

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

WIZ, INC.,
Petitioner,

v.

ORCA SECURITY LTD.,
Patent Owner.

IPR2024-01109

PATENT NO. 11,726,809

DECLARATION OF ANGELOS STAVROU

Table of Contents

I. Background	1
II. Materials Considered	2
III. Level of Ordinary Skill in the Art.....	5
IV. Legal Principles.....	5
V. Claim Construction	5
A. “location”	5
B. Analyze/Analyzing “the at least one snapshot”	6
VI. Ground 1.....	6
A. Dr. Kaeli’s Embodiment-Mixing Argument.....	6
B. The Veselov-Mohanty Combination.....	11
C. The Combination Satisfies the Claims.....	25
1. Elements 1.5/16.5/19.5.....	26
2. Elements 1.6/16.6/19.6.....	28
3. Elements 1.7/16.7/19.7.....	30
4. Elements 1.8/16.8/19.8.....	30
5. Claims 15/22	32
VII. Ground 2.....	41
A. The Further Combination with Czarny	41
B. Element 1.8.....	44
A. Claims 6-7 and 17-18	44
VIII. Grounds 3-4.....	45
IX. Secondary Considerations.....	48
A. No Presumption of Nexus	50
B. Orca’s Evidence	56
1. Known Techniques/Features.....	56
2. Long-Felt Need	68
3. Skepticism	71
4. Praise	74

5.	Copying	80
6.	Commercial Success	92
X.	Concluding Statements	95

I, Angelos Stavrou, declare as follows.

I. Background

1. As noted in my previous declaration, I have been retained by Wilson Sonsini Goodrich & Rosati Professional Corporation on behalf of Petitioner Wiz, Inc. (“Wiz”) to provide my opinions concerning the validity of claims 1-23 of U.S. Patent No. 11,726,809 (“’809 Patent” or EX1001) in IPR2024-01109. I understand that this IPR was instituted. I also understand that Orca Security Ltd. (“Patent Owner” or “PO”) submitted a Patent Owner Response (“POR”) that was accompanied by declarations from Dr. David Kaeli (“Kaeli-Dec” or EX2073) and Mr. Avi Shua (“Shua-Dec” or EX2074).

2. Although I am being compensated at my standard consulting rate for the time I spend on this matter, no part of my compensation depends on the outcome of this proceeding, and I have no other interest in this proceeding, any of the parties, or the challenged patent.

3. The statements made and opinions provided in this Declaration are based on my own personal knowledge and, if called as a witness, I could and would testify in a manner consistent with this Declaration.

4. In this Declaration, all emphasis has been added unless otherwise noted.

II. Materials Considered

5. In forming my opinions, I have reviewed and am familiar with the challenged patent, its prosecution history, and the other materials cited in my previous declaration and in this declaration. I have also considered the POR and Orca's exhibits, including the declarations of Dr. Kaeli and Mr. Shua. In addition to the list of exhibits provided in my previous declaration and Orca's exhibits, some of which I discuss below, I cite the following additional exhibits in this declaration.

Exhibit No.	Description
1115	Avi Shua deposition exhibit of "Cloud Security: Vendor Landscape and TCO Analysis," Orca Security Ltd. (June 12, 2025) https://orca.security/lp/tag-cyber-cloud-security-vendor-landscape-tco-analysis (accessed June 11, 2025)
1118	Dr. Kaeli deposition exhibit of US Patent No. 10,834,050 (June 19, 2025)
1119	Dr. Kaeli deposition exhibit of US Patent No. 9,805,202 (June 19, 2025)
1120	Dr. Kaeli deposition exhibit of US Patent No. 10,956,242 (June 19, 2025)
1121	Dr. Kaeli's June 19, 2025, deposition exhibit of an archived copy of "CreateSnapshot," archived by the Internet Archive's "Wayback Machine" https://web.archive.org/web/20141226192222/https://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_CreateSnapshot.html (December 26, 2014)
1122	Nexpose User's Guide, Product Version: 5.16, Rapid7, LLC. Copyright (2015)

1123	Nessus Enterprise for Amazon Web Services (AWS) Installation and Configuration Guide, Copyright © 2014. Tenable Network Security, Inc., Revision 2 (July 16, 2014)
1124	Configuring AWS EC2 Connectors with Cross-Account Role Authentication (ARN), Hari Srinivasan (/discussions/s/profile/0052L000003olee), (April 24, 2018) (last modified April 27, 2018)
1125	Redefining CNAPP: A Complete Guide To the Future of Cloud Security, Francis and James Berthoty, https://softwareanalyst.substack.com/p/redefining-cnapp-a-complete-guide
1126	How Orca Security Compares to Wiz, https://orca.security/resources/comparisons/wiz/ , Orca Security (2025)
1127	Veselov LinkedIn, https://www.linkedin.com/in/vladimir-veselov-2769753/details/experience/ (personal information redacted)
1129	Michael et al. US Patent No. 8,011,010 B2 (April 17, 2007)
1131	Amazon EC2 Backup and Recovery with Snapshots and AMIs - AWS Prescriptive Guidance, https://docs.aws.amazon.com/prescriptive-guidance/latest/bachup-recovery/ec2-backup.html , (2025)
1132	Creating an Amazon EBS-Backed Linux AMI, https://web.archive.org/web/20170601003305/https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/creating-an-ami-ebs.html , Amazon Web Services, Inc. (2017)
1133	Amazon EBS volumes, Amazon EBS, https://docs.aws.amazon.com/ebs/latest/userguide/ebs-volumes.html , Amazon Web Services, Inc. (2025)
1134	Symantec Data Loss Prevention Solution, Symantec Corporation (2015)
1135	“Deployment Guide,” Forcepoint, DLP ver., 8.4.x (2017)
1136	McAfee Skyhigh Security Cloud, Data Security for the Cloud Era, McAfee, LLC. (2018)

1137	“Google Announces Agreement to Acquire Wiz” Inside Google, Google LLC., https://www.prnewswire.com/news-releases/google-announces-agreement-to-acquire-wiz-302404500.html/
1138	“Security Audit” Reference Guide by Amazon https://docs.aws.amazon.com/aws-managed-policy/latest/reference/SecurityAudit.html)
1139	Scarfone et al., “Technical Guide to Information Security Testing and Assessment: Recommendations of the National Institute of Standards and Technology,” NIST Special Publication 800-115 (Sept. 2008) https://doi.org/10.6028/NIST.SP.800-115
1140	“Vuls: Agentless Vulnerability Scanner for Linux/FreeBSD,” archived on September 1, 2018, by the Internet Archive’s “Wayback Machine” (https://web.archive.org/web/20180901153521/https://vuls.io)
1141	Rapid7 Research, “Quarterly Threat Report: Q3 2018” (Nov. 13, 2018) https://www.rapid7.com/globalassets/_pdfs/research/rapid7-threat-report-2018q3.pdf
1403	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1404	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1405	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1406	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1408	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1413	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1414	[REDACTED] [PROTECTIVE ORDER MATERIAL]

6. I have also relied on my education, skill, training, and experience in the relevant fields of technology. In forming these opinions, I have applied the viewpoint of a person of ordinary skill in the art (“POSA”) as of January 28, 2019.

III. Level of Ordinary Skill in the Art

7. I understand that Dr. Kaeli provides a slightly different description of the POSA compared to the one I set forth in my original declaration. *Compare* Kaeli-Dec, ¶30, with EX1002, ¶22. The differences in these descriptions of the POSA do not affect the specific descriptions of the POSA’s understanding that I discussed in my previous declaration or in this declaration, including the POSA’s familiarity with the relevant technology, the POSA’s approach to combining teachings, and the POSA’s understanding of the references. Therefore, I will continue to apply the level of ordinary skill that I applied in my original declaration, but my opinions would not change if were to apply Dr. Kaeli’s version.

IV. Legal Principles

8. In forming my opinions for this declaration, I have applied the same legal principles that I described in my original declaration. I further note specific principles involving secondary considerations below in line with my discussion of the evidence. *Infra*, §IX.

V. Claim Construction

A. “location”

9. I understand that Dr. Kaeli does not dispute the POSA's understanding of the term "location" that I provided in my previous declaration. Kaeli-Dec, ¶¶42-47.

B. Analyze/Analyzing "the at least one snapshot"

10. I understand that Dr. Kaeli argues that this limitation does not encompass analyzing a duplicate VM created from the snapshot. Kaeli-Dec, ¶¶43-47. Even if his view were adopted, this would not impact my previous conclusion that the claims were obvious in view of the Veselov combinations since Veselov teaches analyzing duplicate VMs instantiated from snapshots and directly analyzing snapshots as data files as alternative ways of conducting the analysis step. EX1007 (Veselov), 3:20-4:4, 10:1-16. In other words, there does not appear to be any dispute that Veselov's snapshot-as-a-data-file analysis option satisfies this limitation.

