites without explanation) relate to a CVE evaluator and malware evaluator but again are silent about “sensitive data.” Ex.1078 (Hufsmith) ¶¶113-122. At best, Hufsmith discloses what Dr. Stavrou calls “avoiding the potential exposure of sensitive data.” Ex.1002 (Stavrou Decl.) ¶108 (citing Ex.1078 (Hufsmith) ¶¶117-118, 181). But this does not support Dr. Stavrou’s characterization of Hufsmith as *detecting* and *reporting* sensitive data, which is what is claimed. Rather, Hufsmith only discloses, e.g., indicators of *potential exposure* of restricted information, including “[i]mpact metrics” such as “a confidentially impact (e.g., ... (N) no loss of restricted information, (L) some restricted information may be obtained, or (H) total loss of restricted information).” Ex.1078 (Hufsmith) ¶118.

106. Dr. Stavrou further asserts that “[i]t was well known that prioritizing and filtering alerts allowed for rapid identification of which vulnerabilities are the most dangerous to a virtual asset.” Ex.1002 (Stavrou Decl.) ¶110 (citing Ex.1078 (Hufsmith) ¶138). But Hufsmith only mentions alerts once, in paragraph 15 in relation to Figure 3A, specifically as it relates to “a plugin of a integrated development environment to annotate code specifying container images with alerts relating to potential security vulnerabilities.” Ex.1078 (Hufsmith) ¶15. As

discussed above, annotating code in pre-deployment systems is not relevant to Veselov, at least because source code is not available in Veselov's systems which analyze deployed virtual machines—but even so, this passage does not disclose “prioritizing and filtering alerts” at all. Cf. Ex.1002 (Stavrou Decl.) ¶110. And paragraph 138, which Dr. Stavrou cites, similarly discusses modifying “weights” associated with vulnerabilities based on revisions of pre-deployment container images. Ex.1078 (Hufsmith) ¶138.

107. Additionally, Dr. Stavrou asserts that “Hufsmith itself further teaches the benefits of alert prioritization and filtering.” Ex.1002 (Stavrou Decl.) ¶111 (citing Ex.1078 (Hufsmith) ¶¶6, 165, 180, 183, 192, 155. Dr. Stavrou does not explain how any concern relating to alleged alerts of Hufsmith applies to Veselov. These passages explicitly relate to annotating source code in a development (pre-deployment) environment and are not applicable to Veselov. Ex.1078 (Hufsmith) ¶153. Additionally, as discussed above, Hufsmith analyzes “container images in complex environments.” See Section VII.B. Such containers, according to Wiz's webpage, were known to be “less secure” than virtual machines “because there is no network isolation” and “containers [are] more prone to vulnerabilities.” Ex.2075, Wiz Academy, *Containers vs. VMs: What's the difference?* (April 8, 2024) at 6. Dr. Stavrou does not address how the same concerns apply to Veselov, which assesses particular volume images and provides assessment results to an

“owner or administrator of the virtual machine.” Ex.1007 (Veselov) 4:5-8. Nor does Veselov provide any indication that alert fatigue is a concern in its system. Hufsmith focuses on environments with multiple scanner applications that can each report their own results. Ex.1078 (Hufsmith) ¶¶71, *see also id.* ¶¶68 (“heterogeneous set of scanner applications”), 73 (“each scanner application may have different criteria”), 79-85 (describing plurality of “scanner applications”), 85-86 (describing each scanner returning results). That is not the case in Veselov. There is no reason to implement the score-based system in Veselov without these multiple scanners, which would add significant complexity to Veselov. In fact, Veselov leads a person of ordinary skill in the art away from modifications that add such complexity. Veselov teaches that its disclosed system “overcome[s] the drawbacks of existing security systems” including “complexity of installation and monitoring of the assessment system.” Ex.1007 (Veselov) 2:42-46. And Hufsmith’s disclosures relating to source code annotations in a development environment teach applications in a pre-deployment scenario where products are not being assessed in a live setting, unlike Veselov. Ex.1078 (Hufsmith) ¶¶153-156, 180-183.

108. To the extent that Hufsmith even discloses prioritized alerts in this scenario—which I don’t agree with (see Section VIII.A.2.c)—a person of ordinary skill in the art would have understood such prioritized alerts to be only applicable

in the pre-deployment development context, i.e., during coding of a container, so that the container has as few major vulnerabilities as possible by the time of deployment. Once the container is deployed, it would be expected to be substantially free of vulnerabilities if one were to follow Hufsmith's techniques. So, the concerns detailed in Hufsmith, where containers are actively being developed to be free of vulnerabilities, would not be the same in a live setting, where the container image is deployed and executing and source code is not available to change. Veselov's goal of keeping the VMs active during assessment is essentially moot in a pre-deployment scenario where a VM is not actively deployed. *See* Section VII.A. Nor would a person of ordinary skill in the art have had a reasonable expectation of success in applying Hufsmith's teachings to Veselov, including because, as discussed above, Veselov's system does not have access to source code to provide such annotations.

109. Dr. Stavrou also asserts that “[a] POSA would have been motivated to obtain... [Hufsmith's alleged] benefits in the types of security scans taught by Veselov, since such scans would typically detect numerous security issues of varying importance.” Ex.1002 (Stavrou Decl.) ¶111 (citing Ex.1007 (Veselov) 2:21-41, 10:17-36, 14:34-36, 16:47-55). This is not how a person of ordinary skill in the art would characterize Veselov. For example, the column 10 passage cited by Dr. Stavrou discloses “perform[ing] the security assessment on the scannable

volume,” which involves “applying the rules of the assessment to the properties of the reproduced target resource to produce assessment results that describe ‘findings’ relative to the assessed properties.” Ex.1007 (Veselov) 10:17-22. Nothing in this passage suggests combination with Hufsmith. For example, applying such “rules” to obtain “results” does not teach a person of ordinary skill in the art that Veselov’s scan would “typically detect numerous security issues of varying importance” as Dr. Stavrou asserts. Ex.1002 (Stavrou Decl.) ¶111. In fact, the application of “rules” suggests a finite set of predictable “results” that do not need prioritization and such that Veselov’s system would not suffer from alleged alert fatigue. *See* Ex.1007 (Veselov) 14:34-36 (“The scanning service 310 may compare the received outcomes to the rules in the rules package to produce the assessment results.”); 2:21-41 (discussing usage of “rules packages” that define the security assessments). Similarly, the column 16 passage that Dr. Stavrou cites, the security assessment is performed “by activating the security agent program to execute one or more *tasks* of the security assessment” and “receiving outcomes of the tasks.” Ex.1007 (Veselov) 16:47-53. Again, the use of such “tasks” does not suggest detecting “numerous security issues of varying importance” for the same reasons as the “rules” of column 10. And Veselov confirms that its “assessment results” can be evaluated by an “owner or administrator” of the VM without any issue. Ex.1007 (Veselov) 4:5-8.

110. Dr. Stavrou discusses various references that are not in the combination, and in doing so repeatedly equates “scores” or “risk” with alerts. Dr. Stavrou asserts that “[a]lerts were commonly prioritized in a variety of ways.” Ex.1002 (Stavrou Decl.) ¶68 (citing Ex.1075 (Mohanty) 6:45-9:26). Dr. Stavrou characterizes Mohanty as “discussing multi-factor calculation of prioritization score.” *Id.* This passage of Mohanty does not prioritize alerts at all. To the contrary, Mohanty discloses that when a “prioritization score” reaches a “predetermined threshold,” an “alert” can be produced. Ex.1075 (Mohanty) 9:11-19. In other words, the prioritization score is not an alert, nor is it a prioritization of alerts. Instead, the alerts that are generated are based on a simple “predetermined threshold” such that the alerts are explicitly *not* prioritized, but rather they all indicate the same thing—the score surpassing a predetermined threshold.

111. Dr. Stavrou further states that “some systems prioritized alerts based on the probability that an asset had a particular vulnerability.” Ex.1002 (Stavrou Decl.) ¶68 (citing Ex.1008 (Basavapatna) Abstract). Again, this cited passage of Basavapatna says nothing about prioritized alerts at all, but rather discloses a “vulnerability detection score” that “indicates an estimated probability that a particular asset possess[es] the particular vulnerability.” Ex.1008 (Basavapatna) Abstract.

112. Dr. Stavrou also asserts that “[a]nother well-known approach was to prioritize by the potential impact or harm that would occur in the event the associated vulnerability was exploited.” Ex.1002 (Stavrou Decl.) ¶68 (citing Ex.1044 (Giakouminakis) 11:51-67. This cited passage, again, is silent about prioritized alerts, and instead discusses “Prioritization of risk” and “priority of remediation of vulnerabilities.” Ex.1044 (Giakouminakis) 11:51-67.

113. Dr. Stavrou’s other tertiary references, not part of the asserted combination, have similar deficiencies and do not motivate combination of Veselov and Hufsmith. Basavapatna, Klemperer (Ex. 1014), Giakouminakis, Canzanese (Ex. 1045), Crabtree (Ex. 1046), Hibbert (Ex. 1063), and Mohanty do not even mention containers or container images and therefore provide no support for why a person of ordinary skill in the art would consider Hufsmith when analyzing the techniques of Veselov. Wysopal (Ex. 1036) discusses container files that are implemented by virtual machines to perform a security analysis, consistent with my statements above, but Wysopal does not motivate combination of Hufsmith’s container-image-based analysis with Veselov’s snapshot-based analysis.

114. Dr. Stavrou further asserts that persons of ordinary skill in the art “would have had a reasonable expectation of success” in making the purported combination. Ex.1002 (Stavrou Decl.), ¶112. But, for the reasons I discussed

above, that is not true. Veselov and Hufsmith do not have “similar technical contexts” (Ex.1002 (Stavrou Decl.), ¶112), nor do Veselov or Hufsmith even teach a majority of the limitations. Additionally, combining the Hufsmith techniques for analyzing container images with Veselov’s practice for instantiating virtual machine instances with snapshots would render unexpected results.

115. Finally, evidence from Wiz and others suggests that Dr. Stavrou’s assertions about combining Veselov and Hufsmith are driven by hindsight. I understand that it is Dr. Stavrou’s burden to establish a motivation to combine cited references in a combination and that it is not sufficient to simply state what is possible for a system to accomplish. Additionally, the motivation cannot rely on the hindsight teachings of the challenged patent.

116. For example, I have reviewed documents that appear to show Wiz making public statements that the problem of detecting and managing sensitive data is challenging and not obvious. For example, I have reviewed Exhibit 2155, which appears to be an article from November 21, 2022, announcing Wiz’s launch of Data Security Posture Management (DSPM) on Wiz’s blog. Ex.2155 at 1. In that article, Wiz asserted that “[a]ttackers are aware of the value of sensitive data and the increasing difficulties in securing it.” *Id.* at 2. “[S]ecuring cloud data is a difficult problem” according to Wiz in this 2022 article. *Id.* “An organization with hundreds of data assets and tens of thousands of data items could have millions of

individual configurations, permissions, and lifecycle policies to manage.” *Id.* Even “[p]rotecting a database from just simple network exposure is a complex problem that shows the limitations of taking siloed approaches to understanding cloud data risks.” *Id.* These statements by Wiz, in 2022, contradict Dr. Stavrou’s assertions about the protection of sensitive data being “well known and routinely performed at the time” (i.e., in January 2019) and support that a person of ordinary skill in the art would not have combined Veselov and Hufsmith in the manner proposed by Dr. Stavrou. Ex.1002 (Stavrou Decl.) ¶104.

117. Similarly, I have reviewed Exhibit 2154, which also appears to be from November 2022 and also relates to Wiz’s launch of DSPM. That article states that “[t]he field of identifying and protecting information in the cloud is relatively new, emerging only in recent years.” Ex.2154 at 3. The article cites a research report finding that “less than 1% of companies around the world use a DSPM solution.” *Id.*; *see also id.* (“The field aims at automating currently manual processes, primarily by identifying where critical information – such as customers’ financial and personal information or commercial secrets – is stored in the cloud.”).

118. 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-Hufsmith 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-Hufsmith combination.

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

119. Even if a person of ordinary skill in the art would have made the Veselov-Hufsmith combination that Dr. Stavrou and Petitioner urge, it is my opinion that none of the claims of the '685 patent would have been obvious over that combination (or other combinations depending on this combination) at least because Veselov and Hufsmith do not disclose or suggest multiple limitations in the claims.

a. The Veselov-Hufsmith Combination Does Not Render Obvious Using Cloud Computing Platform APIs to Identify and Locate Virtual Disks (Limitations 1.2-1.4, 13.2-13.4, 22.2-22.4)²

120. Limitations 1.2-1.4 (and their corresponding limitations in claims 13 and 22) recite features relating to using “cloud computing platform APIs” to identify virtual disks of a virtual machine, as well as to query and receive the location of at least one such virtual disk. The use of the cloud computing platform

² In this declaration, I use the claim limitation labels that Dr. Stavrou and Petitioner used. A listing of the claims of the '685 patent, with those labels, is presented at the beginning of the Petition, in a “Listing of Challenged Claims.” Pet. at iv-vii.

APIs for these functions facilitates the taking or requesting the taking of snapshots and the analysis of such snapshots for vulnerabilities as recited in other limitations, discussed in subsequent sections.

121. Specifically, limitations 1.2 and 22.2 recite, “using the interface, utilize cloud computing platform APIs to identify virtual disks of a virtual machine in the client environment.” Limitation 13.2 recites “using the interface to utilize cloud computing platform APIs to identify virtual disks of a virtual machine in the client environment.” The “interface” referred to in these limitations is the “interface between a client environment and security components” recited in limitations 1.1, 13.1, and 22.1. Limitations 1.3 and 22.3 recite, “use the computing platform APIs to query a location of at least one of the identified virtual disks.” Limitation 16.3 recites “using the computing platform APIs to query a location of at least one of the identified virtual disks.” Limitations 1.4 and 22.4 recite, “receive an identification of the location of the virtual disks of the virtual machine.” Limitation 16.4 similarly recites “receiving an identification of the location of the virtual disks of the virtual machine.”

122. Dr. Stavrou and Petitioner do not rely on any disclosures from Hufsmith for these limitations (*see, e.g.*, Ex.1002 (Stavrou Decl.), ¶¶127-147), and it is my opinion that Veselov does not teach or suggest these limitations.

123. First, Veselov does not disclose “utilizing cloud computing platform APIs to identify virtual disks of a virtual machine.” Dr. Stavrou asserts that “Veselov teaches that [its] scanning service uses [a communication] interface to ‘identify virtual disks of a virtual machine in the client environment’ in order to facilitate its generation of a snapshot of such disks.” Ex.1002 (Stavrou Decl.), ¶128. But the disclosures on which Dr. Stavrou relies simply do not disclose or suggest an *identification of a virtual disk* of a virtual machine made using a cloud computing platform API, as claimed. See Ex.1002 (Stavrou Decl.), ¶128 (citing Ex.1007 (Veselov), 6:30-38, 8:60-9:41). For example, Dr. Stavrou cites the following passage in column 6, which refers to instances of a VM but says nothing about identifying a virtual disk:

“In some embodiments, a desired security assessment may be initiated by receipt of a request. For example, a user device 102 may transmit a request to the scanning service 110 for a particular security assessment of one or more instances of *a virtual machine 112*. The request may be an API call including information corresponding to the requestor, the user associated with the virtual machine 112, the on-demand storage service 108, and/or the particular virtual machine 112.”

Ex.1007 (Veselov) 6:30-38 (italics added). And the nearly column-long passage in columns 8 and 9 that Dr. Stavrou cites discusses “a virtual machine instance,” that “instance itself,” or “a logical volume attached to the corresponding virtual

machine and/or containing the instance.” Ex.1007 (Veselov) 8:60-9:41. And although Veselov refers to an identifier for a VM, it never mentions any type of identifier for, or identification of, a “logical volume,” and a VM identifier does not identify a virtual disk.

124. Additionally, as I discussed above in Section VIII.A.1, Dr. Stavrou mixes disclosures from different examples in Veselov without providing actual motivation to do so, and a person of ordinary skill in the art would not have viewed the various features and disclosures of those different examples to be simply interchangeable. Examples (Figures 1, 3A, 3B) employ “computing environments implemented by a computing resource provider 300, 350” which may use an API to provide a user with access to scanning services within virtual computing environments, but do not locate disks or analyze snapshots, and a person of ordinary skill in the art would have had no reason to add such features to those examples as I discussed previously. Ex.1007 (Veselov) 10:37-41, 14:37-41, 15:23-27, Figs. 3A-3B; *see also* Section VIII.A.1, above. And the use of APIs in these examples relate specifically to providing interfaces to a user, conveying user-specific data, or allowing users to specify information for the system, but never for identifying or locating a virtual disk—nor is there reason to use such APIs to identify or locate a virtual disk at all in these examples.

125. By contrast, the Figs. 5A-5B examples are not implemented with computing resource service providers³ or an API, but rather implement scanning services through a provisioning service such as a hypervisor. Ex.1007 (Veselov) Figs. 5A-5B, 15:56-17:66. But as I discussed above, Dr. Stavrou does not explain why a person of ordinary skill in the art would have viewed elements of those two separate examples as interchangeable and compatible or would have wanted to combine or modify them. Veselov therefore does not render these limitations obvious for the same reasons I discuss in Section VIII.A.1 relating to a person of ordinary skill in the art not mixing these examples in the manner Dr. Stavrou proposes.

126. Veselov does not disclose “utiliz[ing] cloud computing platform APIs” for identifying, querying, and receiving virtual disk location information. To the contrary, Veselov’s APIs provide interfaces to a user or otherwise convey user-specific data as input. *See generally* Ex.1007 (Veselov) 4:29-48, 5:1-8, 5:13-22, 6:34-38, 6:59-60, 8:46-51, 14:37-53, 15:23-38. For example, relating to Fig. 1, Veselov describes that “computing resource service provider 104 may provide, or otherwise be compatible with, an environment application programming

³ As stated above, I note that “computing resource service provider” is not typically used in the art to refer to a cloud computing platform.

interface (API) 120 through which a user device 102 can connect to one or more virtual computing environments of the computing resource service provider.” Ex.1007 (Veselov) 4:29-48. Additionally, “the scanning service 110 may use the API 120 to provide user interfaces to the user device 102, enabling the user to configure settings of the scanning service 110.” Ex.1007 (Veselov) 5:1-8. The user can “provide some or all of the snapshot data 146 to the scanning service” using the API 120. Ex.1007 (Veselov) 5:13-15. Dr. Stavrou ignores these descriptions of Veselov’s APIs and Veselov’s usage of the terms “environment” or “web interface” APIs to describe them, and instead attempts to broadly attribute API usage to multiple unrelated aspects of Veselov. But Veselov’s usage of “environment” and “web interface” is consistent with their limited usage in Veselov, which is how a person of ordinary skill in the art would have understood Veselov—not extremely broadly as Dr. Stavrou implies. Ex.1007 (Veselov) 4:29-48, 13:37-41.