VI. Ground 1

A. Dr. Kaeli's Embodiment-Mixing Argument

11. I understand that Dr. Kaeli argues that my previous declaration mixes disclosures from Figures 1, 3A/B, and 5A/B without explaining why a POSA "would have wanted to combine or modify those varied implementations." Kaeli-Dec, ¶72. Respectfully, Dr. Kaeli is not accurately characterizing my analysis and or Veselov's disclosure. As I previously explained, "Veselov makes it clear that

when a portion of the disclosure associated with a particular figure/embodiment describes certain techniques or components, those teachings are broadly applicable to the rest of Veselov's disclosure." EX1002, ¶91. This is not describing a modification of one system in view of another system. Rather, I explained what a POSA would have understood Veselov as teaching/suggesting for each of these embodiments. In other words, a POSA would have recognized that, aside from certain aspects of the analysis steps themselves—e.g., causing processing activity in the assessment VM of Figs. 3A/B, an example Dr. Kaeli focuses on (Kaeli-Dec, ¶76)—Veselov is describing features common to all embodiments and then simply not repeating itself later on to minimize unnecessary redundancy. Veselov is explicit about this, as I previously explained (EX1002, ¶91), and this would have made perfect sense to a POSA since the steps before the analysis (e.g., receiving scan requests, establishing interfaces, using cloud APIs to identify/locate virtual disks for the snapshots, and creating the snapshots) and after the analysis (e.g., processing/reporting the results as alerts) do not depend on whether the analysis step uses the snapshot to instantiate an assessment VM or reads the snapshot as a data file. *See, e.g.*, EX1007, 3:20-4:4, 8:52-10:1-36 (discussing Fig. 2 process, including alternative analysis options for step 208), Fig. 2. It is also consistent with the POSA's understanding of the relevant technology, since the POSA would have recognized that the well-known cloud-API usage that I discussed in my prior

declaration—including the different equivalent ways to invoke the cloud APIs (web, console, programmatically, etc.) to achieve the same functionality—were equally applicable regardless of which of Veselov’s analysis options were used. I discuss specific issues with Dr. Kaeli’s embodiment-mixing arguments below.

12. Dr. Kaeli’s assertion that APIs are only used in the embodiments associated with Figs. 1 and 3A/B, not those of Figs. 5A/B, is based on a fundamental misunderstanding of the computing environment in Figs. 5A/B. He states that the computing environments of Figs. 1 and 3A/B are implemented by a computing resource service provider that provides the APIs, whereas the environments of Figs. 5A/B are supposedly *not* implemented by this service provider. Kaeli-Dec, ¶¶73-79, 133. This characterization of Figs. 5A/B is incorrect. First, Veselov identifies the service provider as an undepicted aspect of the computing environment illustrated in Figs. 5A/B when, in discussing operations by scanning service 510 involving volume image 508, it refers to using a “security module or security service (not shown) of the computing resource service provider.” EX1007, 18:1-15. That the 5A/B environment is hosted by a cloud provider (and therefore uses the same APIs Dr. Kaeli acknowledges are associated with this provider) is also made clear by Veselov’s express description of that environment. For example, Dr. Kaeli acknowledges that the 3A/B environment is implemented by the service provider (Kaeli-Dec, ¶74), and Veselov

states that the 5A/B environment has the same “implementation” as the 3A/B environment. EX1007, 10:37-41, 16:56-67. Veselov also states that the 5A/B environments have the same “interactions” as the 3A/B environments (*id.*, 16:56-67), which Dr. Kaeli similarly acknowledges uses API interactions. Kaeli-Dec, ¶74.

13. I disagree with Dr. Kaeli’s suggestion that direct snapshot analysis is limited to Figs. 5A/B. He treats Figs. 5A/B as if it is the only example of direct snapshot analysis in Veselov. *See* Kaeli-Dec, ¶¶72-73. This is incorrect. Veselov discusses the analysis of the snapshot as a data file when giving an initial overview of its techniques. EX1007, 3:51-67. Veselov states that the API-using Fig. 1 environment may be used by “embodiments of the present disclosure.” *Id.*, 4:19-29. Veselov states that the process of Fig. 2 may use that environment (*id.*, 8:52-62), and that process again includes analysis of the snapshot as a data file (*id.*, 10:1-16).

14. I also disagree with Dr. Kaeli’s suggestion that Veselov’s APIs are not “cloud” APIs. He states that the term “computing resource service provider” is not typically used to refer to cloud environments (Kaeli-Dec, ¶133), but it would have been perfectly clear to a POSA that Veselov is describing a cloud provider. EX1007, 2:22-3:19 (techniques “particularly suited for security assessments performed in a data center” and address concerns of “cloud” environments); *see*

also EX1118 (Call), 4:32-41; EX1119 (Medeiros), 2:41-52; EX1120 (Kumar), 1:10-14, 1:64-2:5 (using “computing resource service provider” to refer to cloud provider). That is hardly surprising since Veselov is a patent by the world’s largest cloud provider (Amazon, whose cloud service is called AWS), and Veselov himself was an AWS engineer. EX1007, Face; EX1127 (LinkedIn page).¹

15. Dr. Kaeli also presents an overly narrow view of how APIs are used in Veselov. He later states that APIs are not needed to obtain snapshots in Veselov, and that snapshots would be “obtained by means other than APIs even in Figures 1 and 3A-B, and certainly within Figures 5A-5B, where no API even exists.” Kaeli-Dec, ¶133. However, a POSA would not have shared this limited view of API usage in Veselov. Veselov explains that the “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.” EX1007, 5:1-22. In other words, Veselov broadly teaches that the scanning service interacts with the service provider and its resources via API calls. Indeed, that is the paradigm a POSA would have expected since cloud-API calls have long been

¹ EX1127 is a copy of Mr. Veselov’s page on LinkedIn at <https://www.linkedin.com/in/vladimir-veselov-2769753/>. I have confirmed that EX1127 is an accurate representation of his bio on that webpage.

the standard way that virtually every cloud-security platform—and cloud-based products in general—interact with cloud providers, as I discuss in more detail below. *Infra*, §IX.B.1.

B. The Veselov-Mohanty Combination

16. Dr. Kaeli presents several arguments against the combination of Veselov and Mohanty. As I explain below, I disagree with his assessment.

17. I disagree with Dr. Kaeli’s argument that a POSA would not have combined Mohanty with Veselov. He claims Mohanty criticizes snapshot-based technologies. Kaeli-Dec, ¶81. First, the passage of Mohanty’s background section that Dr. Kaeli quotes refers to a “snapshot **of a vulnerability assessment report**” eventually becoming obsolete. EX1075, 1:14-17 (emphasis added); Kaeli-Dec, ¶¶81, 87 (citing same). In other words, Mohanty is not referring to the types of snapshots at issue in Veselov. Regardless, a POSA would have recognized that Veselov’s approach is reasonably current. For example, Veselov describes taking fresh snapshots rather than using old snapshots. EX1007, 8:62-9:25. The assessment can also be initiated on-demand (or triggered by an event) such that the user is getting a reasonably current view of the system’s security status. *E.g., id.*, 6:25-49, 14:51-52. Veselov also describes periodic assessments (*id.*, 4:11-18), meaning such assessments would be reasonably current and continuous. In any event, Dr. Kaeli’s point does not address the net benefit provided by applying

Mohanty's techniques to snapshot-based assessments. EX1002, ¶¶111-15. In other words, Veselov explains several benefits of snapshot-based assessments and explains that they yield results comparable to assessments performed on the target itself (i.e., agent-based assessments). EX1007, 2:42-4:4. Applying the teachings of Mohanty that I previously discussed then provides further benefits to that snapshot-based assessment without compromising Veselov's existing benefits. EX1002, ¶¶111-15. A POSA thus would have been motivated to make this combination as I previously discussed.

18. I disagree with Dr. Kaeli's additional assertion that a POSA considering Veselov would not have looked to Mohanty. He claims that because Veselov overcomes certain drawbacks of agent-based assessments, a POSA "would have therefore not looked to systems that rely on agent-based assessments, like Mohanty." Kaeli-Dec, ¶82; *see also id.*, ¶86. As an initial matter, Dr. Kaeli appears to assume Mohanty is limited to agent-based scanners, but Mohanty does not teach this. It discusses vulnerability scanners generally, and these were not limited to agent-based scanners. EX1075, 4:59-5:15; *see* EX1040 (Ammons), Abstract (agentless scanning), ¶¶43-46 (snapshot-based assessment performing tasks such as compliance checks, installed-application checks, file-system

inspection, drift detection, malware analysis, etc.); EX1139 (2008 NIST Guide),² 4-4 to 4-6 (“Vulnerability scanners can be run against a host either locally or from the network.”); EX1140 (Vuls) (discussing “agent-less vulnerability scanner” that matches installed software against vulnerability databases)³. It discusses Qualys and Rapid7 later as “example[s].” *Id.*, 6:54-63. Even if Mohanty were so limited, the applied teachings of Mohanty—e.g., correlating vulnerabilities with network location, assessing risk, and reporting prioritized alerts (EX1002, ¶¶105-06, 109)—do not depend on the use of agents. As I previously explained, the data that is ultimately considered in Veselov—e.g., applications, files, configurations, etc.—is the same type of data collected by Mohanty for its analysis, and that both references are addressing the same types of security issues. EX1002, ¶110. A POSA therefore would have recognized that the teachings of Mohanty I discussed

² EX1139 is a correct copy of a September 2008 guide by the National Institute of Standards and Technology (“NIST”) titled “Technical Guide to Information Security Testing and Assessment” and obtained from <https://doi.org/10.6028/NIST.SP.800-115>.