127. Dr. Stavrou asserts that “[u]sing the interface to utilize cloud computing platform APIs to identify the virtual disks is further taught or suggested by Veselov’s discussion of tailoring snapshots.” Ex.1002 (Stavrou Decl.), ¶134. Dr. Stavrou cites two examples of what Dr. Stavrou asserts are explanations by Veselov “that the scanning service can determine what elements will be included in the snapshot.” Ex.1002 (Stavrou Decl.) ¶134 (citing Ex. 1007 (Veselov) 8:16-

45, 14:37-57). Neither of these passages disclose using cloud computing platform APIs to identify (let alone query or receive a location of) virtual disks. Rather, they disclose identifying “data elements of the virtual machine” that are needed to “perform the security assessment” (Ex.1007 (Veslov) 8:16-45) and sending “assessment information” to the “user device 302 to be displayed, selected, changed, etc.” (Ex.1007 (Veslov) 14:53-57), neither of which suggest to a person of ordinary skill in the art to modify Veslov to perform the claimed functions utilizing cloud computing platform APIs. To the contrary, both passages suggest that the user is *providing* information, such as locations, to the scanning service, which is the opposite of what is claimed in the '032 patent.

128. Dr. Stavrou also relies on a handful of references other than Veslov, including Ex.1051 (Workstation Pro), Ex.1053 (Prahlad), and Ex.1020 (Sancheti), purportedly to show that identifying virtual disks was “routine[]” or “commonplace.” Ex.1002 (Stavrou Decl.), ¶¶128, 133-134. 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 Veslov’s systems and methods to identify virtual disks of a VM utilizing cloud computing platform APIs, as claimed, as opposed to just identifying a VM generally. Veslov’s systems and methods lack such identifications, and I have been informed and understand that, in an obviousness analysis, merely collecting disparate disclosures

from a handful of references, to purportedly show that a person of ordinary skill in the art *could have* made the claimed inventions because such disclosures allegedly show that each piece of a claim was known or routine or commonplace is not enough. Rather, I understand that obviousness analyses look to what a person of ordinary skill in the art *would have* been motivated to make or do. In my opinion, nothing in the other references Dr. Stavrou cites show that a person of ordinary skill in the art would have been motivated to change Veselov's systems and methods, at all, let alone to modify those systems and methods as claimed in the '685 patent's claims.⁴ In fact, this is an example of where, in my view, Dr. Stavrou is improperly working backwards from the '685 patent's claims.

⁴ Dr. Stavrou's tertiary references do not demonstrate that the claimed features are routine or commonplace. The disclosure Dr. Stavrou cites from the Workstation Pro reference (Ex.1051) simply discloses excluding virtual disks from snapshots and makes no mention of using location data. Ex.1051 (Workstation Pro), 125. The VirtualCenter reference (Ex.1052), relied on in Ex.1002, Stavrou Decl. ¶133, shows that there may be an identification for a virtual machine location, but does not demonstrate using cloud computing APIs to identify virtual disk locations. Ex.1052 (VirtualCenter), 445-46. The disclosures Dr. Stavrou cites from Prahlad (Ex.1053) similarly do not demonstrate the use of the '685 patent's cloud

129. Dr. Stavrou also asserts that in Veselov's system and methods, "the user's request may identify a VM that will be the target of the scan, and the scanning service may then be tasked with taking a snapshot containing multiple virtual disks associated with that VM." Ex.1002 (Stavrou Decl.), ¶130. But the fact that Veselov's systems and methods work with and from the identification of VMs (*not virtual disks*) belies Dr. Stavrou's assertions about the virtual disk identification limitations. Dr. Stavrou offers no explanation for why a person of ordinary skill in the art would have changed those systems and methods to instead identify *virtual disks* of a VM, let alone identify such virtual disks using cloud computing platform APIs, as claimed. And, even if a Veselov VM can "contain[] multiple virtual disks" as Dr. Stavrou states, *see* Ex.1002 (Stavrou Decl.), ¶130; *see also* Ex.1007 (Veselov), 6:7-10, 11:28-33, 17:35-37, that still fails to show that the claimed virtual disk identification would have been obvious. As the '685 patent

computing APIs, nor is it relied upon in Dr. Stavrou's grounds. Ex.1053 (Pahlad), ¶¶81-82, 87-88, 92. Nor does Ex.1020 (Sancheti) demonstrate routineness of identifying virtual disks—in fact, its disclosure (which is not part of Dr. Stavrou's proposed grounds), supports that merely identifying a VM does not identify a virtual disk. *See* Ex.1002 ¶133 (quoting Sancheti 21:36-41). And Sancheti does not disclose identifying virtual disks for the purpose of security assessments.

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.*, '685 patent, Fig. 1B, 3:52-53 (“[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). Thus, the identification of a VM does not necessarily or obviously identify a virtual disk of that VM.

130. Dr. Stavrou further asserts that Veselov “suggests” the claimed identification of virtual disks because Veselov’s “API-based UI” can be used “to select assessment targets and other assessment parameters of [a] scan” and “remotely manage the user’s resources.” Ex.1002 (Stavrou Decl.), ¶¶135-137. As an initial matter, Veselov discusses the user identifying “the data elements of the virtual machine 112 that are needed to perform the security assessment” – not that API elements may be used for identifying virtual disk locations. Ex.1007 (Veselov), 8:16-45. Nothing in Veselov suggests that a “data element” is a virtual disk. And, in fact, Veselov suggests otherwise, because in other disclosures, virtual disk storage is referred to as a “virtual computer component.” So, the tailoring snapshots discussion does not teach that Veselov identifies virtual disks. But, in any event, none of the disclosures Dr. Stavrou cites actually teaches or suggests the claimed identifying virtual disks using cloud computing platform APIs. Dr. Stavrou cites Veselov at 2:22-3:19 (Ex.1002 (Stavrou Decl.) ¶135), but that

disclosure merely cautions that security assessments on distributed environments (such as cloud environments) could detract from the target machine's "normal workload" and explains the need to provide for security assessments that do not interfere with the activity of the target virtual machine. Ex.1007 (Veselov), 2:22-3:29. Dr. Stavrou also cites Veselov at 19:30-37 (Ex.1002 (Stavrou Decl.) ¶135), but that disclosure only addresses Veselov's definition of "data store" in that it involves accessing and retrieving data with various components. Ex.1007 (Veselov), 19:30-37. Dr. Stavrou also cites Veselov at 19:60-64 (Ex.1002 (Stavrou Decl.) ¶135), but that disclosure merely provides that single or multiple devices may perform operations. Ex.1007 (Veselov), 19:60-64. Dr. Stavrou also cites Veselov at 20:46-56 (Ex.1002 (Stavrou Decl.) ¶135), but that disclosure merely states that one embodiment of Veselov is a distributed or virtual environment. Ex.1007 (Veselov), 20:46-56. Dr. Stavrou also cites to Veselov at 4:19-48, 5:1-8, 14:37-53, and 14:53-15:29 (Ex.1002 (Stavrou Decl.) ¶¶136-137), but these disclosures merely highlight Dr. Stavrou's improper combination of distinct embodiments. That is, Veselov at 4:19-48, 5:1-8, 14:37-53, and 14:53-15:29 make clear that its APIs exist for a user to provide input to the scanning environment or to allow a scanning service, outside of the computing resource environment, to access resources within the computing resource environment. Ex.1007 (Veselov), 4:19-48, 5:1-8, 14:37-53, 14:53-15:29. Accordingly, Veselov's disclosures, taken

individually or in an unclear amalgamation (i.e., as Dr. Stavrou has done), do not teach or suggest identifying virtual disks using cloud computing platform APIs.

131. Regarding limitations 1.3, 13.3, and 22.3, Dr. Stavrou does not identify any disclosures in Veselov that actually relate to using cloud computing platform APIs to query a location of at least one virtual disk. Nor could he. Veselov only discusses queries in the context of querying a “snapshot storage service” to find a stored snapshot. Ex.1007 (Veselov), 7:34-61. Veselov does not contain any other disclosures relating to queries.

132. Dr. Stavrou asserts that “[i]t would have been apparent to a POSA that, for the scanning service to generate a snapshot of certain virtual disks... the scanning service would first determine the location of the disks it will be including in the snapshot.” Ex.1002 (Stavrou Decl.) ¶140. But Veselov only discloses querying a “snapshot storage service” to find a *stored* snapshot (Veselov 7:43-55), which obviates the need for the scanning service to query the location of a virtual disk. The claim requires using “the computing platform APIs to query a location of at least one of the identified virtual disks.” In other words, because Veselov queries the snapshot storage service for a stored snapshot, no further querying for a location of a **virtual disk** is needed to take a snapshot or is obvious and thus Veselov does not suggest using the claimed cloud computing platform API to query for the location of a virtual disk.

133. Dr. Stavrou again cites a handful of references other than Veselov, asserting that querying for locations of virtual disks was “standard,” “routine[,],” and “commonplace.” Ex.1002 (Stavrou Decl.), ¶¶140-141. But, as with limitations 1.2, 13.2, and 22.2, with limitations 1.3, 13.3, and 22.3, Dr. Stavrou does not identify anything in the references he cites that would explain why a person of ordinary skill in the art would have been motivated to change Veselov’s systems and methods to use cloud computing platform APIs to query a location of at least one virtual disk, as claimed. Veselov’s systems and methods lack such querying, and 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 to include it.⁵ In my view, this is another example of where Dr.

⁵ As I note in the preceding discussion, the VirtualCenter reference (Ex.1052) discusses properties for VMs but does not demonstrate queries for virtual disk locations. Ex.1052 (VirtualCenter), 445-46. The Virtual Disk reference (Ex.1055) makes no use of *queries* for virtual disk *locations*. Ex.1055 (Virtual Disk), 55-56. Ex.1020 (Sancheti) discloses a “management component” that “discovers and identifies the virtual disks” in a manner unlike anything in Veselov, and Dr. Stavrou does not propose to combine Sancheti with Veselov or otherwise modify Veselov according to Sancheti. Sancheti 9:36-41.

Stavrou is improperly working backwards from the '685 patent's claims, and attempting to fill gaps in Veselov by asserting limitations are "standard" or "routine" despite his references not supporting those assertions.

134. Dr. Stavrou also asserts that Veselov teaches or suggests querying a virtual disk's "more general location" because it discusses an "endpoint/virtual network address of the VM." Ex.1002 (Stavrou Decl.), ¶142. Veselov, however, merely discloses that a "virtual machine 312 may be associated with a virtual network address." Ex.1007 (Veselov), 10:48-11:4. Veselov does not mention a virtual network address anywhere else, let alone querying for such an address, or querying for a location of a virtual disk. A person of ordinary skill in the art would have known that a virtual network address associated with a VM is not the same as a location of a virtual disk, including for reasons I discussed above (e.g., a VM and a virtual disk may not be located at the same location or address). And Dr. Stavrou has not explained how or why Veselov's systems and methods would be modified to incorporate the querying of a location of a virtual disk, let alone modified to require such querying using computing platform APIs, as claimed.

135. Furthermore, for at least the same reasons Veselov does not teach or suggest the querying of limitations 1.3, 13.3, and 22.3, Veselov does not teach or suggest "receiv[ing] an identification of the location of the virtual disks of the virtual machine," as required by limitations 1.4, 13.4, and 22.4.

136. Here too, Dr. Stavrou pulls in the references other than Veselov that he cited for limitations 1.3, 13.3, and 22.3. But, as discussed above, nothing in the other references Dr. Stavrou cites shows that a person of ordinary skill in the art would have been motivated to change Veselov's systems and methods to include the claimed querying. Likewise, nothing in those references shows that a person of ordinary skill in the art would have been motivated to change Veselov's systems and methods to receive an identification of the location of a virtual disk of a VM virtual machine in response to such querying. In my view, this is yet another example of where Dr. Stavrou is improperly working backwards from the '685 patent's claims.

137. Dr. Stavrou also asserts that limitations 1.4, 13.4, and 22.4 would have been obvious because Veselov discusses a "virtual network address" of a VM, including in connection with the VM's "communication with other systems." Ex.1002 (Stavrou Decl.), ¶146. But as I explained above, a person of ordinary skill in the art would have known that a virtual network address associated with a VM is not the same as a location of a virtual disk. And Dr. Stavrou has not explained how or why Veselov's systems and methods would be modified to incorporate the querying and receiving of location information for virtual disks.

138. For at least the foregoing reasons, it is my opinion that Petitioner and Dr. Stavrou have not shown that limitations 1.2-1.4, 13.2-13.4, and 22.2-22.4 would have been obvious.

b. The Combination Does Not Render Obvious Detecting Sensitive Data (Limitations 1.6, 13.6, 22.6)

139. Limitations 1.6, 13.6, 22.6 recite, “[analyze/analyzing] the at least one snapshot to detect vulnerabilities *and sensitive data*, wherein analyzing the at least one snapshot requires no interaction with the virtual machine.”

140. Dr. Stavrou asserts that these limitations would have been obvious over Veselov and Hufsmith. Ex.1002 (Stavrou Decl.), ¶¶153-163. I disagree. It is my opinion that Veselov and Hufsmith do not teach or suggest these limitations.

141. Dr. Stavrou asserts that these limitations would have been obvious over Veselov because Veselov’s “analysis can be performed by analyzing an assessment VM instantiated based on the snapshot, or the snapshot data can be analyzed directly (i.e., without instantiating an assessment VM),” and that each type of analysis “satisfies the construction of analyzing the at least one snapshot.” Ex.1002 (Stavrou Decl.), ¶155. I disagree. As I explained above, Dr. Stavrou’s proposed claim construction for the “[analyze/analyzing] the at least one snapshot” terms in the ’685 patent’s claims is incorrect because a person of ordinary skill in the art would not understand those terms to encompass evaluations of VMs instantiated from snapshots, and Dr. Stavrou’s contrary assertion is based on only

an unsupported assumption Dr. Stavrou was “asked” to make—nothing more. *See* Section VI.B.1 (“Claim Construction”).

142. Dr. Stavrou contends that Veselov in combination with Hufsmith “teaches performing [the] snapshot analysis to detect ‘sensitive data.’” Ex.1002 (Stavrou Decl.), ¶159. I disagree. First, as discussed above, Hufsmith is devoid of any teaching of a snapshot analysis, and does not provide any indication that it detects sensitive data as is required by the challenged claims. *See* Section VIII.A.1. Dr. Stavrou cites to multiple disparate passages of Hufsmith (Ex.1002 (Stavrou Decl.), ¶160 (citing Ex.1078 (Hufsmith), ¶¶82, 105, 228, claim 9) that discuss various scan configurations (i.e., static, dynamic), weight determinations for context properties of container images, and layer assessments for the container images. *See* Ex.1078 (Hufsmith) ¶¶82, 105, 228, claim 9; *see also* Ex.1078 (Hufsmith), ¶¶8, 33, 84, 118, 129, 156, 160, 181 (discussing CWE and CVE vulnerabilities, “software weakness types,” developer annotations, and other plugins to detect vulnerabilities in code development). Although paragraph 118 refers to “restricted information,” that disclosure is in relation to “[i]mpact metrics” such as a “confidentially impact” that only provides an estimate of *potential exposure* of “restricted information” (i.e., “no loss,” “some” loss, or “total loss”). *See* Ex.1002 (Stavrou Decl.) ¶160 (discussing Hufsmith’s alleged detection of “potential exposure” of sensitive data). Such metrics relate to a detected

vulnerability, rather than the detection of vulnerabilities **and** sensitive data, which is what is claimed. *See* Ex.1078 (Hufsmith) ¶116. But even accepting Dr. Stavrou's characterization of Hufsmith, detecting that vulnerabilities potentially expose sensitive data, where that sensitive data may not even exist, does not teach or suggest detecting sensitive data as claimed.

143. Notably, the '685 patent distinguishes detecting vulnerabilities and detecting sensitive data. For example, the '685 patent discloses, in column 4, detection of "certain known vulnerabilities identified in, for example, a CVE database." Ex.1001 ('685 patent) 4:40-42. Separately, the '685 patent discloses "detect[ing] cyber threats that do not represent vulnerabilities" including "detect[ing] and alert[ing] on sensitive data not being encrypted on the logical disk, private keys found on the disks, [and] system credentials stored clearly on the disk." Ex.1001 ('685 patent) 6:19-30; *see also id.* 6:46-53 ("[D]etected cyber threats related to sensitive data (including personally identifiable information, PII) is reported at a higher priority.... [S]uch data is determined by searching for the PII, analyzing the application logs to determine whether the machine accessed PII/PII containing servers, or whether the logs themselves contain PII...."). Hufsmith does not teach or suggest anything similar to the '685 patent's detection of sensitive data, as claimed.

144. Dr. Stavrou also states that “the detection of unsecured sensitive data was a common best practice enforced by security assessments performing static and runtime/dynamic analysis.” Ex.1002 (Stavrou Decl.) ¶160. In support, Dr. Stavrou cites Ex.1036 (Wysopal) and Ex.1056 (Chari). For example, Dr. Stavrou characterizes column 12 of Wysopal as disclosing “detecting personally identifiable information in virtual machine image.” Ex.1002 (Stavrou Decl.) ¶160. But what the cited passage of Wysopal actually discloses is that “a virtual machine that includes an *application* that processes healthcare data including personally identifiable information may accord a rating of ‘serious’ to factors such as damage to reputation, liability, unauthorized release of sensitive information and criminal liability, but ‘n/a’ for personal safety.” Ex.1036 (Wysopal) 4:40-45. Thus, like Hufsmith, Wysopal at best assigns a rating after determining *potential exposure* of sensitive data (e.g., the application has the potential to expose sensitive data, but nothing indicates that the detection of sensitive data has taken place or would have been obvious).

145. Regarding Chari, Dr. Stavrou separately includes Chari in the combination in Ground 2. I discuss Ground 2 in Section VIII.B, below. His reliance on Chari as a tertiary reference in Ground 1 is misplaced for the same reasons discussed below for Ground 2.

146. For these reasons, Dr. Stavrou fails to show that limitations 1.6, 13.6, and 22.6 would have been obvious over Veselov and Hufsmith.

c. The Combination Does Not Render Obvious Reporting Detected Vulnerabilities As Alerts (Limitations 1.8, 13.8, 22.8)

147. Limitations 1.8, 13.8, and 22.8 recite, “report the detected vulnerabilities and sensitive data as alerts, wherein the alerts are filtered and prioritized based on the determined risk level of the virtual machine.”

148. First, as I discussed in the previous section, the combination does not render obvious detecting sensitive data. For the same reasons, the combination does not render obvious reporting the detected sensitive data as alerts.

149. Additionally, Veselov does not mention alerts, let alone filtered and prioritized alerts. Veselov only discloses that its “scanning service 110 may be configured to perform security assessments and produce assessment results based at least in part on snapshot data.” Ex.1007 (Veselov), 5:9-13. Veselov also provides that the scanning service may “process the snapshot data” and “generate one or more assessment results.” Ex.1007 (Veselov), 8:19-21, 18:53-56 (similar).