³ EX1140 is a correct copy of a webpage archived at <https://web.archive.org/web/20180901153521/https://vuls.io/>.

in my previous declaration were applicable and useful to snapshot-based assessments taught by Veselov.

19. Dr. Kaeli also mistakenly assumes that the combination involves the installation of agents on the target. He states that Mohanty runs counter to Veselov's goal of avoiding installation of agents on targets and associated interference. Kaeli-Dec, ¶82. He argues that material differences would have led a POSA away from the combination, namely that Veselov disparages Mohanty's supposed process of installing agent-based scanners "*on the target assets.*" *Id.*, ¶¶83-86. He argues that "the only implementation in Veselov that could be combined with Mohanty's teachings is" the 3A/B approach because "it is not possible to install an agent on a snapshot." *Id.*, ¶87. However, at no point did I discuss installing agents on the target asset as part of the combination. EX1002, ¶¶104-16. Rather, I explained the POSA's motivations and reasonable expectation of success in applying Mohanty's teachings—e.g., correlating network location with vulnerabilities, measuring context-sensitive risk, and reporting prioritized alerts—in the context of snapshot-based assessments. *Id.*

20. I disagree with Dr. Kaeli's argument that Veselov and Mohanty supposedly do not consider the same types of data and security issues. He does not appear to dispute my actual point: that Mohanty considers the target system's files, applications, and configurations, which are the same types of data stored and

readily analyzed in snapshots. EX1002, ¶110. Instead, Dr. Kaeli focuses on the *process* for collecting data (or the component that does the collecting), rather than the data itself. Kaeli-Dec, ¶88. For example, he emphasizes that snapshotting “does not require guest cooperation, whereas agent assistance access requires guest cooperation” and that agents run inside the target. *Id.* (citing EX1038 (Suneja), 134). He also argues that agents “could not be installed on Veselov’s snapshots” and that snapshots are “not an active, executing computing device, like a virtual machine.” Kaeli-Dec, ¶¶88-89. None of this addresses the collected data itself, which was widely understood to include things like files, applications, and configurations, as I previously explained. EX1002, ¶¶54, 110, 155; EX1030 (Bairwa), 116; EX1031, 27-31 (distinguishing “Data Collection”—which identifies “software and system configuration”—and “Data Analytics” of the collected data); EX1032 (Li), ¶¶27, 43; EX1075, 4:59-5:15. This is basic knowledge a POSA would have had, and it had been well known for many years. *E.g.*, EX1139 (2008 NIST Guide), ES-1 to ES-2 (guide discusses “basic technical aspects” of security assessments), 4-4 to 4-6 (“[V]ulnerability scanning identifies hosts and host attributes (e.g., operating systems, applications, open ports)...”). The process by which an agent collects that data does not address the POSA’s understanding that snapshots contain the same types of data and that this data is readily collected and analyzed in the snapshot, particularly given Veselov’s express disclosure of robust

snapshot content and straightforward ways of collecting that data for analysis.

EX1007, 3:61-67 (mount logical volume as hard drive or transform snapshot into data file “then perform the security assessment on the data at the file system level”), 8:30-45, 9:42-10:16, 17:9-29 (mounting volume image and accessing it “as if it were a disk drive”), 18:32-53 (parse raw data using string/pattern comparison).

21. Dr. Kaeli argues that the industry viewed agentless scanning as “rudimentary” compared to agent-based scanning. Kaeli-Dec, ¶90 (citing EX2106, 1-2). But the quoted exhibit is a Palo Alto Networks blog post touting their own product. EX2106. In other words, it is marketing, not an objective/disinterested description of the field. Regardless, even if Dr. Kaeli’s characterization (based on a marketing document’s discussion of certain competing products) were accurate, this would not meaningfully address the straightforward nature of applying Mohanty’s teachings involving context-sensitive risk assessments (including network-location correlations) and prioritized alerts to snapshot-based assessments, nor does it address the benefits of the combination that I explained. EX1002, ¶¶109-16. This combination does not involve agents, nor does it involve data-collection or data-analysis techniques that require agent-based scanners.

22. Dr. Kaeli’s arguments are also directly contradicted by Veselov. Dr. Kaeli argues that Veselov is consistent with the industry’s supposed view of agentless scanning as rudimentary and that Veselov “makes clear that agent-based

security analyses (Veselov's Figs. 3A-3B) result in significantly different data compared to analyzing a snapshot (Veselov's Figs. 5A-5B)." Kaeli-Dec, ¶91. This is incorrect. Veselov describes both the assessment-VM approach and the direct-snapshot-analysis approach as yielding results comparable to assessments performed on the target asset. *Id.*, 3:41-4:4. Dr. Kaeli also claims the data detected when analyzing the snapshot in Figs. 5A/B is "limited" compared to Figs. 3A/B. *Id.*, ¶¶91-92. In particular, he claims the snapshot does not include information about network connections, communications, "or any of the other information available when using a security agent as in the Figure 3A-3B embodiments." *Id.*, ¶92. This is also incorrect. Veselov states that the snapshot is the "same" in both Figs. 3A/B and 5A/B. EX1007, 16:56-67. While additional assessment tasks are described for Figs. 3A/B, Veselov explains that the "activity and connections" are captured "in the snapshot data" in that example.⁴ *Id.*, 15:41-

⁴ Dr. Kaeli includes a footnote arguing that Veselov's assessment-VM approach requires agents installed inside the assessment VM. Kaeli-Dec, ¶92 n.3. Even if this were true, it would not impact the combination I discussed or the claim-element mappings. In any event, his assertion is incorrect. Veselov describes examples in which an agent is installed inside the assessment VM, as I noted, but

57; *see also id.*, 15:57-59 (statically captured snapshot “may include data that describes the expected activity and/or connections”), 16:3-6 (scanning system may intercept traffic “before or during generation of the snapshot data, in order to create a profile of the traffic to the target virtual machine 312”). Regardless, Veselov’s general description of the snapshot used for both approaches shows that the snapshot includes the relevant data for purposes of the combination, including network-location information and various types of information commonly used to support risk assessments. EX1007, 3:20-40, 9:42-10:16 (snapshot may include data such as file/OS information, “configurations of network protocols and services; properties of connections to other computing resources; permission settings; software packages installed, and configurations thereof; system log properties; and the like”); *see also id.*, 2:35-41 (discussing assessment of “host configuration”).

23. Dr. Kaeli’s arguments involving Veselov’s discussion of string/pattern comparison in column 18 is misleading for additional reasons. He points to this description to support his assertion that the 5A/B snapshot data is “limited” compared to Figs. 3A/B. Kaeli-Dec, ¶92. In this passage, Veselov is discussing an

such agents are not required. In fact, Veselov expressly describes an “external agent” for Figs. 3A/B. EX1007, 16:10-14.

example of how data in the snapshot may be identified when performing direct snapshot analysis, not the assessment tasks themselves. EX1007, 18:32-53 (identify data, then determine “which assessment tasks to run based on the installed software packages”). Veselov is also describing examples of data that may be collected, not all data that is available for the assessment. For example, Veselov’s general discussion of snapshot content explains that the snapshot may also include machine-memory state—a routine aspect of snapshots in some platforms—and I previously discussed well-known analysis techniques for such data. EX1002, ¶¶32-36 (discussing snapshots), ¶¶54-58 (discussing analysis techniques, including VMI); EX1014 (Klemperer), 24-25. Regardless, Veselov’s exemplary description of identifying data in the snapshot does not describe “limited” content, as Dr. Kaeli claims. It is an open-ended list, and the provided examples themselves cover a vast range of relevant data. *Id.*, 18:32-53 (identifying “files, software packages, application versions, and the like”). Beyond the applications/versions, numerous types of information are stored as “files,” including much of the previously discussed network information. *See id.*, 3:20-40, 7:16-22, 9:42-67.

24. I disagree with Dr. Kaeli’s arguments addressing the background exhibits that I cited in my original declaration. Kaeli-Dec, ¶¶93-96. I discussed these, for example, to illustrate the POSA’s understanding of the routine and beneficial nature of the relevant teachings, the POSA’s understanding of the

relevant context of the references’ teachings, and the POSA’s reasonable expectation of success. EX1002, ¶¶109-16. Dr. Kaeli argues that many “do not implement Mohanty’s agent-based scanning and therefore provide no support” for the combination, while he argues that others do not support the combination because they *do* use agents. Kaeli-Dec, ¶93 (citing EX1014, EX1025, EX1034, EX1044, EX1045, EX1046, EX1076, EX1078, EX1085; EX1063; EX1008); *see also* Kaeli-Dec, ¶97 (dismissing EX1077 because it involves agents). He similarly dismisses EX1025 (NIST Incident Handling). Kaeli-Dec, ¶93. Dr. Kaeli does not explain these arguments, nor does he substantively respond to the particular points for which I cited those exhibits.

25. I also disagree with Dr. Kaeli’s dismissal of EX1048 (Price). I cited this exhibit in the Ground 1 rationale section to illustrate the well-known and predictable nature of prioritized alerts in both agent-based and agentless contexts. EX1002, ¶109; EX1048, ¶43. Dr. Kaeli argues that this exhibit is inapplicable because it involves “images” rather than snapshots. Kaeli-Dec, ¶94. But the term “image” itself was often used to broadly refer to system/disk state. While this is not literally synonymous with “snapshot,” this broader sense of “image” was often used in a way that encompassed snapshots. *See, e.g.*, EX1040 (Ammons), ¶43 (discussing function that “takes a ‘snapshot’ virtual image”); EX1129 (Michael), 2:28-48 (“Also, a copy of the virtual machine image may be stored by taking a

snapshot or a checkpoint of the virtual machine while the virtual machine is on-line.”). Indeed, Veselov itself describes an example snapshot as a volume “image.” EX1007, Abstract. To the extent Dr. Kaeli is focusing on specific machine-image resources vs. snapshot resources as used by cloud/virtualization platforms, while these resource types are not synonymous, they are not the fundamentally distinct concepts Dr. Kaeli implies. For example, in AWS, the machine-image resource (referred to as an “Amazon Machine Image” or “AMI”) is essentially just a wrapper for one or more snapshots. A snapshot in AWS encompasses a single volume/disk, which may or may not be a boot volume (EX1121),⁵ whereas an AMI essentially just includes one or more snapshots (one of which is a boot volume) along with metadata such as launch parameters.

⁵ I understand that EX1121 is a copy of archived AWS documentation obtained from

https://web.archive.org/web/20141226192222/http://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_CreateSnapshot.html and introduced during Dr. Kaeli’s deposition. I have compared it with the webpage at that URL and confirm that it is an accurate representation.

EX1131.⁶ In fact, an AMI can simply be created from a snapshot. EX1132.⁷ This type of close relationship between machine images and snapshots was common in cloud/virtualization platforms. I therefore disagree with Dr. Kaeli's argument that Price is irrelevant.

26. I also disagree with Dr. Kaeli's argument that Mohanty's techniques are "squarely against Veselov's efforts to reduce the complexity of security assessments." Kaeli-Dec, ¶¶97-98 (citing EX1007, 2:42-46). As an initial matter, Veselov's comment addresses how snapshot-based scanning avoids complex installation/monitoring of *agents* on the target asset, not complexity generally. EX1007, 2:42-3:19. Regardless, context-sensitive risk assessment/prioritization was commonplace, and a POSA would not have viewed the applied teachings as

⁶ EX1131 is a correct copy AWS documentation obtained from <https://docs.aws.amazon.com/prescriptive-guidance/latest/backup-recovery/ec2-backup.html>. Based on my experience, its description of AMIs including one or more snapshots is consistent with the state of the art in January 2019.

⁷ EX1132 is a correct copy of AWS documentation archived at <https://web.archive.org/web/20170601003305/http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/creating-an-ami-ebs.html>.

overly complex. In fact, Veselov encourages performing multiple types of assessment tasks. *Id.*, 2:35-41.

27. I disagree with Dr. Kaeli's response to my explanation of motivation "to provide more accurate assessment of risks posed by particular vulnerabilities and the asset as a whole." EX1002, ¶112. Dr. Kaeli repeats his earlier argument that Mohanty "discourages" using snapshots because they are point-in-time captures that become obsolete. Kaeli-Dec, ¶98. As I explained above, however, a POSA would have understood that Veselov's approach is reasonably current and continuous given its teaching of using fresh snapshots (EX1007, 8:62-9:25) and performing assessments on-demand, triggered, or periodically. *Id.*, 4:11-18, 6:25-31, 14:51-52. Regardless, it would have been clear to a POSA that measuring context-sensitive risk would be similarly useful to assessments of live VMs and snapshot-based assessments. Risk levels do not become irrelevant simply because assessments are based on snapshots (let alone the reasonably current snapshots taught by Veselov). Dr. Kaeli repeats his arguments that Mohanty relies on agents that cannot be run in a snapshot and that Veselov discourages the complexity added by Mohanty's teachings. Kaeli-Dec, ¶98. I explained why these arguments are incorrect above.

28. I also disagree with Dr. Kaeli's response to my explanation of motivation "to obtain known benefits of prioritizing vulnerabilities and associated

alerts.” EX1002, ¶113; Kaeli-Dec, ¶¶99-101. Dr. Kaeli argues that Mohanty’s context “is completely divorced from the goals set forth in Veselov” and that Mohanty’s scanning is vastly different from Veselov’s. Kaeli-Dec, ¶99. This is inaccurate. For example, as I previously explained, both references provide cybersecurity for virtual assets using analogous techniques (e.g., collecting and assessing asset data to detect and evaluate potential security problems), and the applied teachings of Mohanty are readily applicable to snapshot-based assessments. EX1002, ¶¶108-16. Dr. Kaeli argues that Mohanty’s techniques would not work with or benefit Veselov because they are supposedly reliant on agents and use “real-time, dynamic data from executing processes on a virtual machine,” whereas Veselov considers static snapshot data. Kaeli-Dec, ¶100. As I explained above, however, even if Mohanty were limited to agent-based scanners, agents collect the same types of data—e.g., files, applications, and configurations—to detect vulnerabilities, measure context-sensitive risk, and prioritize/report alerts based on that risk. Once that data is collected by an agent, the data is used in the same predictable way that it would be used in Veselov. EX1002, ¶110. Nor is the changing nature of the asset a meaningful distinction since the virtual assets in Mohanty and Veselov are both subject to change over time. A POSA would have been similarly interested in measuring context-sensitive risk with Veselov’s reasonably current and continuous snapshot-based

approach. Dr. Kaeli further argues that Veselov does not teach alerts and that it was not obvious to use alerts because Veselov's findings are "finite." Kaeli-Dec, ¶111. But Veselov does teach alerts, as I explain below when discussing element 1.8, and alerts were commonplace and useful in agent-based and agentless scanning alike. EX1002, ¶¶64-68, 181; EX1008 (Basavapatna), Abstract, ¶138; EX1044 (Giakouminakis), 6:9-7:10, 10:3-20, 11:51-12:19, Fig. 2; EX1048 (Price), ¶43.

29. I also disagree with Dr. Kaeli's arguments regarding the POSA's reasonable expectation of success. He again argues that agents cannot be installed on snapshots. Kaeli-Dec, ¶¶102-04. As I previously explained, however, the combination does not involve the incorporation of agents, and the references consider comparable underlying data. Dr. Kaeli's arguments thus continue to misunderstand the combination and conflate the component/process collecting the data with the data itself.

30. It is therefore still my opinion that a POSA had reason to combine the teachings of the references to arrive at the claims with a reasonable expectation of success.

C. The Combination Satisfies the Claims

31. Having reviewed Orca's evidence, it is still my opinion that the references teach all limitations of the claims as recited. Below, I explain why I

disagree with Dr. Kaeli's arguments that the references do not teach certain limitations. For simplicity, I sometimes refer to claim 1 elements, but my analysis applies to the corresponding elements of the other independent claims.

1. Elements 1.5/16.5/19.5

32. I previously explained why the combination teaches or suggests all aspects of element 1.5, which recites "correlating the determined potential cyber vulnerabilities with a network location of the protected virtual cloud asset."

EX1002, ¶¶161-64. Dr. Kaeli makes several arguments in response. Kaeli-Dec, ¶¶107-12. I respectfully disagree with his assessment.

33. I disagree with Dr. Kaeli's response regarding Veselov. He does not dispute that Veselov considers network information, arguing instead that Veselov does not satisfy element 1.5. Kaeli-Dec, ¶108. But I did not address element 1.5 based on Veselov alone. In the arguments he cites, I showed that Veselov considers the same type of information that Mohanty uses in correlating vulnerabilities with a network location of the asset, meaning that the relevant information is available and readily used for that purpose and that Mohanty's teachings are consistent with Veselov's general teaching of considering such information. EX1002, ¶¶162-63. Dr. Kaeli cites Veselov's statement the snapshot may include any properties (including the network information I discussed) that are relevant to the assessment, and he argues that Veselov only teaches using network

information in the 3A/B example. Kaeli-Dec, ¶108. As an initial matter, as I explained above, Dr. Kaeli overlooks Veselov's teachings that the snapshot is the same in both approaches (i.e., using assessment VMs or analyzing the snapshot as a datafile). *Supra*, §VI.B; EX1007, 3:20-4:4, 9:42-10:16, 16:56-67. Even if that were not the case, the applied teachings of Mohanty provide an express reason to use the network information, making the information "relevant to the security assessment" in the combination regardless. EX1002, ¶¶108-15; EX1007, 9:42-10:16.