150. Dr. Stavrou, however, incorrectly equates communicating “assessment results” with “reporting... vulnerabilities as alerts.” Ex.1002 (Stavrou Decl.), ¶169. But a person of ordinary skill in the art would have understood that communicating results does not require or suggest reporting detected

vulnerabilities “as alerts,” as claimed. *See, e.g.,* Ex.1001 (’685 patent), 7:50-51 (“[T]he detected cyber threats and/or vulnerabilities are reported, *for example, as alerts*”) (italics added), cls. 1, 13, 22 (“report the detected vulnerabilities and sensitive data *as alerts...*”). Veselov’s disclosure of simply “communicating” results is not enough to constitute an alert. For example, Veselov’s “assessment results” include “findings” resulting from “applying the rules of the assessment” that do not suggest the urgency implied by an alert. Ex.1007 (Veselov) 10:18-22. Rather, simply “communicating” results in Veselov suggests a standard report. Persons of ordinary skill in the art would have understood that alerts and reports are different and can involve different system components, capabilities, and purposes. *Compare* Ex.2010 [U.S. Pat. No. 9,467,465 (“Hibbert”)], 20:34-44 (“The **alert** module 516 may **generate an alert** based on the assessment. If one or more applications or files identified within the records of the assessment request are classified as vulnerable or highly suspicious, the alert module 516 may generate an alert. The alert may flag the vulnerability, potentially identify the exploit, indicate the degree of danger, or provide any other information. The alert may be provided in any manner. For example, the alert module 516 may provide one or more alerts via email, SMS text message, MMS, web page, intranet alert, extranet alert, or any other way.”) *with id.*, 20:12-21 (“The **report** module 514 may **generate any kind of report** indicating the results of one or more assessments. Information

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

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

151. Alerts serve to notify stakeholders of situations that require attention, and not all reported results are alerts. Dr. Stavrou agreed in his deposition that “an alert is different than a report.” Ex.2077 (Stavrou Tr.) 286:9-287:3. Dr. Stavrou’s opinions about Veselov for these limitations, however, contradict this acknowledgment and how a person of ordinary skill in the art would understand limitations 1.8, 13.8, and 22.8, as those limitations do not simply recite “report[ing] the detected vulnerabilities.” Ex.1002 (Stavrou Decl.) ¶168. To the contrary, they recite that the detected vulnerabilities be reported “*as alerts.*” ’685 patent, cls. 1, 13, 22. Additionally, Dr. Stavrou provides no indication as to why a person of ordinary skill in the art would modify Veselov’s techniques to report its assessments results as alerts as limitations 1.8, 13.8, and 22.8 require.

152. Dr. Stavrou also relies on Hufsmith as teaching alerts, citing to Hufsmith’s “annotations in a development environment,” i.e., displaying information about a vulnerability in order to show a developer particular lines of source code in a pre-deployment container image undergoing development that should be addressed. Ex.1002 (Stavrou Decl.) ¶170 (citing Hufsmith ¶¶153-154, 156, 180-182). However, a person of ordinary skill in the art would have understood this context to be inapplicable to Veselov, which only analyzes deployed VMs or snapshots where such source code would be unavailable for such annotating. Dr. Stavrou does not explain how or why persons of ordinary skill in

the art could have or would have implemented this feature in Veselov's systems. Although he asserts that this disclosure "also appl[ies] to the other embodiments," Dr. Stavrou explicitly relies on "annotations" of the source code that persons of ordinary skill in the art would understand *not* to apply to situations, such as in Veselov, where *source code is not available to be annotated*. Ex.1002 (Stavrou Decl.) ¶¶169-170. And without the source code available, Hufsmith's stated objectives of annotating the source code—e.g., to increase awareness "earlier in the workflow of container development" and identify vulnerabilities "in a more real time manner" (Hufsmith ¶160)—do not apply. A person of ordinary skill in the art would not have had a reasonable expectation of success in combining these cited portions of Hufsmith with Veselov as a result. In addition, a person of ordinary skill in the art would not have had a reasonable expectation of success in combining these cited portions of Hufsmith with Veselov because the targets of Veselov's security assessments are expressly designed to be "read-only" and thus the targets of Veselov's security assessments are not in development, and therefore, they could not be modified; accordingly, Hufsmith's annotations in a development environment would be of no use. Ex.1007 (Veselov) 18:1-3 ("The scanning service 510 may mount volume images as read only to avoid inadvertently modifying the volume corresponding to a volume image."), 18:26-29 ("read-only virtual hard disk"), Fig. 6 ("read-only file system"); *see also id.* 18:38-39 ("volume

image can be read”), 17:5-8 (scanning service can “read the copied files”), 17:23-29 (“read files”). Even so, Dr. Stavrou does not explain how such annotations of source code (see Hufsmith’s Figure 3B, below) would be included in “reporting the vulnerabilities and sensitive data *as alerts*” that “are filtered and prioritized based on the determined risk level of the virtual machine.” *Id.* Simply making it “available to the user even when a partial report is presented” does not constitute an alert as discussed above. Rather, because the pre-deployment context in which Hufsmith’s alleged “alerts” are displayed, where such annotations can trigger a change to source code, is entirely different than the read-only execution environment of Veselov, a person of ordinary skill in the art would not have expected success or have combined Veselov and Hufsmith as Dr. Stavrou alleges. Although Dr. Stavrou asserts that Hufsmith’s “teachings broadly apply to other approaches described elsewhere in Hufsmith” (Ex.1002 (Stavrou Decl.) ¶170), Hufsmith discloses that the techniques may be used “in conjunction with the approaches above” (Ex.1078 (Hufsmith) ¶¶153-154), which does not teach or suggest using these approaches in any setting where source code is not available such as in Veselov. And Dr. Stavrou does not explain how or why a person of ordinary skill in the art would have applied these or any other teaching of Hufsmith to Veselov’s system.




FIG. 3B

Hufsmith, Figure 3B

153. Dr. Stavrou further asserts that “Hufsmith teaches prioritizing the alerts based on the global weights.” Ex.1002 (Stavrou Decl.) ¶171. I disagree. Hufsmith’s weights are used to calculate a weighted sum of vulnerabilities and therefore are used to calculate a type of “count of the number of potential vulnerabilities detected.” Ex.1078 (Hufsmith) ¶90. Further, the weights are not reported as alerts, nor are they “filtered and prioritized” as claimed. In paragraph 120, cited by Dr. Stavrou, the weights relate to metrics and values, not to a detected vulnerability itself. See Ex.1078 (Hufsmith) ¶112 (“the configuration evaluator 105 determines a weight for each metric or value of a metric based on whether the

configuration properties implicate that the container is more, less, or neutrally susceptible to that metric or specific value associated with that metric.”).

154. Similarly, Dr. Stavrou cites to Ex.1063 (Hibbert) for support without explanation. Ex.1002 (Stavrou Decl.) ¶171 (quoting Ex.1063 (Hibbert) ¶187). The quoted passage says nothing about filtering and prioritizing alerts “based on the determined risk level of the virtual machine.” *See id.* Dr. Stavrou does not explain why a person of ordinary skill in the art would have interpreted Hufsmith in view of Hibbert, which appears to be unrelated to Hufsmith. Hufsmith focuses on calculating a “Combined Threat Score for Container Images,” particularly for distributed applications (Ex.1078 (Hufsmith), Title, Abstract), and Hibbert focuses on agent-based risk assessments of applications on a device (Ex.1063 (Hibbert) Abstract) and is entirely silent about container images for distributed applications or even virtual machines. Nor does Dr. Stavrou assert that a person of ordinary skill in the art would have made any modifications to Veselov or Hufsmith based on Hibbert. Further, Dr. Stavrou’s reliance on Basavapatna and Hibbert (Ex.1002 (Stavrou Decl.) ¶173) does not cure the deficiencies of Hufsmith. Dr. Stavrou cites paragraphs 138-140 of Basavapatna for his assertion that it was “routine in the art to prioritize vulnerabilities and their associated alerts based on a determined risk level....” Ex.1002 (Stavrou Decl.) ¶173. But the cited disclosure of Basavapatna merely discloses that “the user should be alerted” when an “aggregate risk metric

for any asset rises above a specific threshold.” Ex.1002 (Stavrou Decl.) ¶173 (quoting Ex.1008 (Basavapatna) ¶¶138-140). In other words, the alerts are *not* prioritized at all, but rather all alerts arise from a metric surpassing the same “specific threshold.” Hibbert teaches the exact same approach, which differs from Dr. Stavrou’s characterization and from the claimed limitation. *See* Ex.1002 (Stavrou Decl.) ¶173 (quoting Ex.1063 (Hibbert) ¶¶13, 189) (“generate an alert... based on the relevant risk threshold (e.g., determined that a risk value exceeds a relevant risk threshold).”).

d. Dependent Claims 2 and 14

155. Claim 2 depends from claim 1 and recites “wherein during the analysis of the at least one snapshot, the virtual machine is inactive wherein the snapshot includes a page file of memory associated with the virtual machine.” Claim 14 depends from claim 13 and recites the same additional limitation. Dr. Stavrou asserts that Veselov teaches this limitation. Ex.1002 (Stavrou Decl.) ¶177.

156. Dr. Stavrou has not shown that Veselov discloses a virtual machine that is “inactive” during analysis of a snapshot. Dr. Stavrou asserts that “most of Veselov’s snapshot-based assessment techniques require no interaction with the original VM.” Ex.1002 (Stavrou Decl.) ¶178. But the explicitly stated goal of Veselov’s approach of minimizing interaction with the VMs is to *keep them active*. *See, e.g.,* Ex.1007 (Veselov) 2:42-46 (“The present systems and methods overcome

drawbacks of existing security assessment systems, such as... unavailability of resources services that are part of the assessment....”). Dr. Stavrou’s conclusion that “a POSA would have expected many target VMs to already be inactive at the time of the security assessment” is wholly unsupported by, and contradictory to, Veselov’s explicit teachings.⁷ *See, e.g.*, Ex.1007 (Veselov) 12:21-27 (security assessments performed on different virtual machine to allow the “customer’s virtual machine” to “servic[e] requests from its customers or perform[] other useful work”). Dr. Stavrou also admitted in his deposition that Veselov does not disclose a VM being offline when the snapshot is taken. *See* Ex.2077 (Stavrou Tr.) 203:21-205:7; *see also id.* 263:16-264:10 (“If [a VM] is at rest, it wouldn’t be able to run its normal workload, that’s true.”). Thus, it is my opinion that Veselov does not teach this limitation.