34. Dr. Kaeli further argues that a POSA would not have thought analyzing the snapshot as a data included determining network location because he reviewed a document (EX2106) indicating that the industry viewed snapshot-based analysis as rudimentary. Kaeli-Dec, ¶109. I disagree with this argument for the same reasons I discussed above. *Supra*, §VI.B. For example, Dr. Kaeli's assertion does not contend with Veselov's discussion of robust snapshot content and comparable findings compared to assessments performed on the target asserts, does not address the state of the field that I discussed in my original declaration, and does not address the fact that the cited document is a marketing blog touting Palo Alto Networks' product over competing products, not an objective industry assessment. *E.g.*, EX1002, ¶¶32-36, 54-58; EX1007, 3:20-4:4, 9:42-10:16; EX2106.

35. I also disagree with Dr. Kaeli's arguments against Mohanty, which go to the reasons for making the combination rather than whether the combination, if made, satisfies element 1.5. Dr. Kaeli refers to his previous arguments based on his assertion that Mohanty is agent-based. Kaeli-Dec, ¶110. I disagree with those arguments for the previously discussed reasons. *Supra*, §VI.B. Dr. Kaeli argues that I did not explain how Mohanty's workload context would apply to Veselov or how Veselov would have the same workload context to provide the recited correlation. Kaeli-Dec, ¶110. As I explained above, however, I showed that Veselov contains the same types of underlying data that Mohanty uses, including information indicating network location to facilitate the correlation. EX1002, ¶¶110, 114-16, 162-63. Dr. Kaeli argues that I did not explain how to use Mohanty's tags in Veselov (Kaeli-Dec, ¶111), but the combination does not require wholesale incorporation of Mohanty's system, tags and all. Again, Veselov's snapshot-based assessment has access to various types of network-location information to facilitate the correlation taught by Mohanty. *Supra*, VI.B; EX1007, 3:20-4:4, 9:42-10:16.

2. Elements 1.6/16.6/19.6

36. Dr. Kaeli largely repeats previous arguments in his brief discussion of element 1.6. He states that his arguments for element 1.5 apply here, but I previously explained why those arguments are incorrect. Kaeli-Dec, ¶114; *supra*,

§VI.C.1. Echoing his previous arguments that agents cannot be installed on snapshots and that Mohanty and Veselov use different data, Dr. Kaeli states that my analysis “ignores the distinction between agent-based analysis in Mohanty and snapshot-based analysis.” *Id.*, ¶115. I explained why this argument is incorrect when discussing the reasons to combine. *Supra*, §VI.B. Dr. Kaeli further argues that one of the background exhibits I cited—EX1078 (Hufsmith)—does not support my conclusions because it “relates to container images,” which is supposedly inapplicable to Veselov’s snapshot analysis. Kaeli-Dec, ¶115. Dr. Kaeli does not explain this assertion, but regardless, it is incorrect for multiple reasons. For example, Hufsmith expressly teaches applicability to VMs and VM images, not just container images, and while containers and VMs are not synonymous, they are closely related virtual resources. EX1078, ¶¶26, 149, 154, 159; EX1007, 1:34-40, 2:46-54. In any event, Hufsmith supports the point for which I cited it, which was that “determining this type of score/risk using determined vulnerabilities and an asset’s network location would have been obvious in view of the references’ teachings, since an asset’s risk was routinely determined using such factors.” EX1002, ¶169.

3. Elements 1.7/16.7/19.7

37. For element 1.7, Dr. Kaeli refers to his previous arguments for elements 1.5 and 1.6. I previously explained why those arguments are incorrect. Kaeli-Dec, ¶114; *supra*, §§VI.C.1-2.

4. Elements 1.8/16.8/19.8

38. I previously explained why the combination teaches or suggests all aspects of element 1.8. EX1002, ¶¶177-83. In response, Dr. Kaeli makes several arguments against the teachings of both Veselov and Mohanty. Kaeli-Dec, ¶¶121-29. I respectfully disagree with his assessment.

39. I disagree with Dr. Kaeli's arguments involving Veselov. He argues that Veselov does not use alerts and implies that Veselov teaches no more than simply communicating results. Kaeli-Dec, ¶¶122-25. But Veselov does not merely describe communicating results generally. As I explained in my previous declaration, Veselov describes reporting identified vulnerabilities in a way that brings specific problems (and identified remediations) to the user's attention. EX1002, ¶178. For example, it describes reporting information that describes the vulnerability, identifies a remediation, and conveys this information "in a user-readable format to a user interface to be displayed on the user device." EX1007, 10:17-36, 20:57-21:13, 23:19-47. This is what Dr. Kaeli says an "alert" does. Kaeli-Dec, ¶124 (alerts "inform stakeholders of a situation that needs attention").

Dr. Kaeli's argument also fails to consider the perspective of the POSA in reading these descriptions in Veselov. As I previously explained, reporting vulnerabilities as alerts was commonplace in security platforms (EX1002, ¶¶64-68), and a POSA would have recognized that Veselov's descriptions that I noted above are consistent with this well-known, commonplace technique.

40. I also disagree with Dr. Kaeli's arguments addressing Mohanty. Dr. Kaeli repeats arguments that he made against the reasons to combine, arguing that a POSA would have understood that Veselov "should not issue 'alerts' on the vulnerabilities that it detects" because Mohanty uses agents to perform continuous monitoring, whereas Veselov uses snapshots that will eventually become obsolete. Kaeli-Dec, ¶¶127-28. But Veselov already teaches alerts, and regardless, I previously explained why a POSA would have understood prioritized alerts to be beneficial to Veselov's snapshot-based assessment. EX1002, ¶¶113-15. The fact that snapshots represent system state at a given time does not make context-sensitive prioritization irrelevant, and it was well known that such prioritization was useful for assessments based on images/snapshots. *Id.*, 64-69; EX1048 (Price), ¶43 (prioritization for VM image-based assessment); EX1078 (Hufsmith), ¶¶180-83 (prioritization in container image example); *see also id.*, ¶¶149, 154, 159 (applicability to VMs/VM images). Moreover, as I previously explained, a POSA would have recognized that Veselov's approach is reasonably current since

Veselov teaches using fresh snapshots taken for on-demand, triggered, or periodic assessments. *Supra*, §VI.B.

41. Dr. Kaeli refers to his later discussion of secondary considerations, which he says shows that the combination relies on hindsight. Kaeli-Dec, ¶128. In particular, he argues that Orca’s targeted vulnerability alerts were praised. *Id.* As I explain below, such alerts were well known, including for agentless scanning (*infra*, §IX.B.1), and I disagree that the evidence indicates that the combination’s application of Mohanty’s prioritized-alert teachings relies on hindsight. *Infra*, §IX; *see also infra*, §IX.B.4 (discussing alleged praise).

5. Claims 15/22

42. I previously explained why the combination of Veselov and Mohanty teaches or suggests all aspects of claims 15 and 22. EX1002, ¶¶245-54. These claims recite that determining the location of the snapshot includes “querying a cloud management console of the cloud computing environment for the location of the snapshot and the location of the virtual disk.” Dr. Kaeli makes several arguments in response (Kaeli-Dec, ¶¶132-40), but I respectfully disagree with his assessment.

43. Addressing the “cloud” aspect, he argues that “computing resource service provider” does not typically refer to a cloud provider. Kaeli-Dec, ¶133. However, it would have been perfectly clear to a POSA that Veselov is describing

a cloud provider. EX1007, 2:22-3:19 (techniques “particularly suited for security assessments performed in a data center” and address concerns of “cloud” environments); *see also* EX1118 (Call), 4:32-41; EX1119 (Medeiros), 2:41-52; EX1120 (Kumar), 1:10-14, 1:64-2:5 (using “computing resource service provider” to refer to cloud provider).

44. I also disagree with Dr. Kaeli’s argument that Figures 5A/B of Veselov do not involve the service provider or APIs. Kaeli-Dec, ¶133. Veselov identifies the service provider as an undepicted aspect of the computing environment illustrated in Figs. 5A/B when, in discussing operations by scanning service 510 involving volume image 508, it refers to using a “security module or security service (not shown) of the computing resource service provider.” EX1007, 18:1-15. That the 5A/B environment is hosted by a cloud provider (and therefore uses the same APIs Dr. Kaeli acknowledges are associated with this provider) is also made clear by Veselov’s express description of that environment. The 3A/B environments are “implemented by a computing resource service provider”—which Dr. Kaeli acknowledges (Kaeli-Dec, ¶133)—and Veselov states that the 5A/B environments has the same “implementation” as the 3A/B environment. EX1007, 10:37-41; 16:56-67. Veselov also states that the 5A/B environments have the same “interactions” as the 3A/B environments (*id.*, 16:56-67), which Dr. Kaeli similarly acknowledges uses API interactions. Kaeli-Dec,

¶133. I explained this in my previous declaration (EX1002, ¶246), and Dr. Kaeli does not appear to respond to that analysis.