⁷ Dr. Stavrou’s citation to Lutas and Thakar do not change the fact that offline VMs are contradictory to Veselov’s teachings. For example, Dr. Stavrou does not suggest that Lutas would have motivated a person of skill in the art to modify Veselov to operate contrary to its purpose and take Veselov’s virtual machines offline. The quoted portion of Thakar merely discusses analyzing an “echo” instead of a target machine, and does not suggest taking Veselov’s VMs offline.

e. Dependent Claims 7 and 19

157. Claim 7 depends from claim 1 and recites “wherein the identification of the location of the virtual disks of the virtual machine includes a virtual address of at least one of the virtual disks.” Claim 19 depends from claim 13 and recites the same limitation.

158. Dr. Stavrou asserts that Veselov “teaches this limitation in several ways.” Ex.1002 (Stavrou Decl.) ¶201. I disagree. Dr. Stavrou relies on Veselov’s disclosure of a virtual network address of a “virtual machine.” Ex.1002 (Stavrou Decl.) ¶202 (quoting Ex.1007 (Veselov) 10:48-11:4). But similar to the above discussion relating to limitations 1.2-1.4, obtaining a virtual address of a virtual *machine* does not teach or suggest identifying a virtual address of a virtual *disk*. Thus, Dr. Stavrou’s statement that, because “a virtual network address of the VM to which the virtual disks are allocated” would be “a virtual address that is associated with the virtual disks (since the virtual disks are allocated to that VM)” is unsupported and incorrect, and does not teach or suggest the limitations of claims 7 and 19.

159. Dr. Stavrou acknowledges that “Veselov does not expressly state that the scanning service receives an identification of a virtual disk’s ‘virtual address.’” Ex.1002 (Stavrou Decl.) ¶203. I agree—but Veselov also does not inherently teach, nor does it suggest at all, identification of a virtual address of a virtual disk.

Dr. Stavrou falls back on an assertion that identification of such an address would have been “routine... in response to such a query”—but Veselov does not even teach querying the location of a virtual disk as discussed above for claim 1. Accordingly, it is my opinion that Veselov does not teach the identification of the location of the virtual disks of the virtual machine includes a virtual address of at least one of the virtual disks.

f. Summary of Ground 1 Opinions

160. For at least the foregoing reasons, it is my opinion that Petitioner and Dr. Stavrou have not shown that any of the independent claims of the '685 patent—claims 1, 13, and 22—as well as dependent claims 2, 7, 14, and 19, would have been obvious. 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 3-6, 8-12, 15-18, 20-21), would have been obvious. Thus, Petitioner and Dr. Stavrou have not shown that any of claims 1-22 of the '685 patent would have been obvious.

B. Ground 2: Claims 1-22 Would Not Have Been Obvious Over Veselov, Hufsmith, and Chari

161. Dr. Stavrou sets forth Ground 2 to build upon Ground 1 and Chari is applied only for its alleged disclosure of “detecting and remediating sensitive data.” Ex.1002 (Stavrou Decl.), ¶¶224-237. But Dr. Stavrou challenges claims 1-22 largely based on only Veselov and Hufsmith. Therefore, I disagree with Dr.


Stavrou's analysis for Ground 2 for at least the same reasons as Ground 1. In particular, I disagree that Veselov and Hufsmith (with or without Chari) discloses the limitations that I identified as missing above for Ground 1.

162. In addition, it is my opinion that a person of ordinary skill in the art would not have been motivated to combine Chari with Veselov and Hufsmith. As an initial matter, a person of ordinary skill in the art would not have combined Veselov and Hufsmith for the reasons described above for Ground 1. Additionally, a person of ordinary skill in the art would not have been motivated to add Chari to the combination. Chari identifies sensitive data by directly interacting with the virtual machine. *See, e.g.*, Ex.1056 (Chari) ¶¶37 (“Virtual machine image manager 218 inserts labeler 222 and sanitizer 224 in each virtual machine image.”), 39 (“labeler 222” uses scripts to “identify and label sensitive data contained within a virtual machine”). Chari's approach to identifying sensitive data is contrary to the proposed combination, which relies on Veselov's lack of interaction with the original VM, and is contrary to the claims, which recite that “analyzing the at least one snapshot [to detect sensitive data] **requires no interaction with the virtual machine.**” *See* Ex.1002 (Stavrou Decl.) ¶244 (emphasis added). Chari's interaction with the virtual machine to identify sensitive data is the opposite of what is claimed in the '685 Patent. A person of ordinary skill in the art also would not have been motivated to combine Chari with Veselov and Hufsmith to analyze

snapshots to detect sensitive data because Veselov counsels *against* “complexity of installation and monitoring” of an assessment system. Ex.1007 (Veselov), 2:42-46. A person of ordinary skill in the art would not have had a reasonable expectation of success in combining Chari with Veselov and Hufsmith in the proposed manner at least because, Chari does not provide the person of ordinary skill in the art with the means for identifying sensitive data without interacting with the VM as required. A person of ordinary skill in the art also would not have been motivated to or have reasonably expected to succeed in combining Chari with Veselov. Veselov analyzes volume images and snapshots when they are “read only to avoid inadvertently modifying the volume corresponding to a volume image.” Ex.1007 (Veselov), 18:1-3, 18:26-29 (“read-only virtual hard disk”), Fig. 6 (“read-only file system”); *see also id.* 18:38-39 (“volume image can be read”), 17:5-8 (scanning service can “read the copied files”), 17:23-29 (“read files”). Chari, in contrast, explicitly adds “labels” and “sanitize[s]” sensitive data, and would thus require modifying data, in a manner that is contrary to and incompatible with Veselov.

163. Dr. Stavrou asserts that Chari is “[s]imilar to Veselov” in that Chari’s analysis “may be performed on VM images/snapshots or on instantiated VMs.” Ex.1002 (Stavrou Decl.) ¶230. But as I discussed previously, the claims of the ’685 patent require analyzing snapshots—not instantiated VMs. *See* Section

VI.B.1, VIII.A.2.b. Further, Chari discloses generating a snapshot “for rollback”—in other words, to be used to instantiate the VM at an earlier state. Ex.1056 (Chari) ¶73; *see also id.* ¶94 (snapshots are generated “for rollback purposes in case the specific instance of the virtual machine image fails or becomes corrupted”). Chari discloses that “the virtual machine image manager labels and sanitizes the snapshot of the specific instance of the virtual machine image” but is silent about analyzing the snapshot to detect sensitive data. To the contrary, in this scenario the system is simply repeating what had already been scanned, labeled, and identified in the “virtual machine instance” in an earlier step of the process—not detecting sensitive data by analyzing the snapshot. *See* Ex.1056 (Chari) ¶¶62, 92-96, Fig. 9 (reproduced below). Chari never discloses analyzing snapshots to detect sensitive data, and a person of ordinary skill in the art would not have been motivated, nor have had a reasonable expectation of success, in trying to do so based on Chari and Veselov.



Ex.1056 (Chari) Fig. 9

C. Ground 3: Claims 2-3 and 14-15 Would Not Have Been Obvious Over Veselov, Hufsmith, Chari, Price, and Huseinović

164. Dr. Stavrou sets forth Ground 3 to address only the additional limitations of claims 2-3 and 14-15. Ex.1002 (Stavrou Decl.), ¶¶238-253. But Dr.

Stavrou challenges claims 1 and 13, on which claims 2-3 and 14-15 depend, based on only Veselov and Hufsmith. *See* Ex.1002 (Stavrou Decl.), ¶238 (Ground 3 “builds upon Ground 2”). Accordingly, for at least the same reasons Dr. Stavrou fails to demonstrate a reasonable likelihood of prevailing against claims 1 and 13, it is also my opinion that Dr. Stavrou fails to demonstrate a reasonable likelihood of prevailing against claims 2-3 and 14-15. In particular, I disagree that a person of ordinary skill in the art would have combined Veselov/Hufsmith/Chari with or without Price and Huseinović as I discussed above for Grounds 1 and 2.

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

165. Additionally, a person of ordinary skill in the art would not have been motivated to combine Price or Huseinović, let alone both, to the Veselov/Hufsmith/Chari combination, to arrive at the challenged claims. Price performs security assessments differently depending on if a VM is powered-down or active. Specifically, Price collects VM images from powered-down VMs, but for active VMs Price installs agents to monitor them. *See, e.g.*, Price (Ex.1048), Fig. 6. Rendering Veselov’s VMs inactive goes expressly against Veselov’s goal of implementing security assessments without “reduc[ing] the target computing device’s capabilities” (Veselov, 2:56-67) as I discussed above, and thus, a person of ordinary skill in the art would not have combined Price with Veselov. *See*

Section VII.A, VII.C. Huseinović, for its part, discloses techniques for collecting VM memory data, and counsels *against* powering down VMs. *See* Huseinović (Ex.1049), 940 (“If the computer is set to off, the data are lost and cannot be retrieved for later analysis.”). These teachings are likewise incompatible with teachings of Price and the recitations of dependent claims 2 and 14 of the ’685 Patent. Thus, it is my opinion that a person of ordinary skill in the art would not have been motivated to make, and would not have reasonably expected to succeed in making, the proposed combination of five references to reconstruct claims 2-3 and 14-15 of the ’685 patent.