45. I also disagree with Dr. Kaeli's argument that Veselov does not teach using APIs to obtain snapshots. He argues that this is the case because, in Figs. 3A/B, 4, and 5A/B, the volume image is provided to a hypervisor or service "that launches an assessment machine." Kaeli-Dec, ¶133 (citing EX1007, 12:37-51; 18:16-21). As an initial matter, it is not clear what he is referring to with respect to 18:16-21 of Veselov, which does not mention launching an assessment machine. In any event, the use of hypervisors and provisioning services to manage virtual resources does not imply that cloud APIs are not used. To the contrary, the publicly exposed cloud APIs operate in parallel with the hypervisors and provisioning services that are part of the cloud platform's native resource-management layer. Dr. Kaeli argues that APIs are not "needed" to obtain snapshots in Veselov (Kaeli-Dec, ¶133), but I did not provide an inherency argument, and I understand that this is not a requirement for obviousness.

46. Regardless, Dr. Kaeli is incorrect in his assertion that a POSA would not have viewed Veselov as using APIs to obtain snapshots. Kaeli-Dec, ¶133. As I previously explained, Veselov broadly states that the scanning service may interface with the provider or any associated environments/resources via API 120 (the environment API) "or another API." EX1007, 5:1-5; EX1002, ¶246. That

broad description would have made sense to a POSA because it was common and well-known that the same cloud APIs were provided for multiple types of interfaces—including Web interfaces, command-line interfaces, and programmatic interfaces using a software-development kit (or “SDK”)—with the cloud-provider customer using whichever type of interface was appropriate based on the customer’s needs. In other words, it was well known that each of these interfaces would use cloud APIs for communication with the cloud provider/environment. I explained the POSA’s understanding of how cloud APIs are used in these well-known types of interfaces and how Veselov’s disclosure therefore suggests using a cloud management console (EX1002, ¶247), and Dr. Kaeli does not appear to dispute to that explanation.

47. With respect to the use of such a console to query the cloud environment for “the location of the snapshot and the location of the virtual disk,” Dr. Kaeli does not appear to dispute the snapshot location aspect. For the “location of the virtual disk” aspect, I explained multiple independent ways that Veselov teaches this feature. First, a POSA would have understood that querying the snapshot’s general location typically gives the disk’s general location, since they would often have the same general location. EX1002, ¶250. Second, I explained how this was taught or suggested by Veselov’s discussion of using an API-based interface to display and control configuration parameters of the security

assessment, such as the choice of the target, including target virtual disks. *Id.*, ¶¶251-52. Dr. Kaeli primarily responds to the first argument and provides only a brief comment on my analysis for the second. Kaeli-Dec, ¶¶134-36. I address each in turn below.

48. I disagree with Dr. Kaeli's arguments with respect to the general location. He states that even if it is true that the snapshot and virtual disk sometimes have the same general location, Veselov still does not suggest determining the virtual disk's location because its snapshots are supposedly stored in and obtained from a snapshot storage location. Kaeli-Dec, ¶134. But this does not address Veselov's full teachings because Veselov makes clear that the snapshot storage service is optional (and I did not rely on that optional aspect). EX1007, 5:9-22, 9:9-25. Dr. Kaeli argues that the snapshot and virtual disk do not "necessarily" have the same general location. Kaeli-Dec, ¶134. Again, however, my argument did not rely on this "necessarily" being the case in every circumstance. Rather, I discussed the POSA's understanding of typical operations. EX1002, ¶¶232, 250. Dr. Kaeli also argues that I did not explain how the combination "would *benefit* from such information." Kaeli-Dec, ¶134. He appears to assume that my analysis for this limitation relied on motivation to modify Veselov, which is not the case. Rather, I explained why this limitation is suggested by Veselov as read from the POSA's perspective. EX1002, ¶¶250-53.

49. I also disagree with Dr. Kaeli's argument that my analysis is "belie[d]" by Veselov's system use of VM identification. Kaeli-Dec, ¶135. In particular, Dr. Kaeli argues that Veselov simply identifies the VM and "has no need" to locate the virtual disk. *Id.*; *see also id.*, ¶136 (VM "identification means the user does not need to query for a location"). Again, however, this does not address Veselov's full disclosure (or my analysis of that disclosure). First, the initial scan may not even identify the VM. EX1007 (Veselov), 6:34-38 ("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." (emphasis added)). Regardless, even if the request does identify a particular VM, Veselov teaches determining which data elements to include in snapshots (*id.*, 9:25-28), and a POSA would have been well aware that a virtual disk was typically the core element of a snapshot. *See, e.g.*, EX1121 (AWS snapshot tied to particular volume). Dr. Kaeli is therefore incorrect that a POSA would have viewed VM identification as obviating the relevance of virtual-disk location. Indeed, in AWS, for example, you cannot generate a snapshot via the standard API call by providing only a VM identifier. *Id.* Veselov expressly discloses using an interface to select an assessment target, which may be a virtual disk. EX1007, 8:52-9:8 (target may be a "logical volume, physical or virtual storage device"), 14:37-63 (API/interface-

based assessment configuration may include “selection of the assessment target(s); parameters of the security assessment(s)”; *see also id.*, 16:54-67 (environment and user device of Figs. 5A/B have same “implementation and interactions” as those of Figs. 3A/B). Additionally, Dr. Kaeli argues that VMs and their virtual disks “can be located apart” (Kaeli-Dec, ¶136), but a POSA would have understood that they typically have the same general location in a cloud environment. For example, AWS, like other major platforms, requires VMs and their associated volumes to be in the same “region” and “availability zone” (which generally meant being the same data center or group of data centers). EX1133⁸ (“The volume and instance must be in the same Availability Zone.”).

50. I also disagree with Dr. Kaeli’s arguments addressing several background exhibits. Dr. Kaeli argues that I did not explain the relevance of EX1015 (VMware vSphere) to Veselov. Kaeli-Dec, ¶137. But I cited this in support of my explanation of how the POSA’s familiarity with the relevant technology—in this case, cloud-management-console querying being a common

⁸ EX1133 is a correct copy of AWS documentation obtained from <https://docs.aws.amazon.com/ebs/latest/userguide/ebs-volumes.html>. Based on my experience in the field, the cited description of the volume/instance co-location requirement is consistent with the state of the art in January 2019.

aspect of the type of VM management discussed in Veselov—informs the POSA’s understanding of Veselov’s teachings. EX1002, ¶¶248-49, 252. Dr. Kaeli argues that EX1031 (NIST IT Asset Mgmt) only refers to networked devices and software, not virtual disks of VMs, (Kaeli-Dec, ¶137), but it provides a broad disclosure of computer security principles and expressly addresses virtual assets. EX1031, 1-2 (Executive Summary), ii (pdf page 9) (addressing “physical **and virtual assets**” (emphasis added)); *see also id.*, 40 (auditing/monitoring processes include “data storage” and “reporting/visualization”).

51. I also disagree with Dr. Kaeli’s arguments addressing the background exhibits from my cross-referenced discussion of claims 10 and 12. In discussing claims 15 and 22, I referred to my discussion of claim 12 involving the determination of virtual-disk location. EX1002, ¶¶231-33, 250. Dr. Kaeli argues that this evidence would not have “motivated” a POSA “to change Veselov’s systems and methods to identify a location of a virtual disk, as claimed, as opposed to locating a snapshot or identifying a VM generally.” Kaeli-Dec, ¶¶139-40. Again, however, Dr. Kaeli assumes that my analysis for this limitation relied on motivation to modify Veselov, which is not the case. Rather, I explained what a POSA—informed by their understanding of the relevant technology—would have understood Veselov’s disclosure to involve or at least suggest. In other words, I understand that references are read from the POSA’s perspective, and I discussed

these background exhibits to illustrate the relevant knowledge that would have informed the POSA's understanding of Veselov. Dr. Kaeli argues that my claim 10 analysis does not address his previous arguments (Kaeli-Dec, ¶138), but he does not substantively dispute my actual point regarding the relevance of my claim 10 analysis. EX1002, ¶253.