2. Even If Combined, Veselov, Hufsmith, Chari, Price, and Huseinović Do Not Disclose or Suggest Multiple Limitations in Claims 2-3 and 14-15

166. Dr. Stavrou asserts that these claims are obvious based on Price and Huseinović “where the VM is powered off at the outset.” Ex.1002 ¶245. I disagree. Price does not disclose or suggest “wherein during the analysis of the at least one snapshot, the virtual machine is inactive” because Price is directed to assessments with VM *images*, not *snapshots*, as I discussed above. *See* Section VII.C. Price does not mention snapshots, which present a different class of security concerns from VM images. *See* Section VIII.A.1.b. Nor does Dr. Stavrou provide any evidence that snapshots could be utilized for analyses of inactive VMs in lieu of Price’s VM images.

167. Dr. Stavrou also asserts that Huseinović teaches “wherein the snapshot includes a page file of memory associated with the virtual machine.” Ex.1002 (Stavrou Decl.) ¶248. According to Dr. Stavrou, “Huseinović teaches that, in a common virtualization platform... a snapshot taken while a VM is suspended/hibernating typically includes a page file representing the state of machine memory.” Ex.1002 (Stavrou Decl.) ¶248. But preserving this “vmem file” is an alternative to a snapshot. Huseinović describes that “[b]y using VMware workstation it is possible to choose two natively supported options to save a physical memory State and data.” Ex.1049 (Huseinović), 940. “The first one is preserving appropriate *vmem* file (by using suspend option).” *Id.* “[T]he second option is to take a virtual machine snapshot.” *Id.* In other words, the virtual machine snapshot is a different option than the suspend option to preserve the appropriate vmem file. Dr. Stavrou does not explain why a person of ordinary skill in the art would have implemented any of the combined systems in a “VMware workstation,” or why it would have combined different alternative approaches.

168. Thus, it is my opinion that Dr. Stavrou has not shown claims 2-3 and 14-15 to have been obvious over the five references asserted in the Ground 3 combination.

IX. Objective Indicia of Non-obviousness

169. I have considered objective evidence that supports that a person of ordinary skill in the art would not have considered the claims of the '685 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 '685 Patent Inventions

170. 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 '685 patent. In that section of his declaration, Mr. Shua maps Orca's SideScanning™ to the claims of the '685 patent. I have further reviewed the Exhibits cited in that Section of Exhibit 2074. Mr. Shua's mapping is consistent with my review of those Exhibits. Thus, I adopt and incorporate by reference Section III of Exhibit 2074. Based on my review of those Exhibits and Mr. Shua's mapping, it is my opinion that Orca's SideScanning™ practices the claims of the '685 patent.

B. Long-Felt, Unresolved Need

171. 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 ’685 patent describes, “[o]rganizations [] increasingly adapted their applications to be run from multiple cloud computing platforms.” ’685 patent, 1:20-21. Virtualization technologies such as virtual machines, software containers, and serverless functions helped enable this shift to commercial cloud computing. *See* ’685 patent, 1:36-39.

172. As applications moved to the cloud, security lagged behind. “In the early day of cloud, [security] vendors rushed to deploy their existing solutions into this new environment.” Ex.2076 at 4. But even as vendors realized that cloud computing posed new security challenges, new solutions had a number of drawbacks. *See* Ex.2136, Cloud Security Report, 2 (2018) (“As more workloads are shifting to the cloud, cybersecurity professionals remain concerned about security of data, systems, and services in the cloud.”); Ex.1029, *Overview of Attacks on Cloud Computing* (April 2012) at 1 (“[S]ecurity is the biggest problem of” cloud computing systems). As the ’685 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. ’685 patent, 1:44-2:14.

173. 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. ’685 patent, 1:49-2:6. The problem of limited coverage, which was especially true with agent-based systems, includes difficulties associated with installing agents on every VM. Thus, while an agent may have better visibility into a VM’s security posture, the difficulty in installing agents throughout a platform left big gaps in coverage. *See* Ex.2076, 11 (“Traditional agents struggle with complex deployment and maintenance.”); ’685 patent, 2:13-14 (“[I]nstalling agents in a large data center may take months.”); Ex.2103, *Security Agents Don’t Belong In Your Cloud* (Aug. 18, 2021) at 3 (“Customers tell me that agent-based approaches are only able to cover up to 50% of the cloud real estate even after several months of deployment.”).

174.

REDACTED

REDACTED

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

176. 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.”).

177. 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 and sensitive data. ’685 patent, 1:46-52, 1:57-59, 2:1-6, 2:15-16.

178. Orca’s claimed inventions solved this need. As discussed in the following sections, Orca’s claimed approach, embodied in its SideScanning™ product, was a “genuinely genius innovation” such that cloud security posture management “can be defined by pre and post-Orca.” Ex.2076 at 7. I have also reviewed articles by Wiz and Palo Alto Networks that describe agentless scanning, and its benefits, in a manner that closely matches the claimed invention. For example, I have reviewed Exhibit 2138, which appears to be a Wiz article from September 28, 2021, titled “Agents Are Not Enough.” Ex.2138 at 1. In this article, Wiz describes: “[A]gentless deep scanning approaches [] achieve a level of

comprehensiveness in terms of coverage and features that is not possible with agent-based approaches due to the tradeoff between performance and impact to the environment.” Additionally, I have reviewed Exhibit 2135, which appears to be a Palo Alto Networks article from November 16, 2021, titled “Palo Alto Networks Shifts Left with Prisma Cloud 3.0.” Ex.2135 at 1. In this article, Palo Alto Networks describes: “agentless scanning... works off cloud provider APIs and disk snapshots” to “extend[] coverage to both active *and dormant* virtual machines”. Ex.2135 at 3.

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

Orca Security made a *big leap* in cloud security with the introduction of its patented agentless SideScanning technique in 2019. This innovation introduced workload visibility without needing security teams to install an agent on endpoints or installing software on workloads for assessments. They operate out-of-band by accessing cloud workloads at the runtime storage level, gathering crucial data through the cloud provider’s APIs, while scanning disk snapshots.⁸

Ex.2076 at 31.

⁸ See ’685 patent, cls. 1, 13, 22 (limitations 1.1-1.6, 13.1-13.6, 22.1-22.6).

180.

REDACTED

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

181. As discussed below, I have reviewed objective evidence that there was industry skepticism of Orca's claimed inventions, but that they were quickly

⁹ See '685 patent, cls. 1, 13, 22 (limitations 1.7, 13.7, 22.7).

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.

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

183. After Orca “introduced the concept of agentless side scanning,” it was “difficult” for “older providers to come to terms with [its] success.” Ex.2076 at 7. That 2024 report states that these providers were “skeptics” because “from a security perspective, agentless scanning was nothing compared to what attacks they were able to detect and protect.” *Id.* “They missed however that security teams were not yet in the weeds of these architectures,” so before Orca, they “didn’t know what they didn’t know.” *Id.* 7. And in 2021, Palo Alto Networks released an article about how Orca’s solutions “fall short.” Ex.2106, *How Agentless Approaches Sideline Cloud Workload Protection Strategies*, Palo Alto Networks (Feb. 23, 2021) at 1 (“Agentless Solutions Fall Short”). By November, Palo Alto Networks announced their own agentless scanning solution. Ex.2135, *Palo Alto*

Networks Shifts Left with Prisma Cloud 3.0 at 2 (announcing “Agentless Security for Flexible Cloud Workload Protection” that “Provide[s] visibility into an organization’s cloud workload and application risks”).

184.

REDACTED

REDACTED

REDACTED

185. Eventually, “everyone needed [agentless scanning] in order to stay competitive.” Ex.2076, 7; Ex.2103, *Security Agents Don’t Belong In Your Cloud* (Aug. 18, 2021) (“[A] new ‘agent-less’ approach to cloud security has recently become possible due to cloud-native APIs now available from CSPs. This approach provides security in a fundamentally different way.”). As I mentioned above, Palo Alto Networks quickly reversed course on its public criticisms of the approach, and Wiz entered the market a year after Orca with the same technology—copied from Orca.

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

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

- CRN's "20 Coolest Tech Startups of 2020." Ex.2119, *The 20 Coolest Tech Startups of 2020*, CRN (Nov. 19, 2020) at 9-10. CRN touted "Orca's SideScanning technology" which "automatically assesses the security state of every discovered asset" to provide "complete visibility." *Id.*
- Forbes's "Top 20 Cybersecurity Startups to Watch in 2021." Ex.2120, *The Top 20 Cybersecurity Startups to Watch in 2021 Based on Crunchbase*, Forbes (Nov. 20, 2020) at 6-7. Forbes found Orca "noteworthy for its innovative approach to providing instant-on, workload-deep security" for cloud computing systems, emphasizing Orca's "patent-pending SideScanningTM technology" that "reads cloud configuration and workloads' runtime block storage out-of-band." *Id.*
- SC Media's Best Cloud Workload Protection Solution. Ex.2123, *SC Award Winners 2023 Orca Security – Best Cloud Workload Protection Solution*, SC Media (Aug. 21, 2023) ("After years of dealing with agents, then seeing how easy it

is to install and use Orca, I knew that its agentless approach was both a major innovation and a game changer”)

187.