52. Lastly, I disagree with Dr. Kaeli's argument that the background exhibits "do not demonstrate that the claimed features are routine or commonplace." Kaeli-Dec, ¶140 n.7. He dismisses several of these exhibits because they are not grounds references in the Veselov-Mohanty combination, but as I explained above, I did not use these exhibits as grounds references or to modify Veselov; rather, I used them to illustrate the POSA's relevant knowledge that would have informed their understanding of Veselov's and Mohanty's teachings. Dr. Kaeli repeats the same argument involving EX1015 (VMware vSphere) that I addressed above. Dr. Kaeli argues that EX1051 (Workstation Pro) shows only that snapshots include disk data, not that disk location is determined. However, it discusses drive selection as part of snapshot-configuration tools. EX1051 (Workstation Pro), 125. Dr. Kaeli argues that EX1052 (VirtualCenter) only addresses locating VMs, not virtual disks. However, it discusses UI-based identification and display of virtual disks, including address information. EX1052 (VirtualCenter), 445-46. Dr. Kaeli similarly argues that EX1053 (Prahlad) and

EX1020 (Sancheti) do not address determining virtual-disk locations. However, EX1053 (Prahlad) provides an extensive discussion of various types of snapshot generation that include determining volumes/disks, and it expressly discusses virtual-disk addresses. EX1053 (Prahlad), ¶¶77-79 (snapshot generation includes file-level, image-level, and disk-level copy), 81-84 (snapshot generation involves “determine[ing] the volumes of the virtual machine 110”), 87-88 (similar teachings for volume-level copy), 92 (similar teachings for disk-level copy), Fig. 7. EX1020 (Sancheti) similarly discusses “discover[ing] and identif[ying] the virtual disks ... associated with the virtual machine.” EX1020 (Sancheti), 21:9-41. For EX1074 (VMware Infrastructure), Dr. Kaeli argues that I did not explain how its disclosure “would be incorporated with Veselov.” However, it is evidence corroborating the POSA’s relevant knowledge, not a disclosure that is “incorporated” into Veselov.

53. It is therefore still my opinion that claims 15 and 22 were obvious.

VII. Ground 2

54. In Ground 2, I explained the reasons for applying certain teachings in Czarny with the previous combination of Veselov and Mohanty, and I explained additional ways that this combination teaches aspects of element 1.8, claim 6/17, and claim 7/18 (beyond those I already addressed in Ground 1). EX1002, ¶¶259-79. I respectfully disagree with Dr. Kaeli’s arguments addressing this ground.

A. The Further Combination with Czarny

55. I disagree with Dr. Kaeli's arguments involving the nature of the data considered in Czarny. Dr. Kaeli contrasts security assessments of VMs (as in Veselov and Mohanty) with what he describes as Czarny's focus: security assessments of "electronic devices." Kaeli-Dec, ¶144. Dr. Kaeli does not explain how this supposed distinction relates to the reasons for combining the references' teachings, nor does he appear to dispute my actual arguments about why the references are analogous. EX1002, ¶261. Dr. Kaeli further argues that neither Veselov or Mohanty use "binary data" as in Czarny. Kaeli-Dec, ¶145. But virtual assets contain the same types of underlying data as physical assets (i.e., information stored as binary data). EX1002, ¶¶24-29. Veselov's snapshot captures disk state (which is binary data). *Id.*, ¶¶32-36; EX1007, 3:20-40, 9:42-67. This is basic computer science knowledge that a POSA would have known. Veselov and Mohanty do not expressly describe the underlying mechanism by which the CVE assessments are performed, and I explained why a POSA had reason to use the standard types of matching techniques described in Czarny. EX1002, ¶¶264, 266. That snapshot-based assessments operate in a virtual-machine context does not make these commonplace matching techniques any less applicable, nor does it meaningfully change how a POSA would use these techniques in CVE assessments.

56. I also disagree with Dr. Kaeli's arguments involving risk level. Dr. Kaeli argues that my discussion in paragraph 265 does not address motivation for applying Czarny's risk-level-reporting teachings. Kaeli-Dec, ¶146. But Dr. Kaeli does not dispute my points in that paragraph, nor does he address my express discussion of motivation for applying those teachings later in that section of my declaration. EX1002, ¶267. He argues that I did not explain why a POSA would "abandon" Mohanty's teachings of reporting risk levels "in favor of Czarny's teachings." Kaeli-Dec, ¶146. But I did not argue that a POSA would "abandon" Mohanty's teachings of reporting risk levels. To the contrary, I explained that Czarny's teachings are consistent with Mohanty's (and the Veselov-Mohanty combination). EX1002, ¶¶259, 265. To be clear, Mohanty does teach reporting risk levels, which I explained in my previous declaration (and which Dr. Kaeli does not appear to dispute). EX1002, ¶180. Czarny's disclosure is simply more explicit in some ways with respect to reporting the risk levels themselves, making the relevant claim limitations even more obvious. *Id.*, ¶¶267, 270. Dr. Kaeli argues that I did not establish how the combination would work (Kaeli-Dec, ¶146), but I explained that the combination uses these commonplace matching techniques according to their known, predictable functions to achieve the same, predictable results. EX1002, ¶¶262-65, 268. This straightforward combination is a square peg in a square hole. Aside from meritless distinctions involving "virtual" assets and

“binary” data, Dr. Kaeli does not identify any aspect of the combined teachings that he thinks do not “fit together.” Kaeli-Dec, ¶146.

B. Element 1.8

57. I disagree with Dr. Kaeli’s arguments for element 1.8. He argues that I did not explain why Czarny’s risk level is a risk of the asset to the “cloud” environment, as recited. Kaeli-Dec, ¶148. But Dr. Kaeli does not address the *combination*, which applies Czarny’s teachings of reporting determined risk levels to the previous combination of Veselov and Mohanty, which teaches performing the assessment in a cloud environment (and therefore determining risk of the asset to the cloud environment). EX1002, ¶¶118-21 (cloud), 165-71 (risk level to “cloud” environment), 259 (Ground 1 analysis applies to Ground 2); *see also supra*, §VI.A (citing EX1007, 2:22-3:19; EX1118, 4:32-41; EX1119, 2:41-52; EX1120, 1:10-14, 1:64-2:5). Dr. Kaeli alleges that there is still a gap when viewing Czarny in combination with Veselov and Mohanty. Kaeli-Dec, ¶149. In particular, he argues that the combination does not teach “alerts,” alerts that are “prioritized,” or a risk level to the cloud computing environment. *Id.* But Dr. Kaeli does not provide any explanation for this assertion. In any event, his assertion is incorrect, and I previously explained why his arguments for Ground 1 involving elements 1.8 are incorrect. *Supra*, §VI.C.4.

A. Claims 6-7 and 17-18

58. Dr. Kaeli does not dispute my analysis showing that the combination satisfies the additional limitations of claims 6-7 and 17-18 beyond referring to his arguments for the independent claims. Kaeli-Dec, ¶150. I disagree with those arguments as discussed above. *Supra*, §VII.A-B.

VIII. Grounds 3-4

59. In Grounds 3-4, I explained how dependent claim 11 was obvious further in view of Hutchins (with or without Czarny). EX1002, ¶¶280-89. Claim 11 further recites “wherein mitigating a potential cyber vulnerability includes at least one of: blocking traffic from untrusted networks to the protected virtual cloud asset, halting operation of the protected virtual cloud asset, or quarantining the protected virtual cloud asset.” In particular, I addressed Hutchins’ teaching of the commonplace mitigation technique of quarantining. Dr. Kaeli presents several arguments for Ground 3 and mostly relies on the same arguments for Ground 4. As I explain below, I disagree with Dr. Kaeli’s assessment.

60. I disagree with Dr. Kaeli’s arguments against the combinations. He argues that using Hutchins’ “resource-intensive” analysis techniques like virus scanning “would require additional resources beyond those provided for in Veselov.” Kaeli-Dec, ¶¶155-56; *see also id.*, ¶162 (referring to previous arguments for Ground 4). But the combination does not involve the incorporation of Hutchins’ analysis step. Rather, it applies Hutchins’ teachings of the

straightforward, well-known mitigation technique of quarantining VMs. EX1002, ¶¶282-86. Regardless, a POSA would not have viewed virus scanning as being too resource-intensive for Veselov's scanning service. Veselov encourages the use of known assessment tasks, and virus scanning was commonplace. EX1007, 2:35-41. Dr. Kaeli argues that a POSA would have been "led away" from combining Hutchins with Veselov's direct-snapshot-analysis approach (discussed for example in the embodiments associated with Veselov's Figures 5A/B) because Hutchins uses snapshots to instantiate duplicate VMs, which is more similar to Veselov's Figures 3A/B. But this again assumes wholesale incorporation of Hutchins' analysis step rather than addressing the combinations' straightforward application of the well-known quarantining remediation. Regardless, Dr. Kaeli's narrow characterization of Hutchins is misleading since Hutchins broadly describes "running a scan operation on the snapshot" and then quarantining the original VM. EX1070, Abstract, claims 1-2. The use of a duplicate VM is an example, and even then, Hutchins scans the attached disk image. *Id.*, 7:3-5. Virus scanning was also known in the context of assessments that directly analyze snapshot/image data. *See, e.g.*, EX1129 (Michael), 2:28-48, 7:47-62. In any event, Dr. Kaeli's argument is beside the point because the combination does not incorporate Hutchins' analysis step and because the use of quarantining as a remediation does not depend on which of Veselov's analysis approaches is used.