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

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

- Exhibit 2076, which I have been informed is a 2024 report published on Substack by Software Analyst Cyber Research (founded by Francis Odum) and James Berthoty that heaped praise on Orca's inventions specific to the claimed combination of features of the '685 patent, such as using cloud computing platform APIs to determine locations of virtual disks and virtual disk snapshots and facilitate snapshot analyses. Select quotes include:

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

Id. at 7

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

***Id.* at 31**

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

***Id.* at 31**

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

***Id.* at 31**

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

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

- Exhibit 2118, which is titled “Vendor Landscape and TCO Analysis for Achieving Cloud-Wide, Workload-Deep Security and Compliance for Public Cloud Environments.” Ex.2118 at 1. This document describes “Orca uses a patent-pending technique called SideScanning™, ... which unobtrusively scans the run-time block storage layer—out of band—to extract valuable data about the environment and the assets communicating in it.” Ex.2118 at 5. The report states that a Senior Cloud Security Engineer from a company called Sisense described Orca’s solution as “unique in that it locates vulnerabilities with precision and delivers tangible, actionable results—without having to sift through all the noise.” *Id.* at 4.

190. In summary, Orca’s extensive praise for the claimed combination of features in the ’685 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™

191. 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 ’685 patent’s claims obvious.

192.

REDACTED

REDACTED

195. REDACTED

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

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

197. 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.

198. REDACTED

199. REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

218. 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™

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

220. REDACTED

221. REDACTED

REDACTED

REDACTED

REDACTED

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

- Exhibit 2026, which I am informed is a webpage on Wiz's website. This webpage touts the benefits of agentless scanning and states that agentless analysis is "based on two fundamental principles" including "Privileged access to customer's cloud environment *via APIs*" and "*Snapshot scanning*." Ex.2026 at 3. "By connecting to cloud APIs, they can discover new resources as they're created, without requiring manual installation of an agent process. This maximizes security coverage from day one, improving the visibility of security issues." Ex.2025 at 8.

What is agentless security?

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

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

Agentless analysis approaches are based on two fundamental principles:

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

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

Ex.2026 at 3

Advantages of agentless security

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

Simple, automatic coverage

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

Ex.2026 at 8

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

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

Ex.2115 at 3

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

Ex.2115 at 1

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED

Exhibit 2155 shows that Orca added “Data Security Posture Management” as a feature that “now detects [sensitive] data such as PII” in November 2022. Ex.2155 at 1, 3. After it copied Orca, Wiz touted the importance of detecting sensitive data on its website. *See* Ex.2156 at 5 (“DSPM prioritizes data security by identifying data-targeted vulnerabilities and enforcing security policies. This is essential for organizations with large amounts of sensitive data and those in regulated industries.”); Ex.2157 at 5 (“With integrated DSPM, Wiz takes a load off your mind with a streamlined approach to data security. It finds and classifies data across all your cloud environments.”).

230.

REDACTED

REDACTED

REDACTED

232. In the discussion that follows, I discuss additional exhibits I have reviewed that further supports the substantial similarity between Wiz and Orca's products as it relates to specific limitations of the claims of the '685 patent. ^{REDACTED}

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

13.1, 22.1: Establish[ing] An Interface Between A Client Environment And Security Components

233. Wiz's Workload Scanner is substantially similar to Orca's SideScanningTM in relation to limitations 1.1, 13.1, and 22.1 of claims 1, 13, and 22 of the '685 Patent (reproduced in full above).

234. For example, similar to Orca's SideScanningTM, Wiz's Workload Scanner is based on assigning privileges to access and scan an environment using cloud APIs. For example, in Exhibit 2115, which I discussed above, Wiz characterizes "Agentless scanning approaches" as being "based on assigning privileges to security vendors to access and scan an environment using cloud APIs." Ex.2115 at 1. Wiz then "outline[s]" its "approach to secure agentless scanning" by describing the "Wiz workload scanner." Ex.2115 at 2. I therefore understand Wiz's description of "Agentless scanning approaches" in this document to encompass its own Workload Scanner.

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

Ex.2115 at 1.

235. 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.

236.

REDACTED

REDACTED

- b. Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.2-1.4, 13.2-13.4, 22.2-22.4: Using The Cloud Computing Platform APIs To Identify Virtual Disks Of A Virtual**

**Machine, and Query and Receive A Location Of The
Virtual Disk(s)**

237. Wiz's Workload Scanner is substantially similar to Orca's SideScanningTM in relation to limitations 1.2-1.4, 13.2-13.4, and 22.2-22.4 of claims 1, 13, and 22 of the '685 Patent.

238. Wiz's Workload Scanner uses the interface between a client's environment and Wiz's security components to utilize cloud computing platform APIs to identify virtual disks of a virtual machine in the client environment, and to query and receive an identification of a location of the virtual disk(s).

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

239. REDACTED

REDACTED

REDACTED

241. In some organizations' security configurations, Wiz's Workload Scanner emulates the virtual disks of a virtual machine to generate the snapshot. For example, Wiz describes the role of hypervisors to emulating hardware.

How VMs work

Virtual machines use virtualization to achieve resource isolation. In virtualization, host machines consist of separate VMs with separate operating systems and dedicated virtual hardware. This level of logical separation is made possible by a hypervisor, a software layer that manages the physical resources and allocates these resources to the VMs as required.

As illustrated below, hypervisors play an important role in the process of emulating underlying hardware.




Figure 3: Virtualization

Ex.2075 at 3-4

- c. **Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.5, 13.5, 22.5: generate at least one snapshot of the virtual disk or emulate the virtual disks to generate at least one snapshot**

242. Wiz's Workload Scanner is substantially similar to Orca's SideScanningTM in relation to limitations 1.5, 13.5, and 22.5 of claims 1, 13, and 22 of the '685 Patent. Wiz's Workload Scanner will generate at least one snapshot of a virtual machine.

REDACTED


REDACTED

- d. **Wiz's Workload Scanner is substantially similar to Orca's SideScanning in relation to limitations 1.6, 13.6, 22.6: Analyz[e/ing] The At Least One Snapshot To Detect Vulnerabilities And Sensitive Data, Wherein Analyzing The At Least One Snapshot Requires No Interaction With The Virtual Machine;**

245. Wiz's Workload Scanner is substantially similar to Orca's SideScanningTM in relation to limitations 1.6, 13.6, and 22.6 of claims 1, 13, and 22 of the '685 Patent (reproduced in full above).

246. 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

247. Wiz’s platform also detects sensitive data. For example, in November 2022 Wiz published an article explain that “its platform [] proactively eliminate[s] attach paths to discovered critical data.” Ex.2155 at 1. Wiz explained that one of its customers “uses Wiz to *discover* and protect mission-critical data.” Ex.2155 at

3. “Wiz scans public buckets..., data volumes, [etc.]... and classifies the data so organizations can easily answer the question of what data is located where.” Ex.2155 at 4.


248. Wiz’s Workload Scanner is “able to scan virtual machines or containers even if the workload is offline” and therefore does not require interacting with the VM while analyzing the snapshot. Ex.2165 at 14.

e. Wiz’s Workload Scanner is substantially similar to Orca’s SideScanning in relation to limitations 1.7, 13.7, 22.7: Determin[e/ing] a Risk Level of the Virtual Machine

249. Wiz’s Workload Scanner uses the features recited in limitations 1.7, 13.7, and 22.7 of claims 1, 13, and 22 of the ’685 Patent (reproduced in full above).

250. For example, Wiz provides “Cloud Security Posture Management” or “CSPM” that it describes as “giv[ing] you a clear risk assessment of your cloud security posture across all providers.” Ex.2156 at 2. Wiz also provides vulnerability analysis identifying, as shown below, numbers of “critical issues,” “high issues,” “medium issues,” and “low issues.”

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

251. Wiz provides an example of a “Attack Path Visualization” in its product of a “Publicly exposed VM with high privileges and high/critical severity network vulnerabilities with a known exploit.” Ex.2165 at 18.




Ex.2165 at 18

- f. **Wiz’s Workload Scanner is substantially similar to Orca’s SideScanning in relation to limitations 1.8, 13.8, 22.8: Report[ing] the Detected Vulnerabilities and Sensitive Data As Alerts, Wherein the Alerts Are Filtered and Prioritized Based on the Determined Risk Level**

252. Wiz’s Workload Scanner uses the features recited in limitations 1.8, 13.8, and 22.8 of claims 1, 13, and 22 of the ’685 Patent (reproduced in full above).

253. 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

254. Wiz’s platform also reports sensitive data. For example, in November 2022 Wiz published an article explain that “its platform [] proactively eliminate[s] attach paths to discovered critical data.” Ex.2155 at 1. Wiz explained that one of its customers “uses Wiz to *discover* and protect mission-critical data.” Ex.2155 at 3. “Wiz scans public buckets..., data volumes, [etc.]... and classifies the data so organizations can easily answer the question of what data is located where.”

Ex.2155 at 4. Wiz’s product includes “DSPM” which “automates discovery” of sensitive data and “then provides... alerting.” Ex.2157 at 4.

255. Wiz’s Workload Scanner reports issues as alerts.

Stage 6: Vulnerability analysis



Example vulnerability dashboard that prioritizes issues by contextual severity

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

Ex.2165 at 8

F. Commercial Success

256. 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

257. 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).

258. It is my opinion that the commercial success of both Orca's and Wiz's products has nexus to the claims of the '685 patent. For example, both systems practice the claims of the '685 patent as discussed above. Additionally, both entities tout their usage of the claimed features in their marketing materials. *See, e.g.,* Ex.2107 at 2-3 (touting Orca's technology including determining which servers and vulnerable and "[w]here is my sensitive data stored" in an

IPR2024-00865 (USP 11,693,685)

Dec'l of David Kaeli

announcement about fundraising).

REDACTED

X. CONCLUSION

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

Executed on 4/7/2025, in Boston, Massachusetts.



Dr. David Kaeli