61. For Ground 4, Dr. Kaeli quotes a statement in EX1084 (Czarny) that computers can “usually” continue to operate after a positive vulnerability assessment. Kaeli-Dec, ¶163 (citing EX1084, 1:8:15). But Dr. Kaeli omits Czarny’s subsequent teaching that quarantining *is* the appropriate response to vulnerability detection in some circumstances (e.g., when the presence of vulnerabilities violates the security standards of the asset’s environment). EX1084, 1:5-35. Nor does Dr. Kaeli appear to dispute my prior explanation that quarantining was well known as a remediation technique in response to detecting vulnerabilities. EX1002, ¶¶70, 283-85; EX1048 (Price), ¶14; EX1081 (Budko), ¶79 (“A detective control may require vulnerability scanning of the VMs at rest to detect vulnerabilities that were recently exposed in the operating system or applications in the VM. A corrective control may require that the VM be quarantined and remediated before being allowed to be activated in a particular environment.”).

62. I also disagree with Dr. Kaeli’s arguments that the combinations do not satisfy claim 11. He argues that Hutchins does not mention “vulnerabilities” and that its quarantining is a response to detecting a virus, which Dr. Kaeli claims is not a vulnerability. Kaeli-Dec, ¶158. In particular, Dr. Kaeli argues that EX1084 (Czarny) describes malware detection and vulnerability assessments as different fields/industries. Kaeli-Dec, ¶158 (citing EX1084, 1:6-7). But Czarny

notes overlapping security concerns for what it characterizes as malware and vulnerabilities. EX1084, 1:5-35. Regardless, the claims do not exclude malware, which was often considered a vulnerability. *See, e.g.*, EX1078 (Hufsmith), ¶93 (discussing “type of vulnerability, such as malware”).

IX. Secondary Considerations

63. Having reviewed the opinions of Dr. Kaeli and Mr. Shua regarding secondary considerations, it is still my opinion that all claims were obvious in view of the grounds that discussed in my original declaration. Below, I discuss my assessment of Orca’s secondary-considerations evidence and explain why that evidence does not support nonobviousness.

64. I was not aware of any secondary-considerations evidence that favored nonobviousness when I prepared my previous declaration. Having now reviewed the secondary-considerations evidence that Dr. Kaeli and Mr. Shua argue supports nonobviousness, I respectfully disagree with their assessment. As I explain in more detail in the following sections, I do not believe this evidence favors nonobviousness.

65. Before turning to the specifics of my assessment, I will summarize the legal principles that I applied in reaching that determination. First, I continued to apply the legal principles involving secondary considerations that I discussed in my previous declaration. EX1002, ¶¶17-19. With respect to nexus, I have been

advised that a patentee can establish a rebuttable presumption of nexus for evidence tied to a given product if the patentee proves that the product (1) embodies the claims and (2) is coextensive with the claims. I discuss this issue in more detail below in section IX.A. I have been advised that if the presumption of nexus is established, the presumed nexus for a given piece of evidence is rebutted if that evidence relates exclusively to known and/or unrecited. I have been advised that without that presumption, nexus requires a patentee to show that the secondary consideration (e.g., the solution to a long-felt, skepticism, praise, copying, or commercial success) is the direct result of the claims' unique characteristics. That is, it must be the result of something that is both claimed and novel. I have also been advised that, even when nexus is shown, it exists on a spectrum. For example, there may be a strong nexus if there is a clear, direct connection between the evidence and claimed/novel features, or there may be a weaker nexus if there is a less direct connection to claimed/novel features.

66. As I explain in the following sections, it is my opinion that product/platform (which Dr. Kaeli refers to as SideScanning™) is not entitled to the presumption of nexus. *Infra*, §IX.A. Even if the product were given this presumption, it would still be my opinion that the evidence lacks nexus because it is merely directed to known features (some of which are also unrecited). Since much the evidence relates to the same handful of features, I provide a summary of

those known features upfront (*infra*, §IX.B.1) so that I can avoid unnecessary repetition when I address the evidence for each type of secondary consideration discussed in Dr. Kaeli's declaration (*infra*, §§IX.B.2-6).

67. Lastly, while I do not believe the evidence has nexus to the claims for multiple reasons, even if it did, it still would not change my ultimate opinion that the claims were obvious. In other words, even assuming nexus, Orca's evidence would not outweigh the very strong evidence of obviousness that I discussed in my original declaration.

A. No Presumption of Nexus

68. As I noted above, I have been advised that a patent can prove entitlement to a presumption of nexus for a product if the patentee proves that the product (1) embodies the claims and (2) is coextensive with the claims. In other words, the product must not only practice the claims but also be reasonably coextensive with them. I have been advised that a product is not coextensive with the claims if the product has significant features that are not recited in the claims. In my opinion, the evidence does not show that Orca's product embodies and is coextensive with the claims.

69. In my opinion, the platform is not coextensive with the claims. Neither Dr. Kaeli nor Mr. Shua contends that it is, so there does not appear to be any dispute on this point. In the interest of clarity, however, I have reviewed

Orca's exhibits describing the product, and in my opinion, the product is not coextensive with the claims since it contains multiple features that are both significant and unrecited. I discuss several of these below.

70. First, Orca's evidence emphasizes the significance of automatically enumerating "every asset" and building an infrastructure map, which is not recited in the claims. EX2067, 8. The process starts with "building a map of your organization's entire estate." *Id.* This covers "every asset in the account" and, as part of the "control plane path," the system "builds an infrastructure map that acts as a guide for its assessment and analysis processes." *Id.*, 8-9. A POSA would have understood this to be a significant feature of the product. Orca's technical brief emphasizes this aspect in its very first page.

The cloud estate of an organization comprises its complete inventory of cloud assets including running and stopped workloads of all types including VMs, containers, storage objects, load balancers, IAM configurations, and more. Organizations are searching for effective ways to scan their cloud estate to look for risks stemming from vulnerabilities, misconfigurations, malware, lateral movement risk, weak and leaked passwords, and improperly secured PII.

Orca Security introduces a radically new approach that secures the entire cloud estate and helps meet compliance mandates, but without disrupting business operations in live environments.

EX2067, 1; *see also id.*, 5 (platform "covers 100% of an organization's assets" and "automatically assesses the security state of every discovered asset"). Likewise, the subsequent description identifies this is part of "one of the most valuable features" that the platform provides. *Id.*, 8. In contrast to the platform's features,

the claims locate and analyze a snapshot of one or more virtual disks of a single virtual asset, meaning they capture only a tiny subset of this cloud estate.

Moreover, the claims focus on a snapshot, whereas Orca emphasizes coverage of many types of resources beyond snapshots. *Id.*, 1 (discussing “containers, storage objects, load balancers, IAM configurations, and more”). Orca includes VMs/virtual disks in just one of 17 different categories of assets. *Id.*, 8 (listing “EC2 instances, volumes, snapshots” as one of 17 asset categories). The platform also “provides definitive alerts” during the control plane path, which is before the “data plane path” where snapshots are created for vulnerability detection. *Id.*, 9-10. By contrast, the claims report vulnerabilities detected by analyzing the snapshot.

71. Numerous aspects of the data plane path are also unrecited. The claims recite analyzing the snapshot by matching installed applications with applications on a known list, then determining “cyber vulnerabilities” based on the matching. Yet Orca states that the platform’s data-plane path performs numerous assessment steps beyond just detecting vulnerabilities. In the technical brief, “Vulnerability Scanning” is one of seven listed steps:

The collector successively mounts each snapshot, then immediately deletes it so customer billing isn't affected. Next, the collector starts reorganizing the volumes. It mounts them in the same way as the scanned OS would've done and runs several data collection steps:

- Vulnerability Scanning
- Configuration Scanning
- Malware Scanning
- Lateral Movement Risk
- Exploitable Keys and Weak Password Detection
- Sensitive Information Scanning
- Container Scanning

EX2067, 10. [REDACTED]

[REDACTED]

[REDACTED]

72. Beyond the significant unrecited features I noted above, it is difficult to know how many more significant features the platform contains, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

73. While I understand that this lack of coextensiveness means that the presumption of nexus has not been established, I also briefly address the additional

nexus-presumption requirement that the patentee prove that the product embodies the claims. Dr. Kaeli adopts Mr. Shua's testimony arguing that Orca's platform practices the claims. Kaeli-Dec, ¶167. I note several issues with the adopted testimony below.

74. First, I note an issue with element 1.6. Element 1.4 recites determining "cyber vulnerabilities," and element 1.6 recites "using" the vulnerabilities and the network location to determine a "risk" of the protected virtual cloud asset to the cloud computing environment. Orca identifies "lateral movement risk" as the "risk" for element 1.6. POR, 36. Yet Mr. Shua's discussion of lateral movement risk ("LMR") does not explain how this LMR is determining "using the determined plurality of potential cyber vulnerabilities" as required by element 1.6. Shua-Dec, ¶¶49-50. For "vulnerabilities," he points to the platforms' use of a vulnerability database to detect vulnerabilities, citing page 11 of EX2067 (SideScanning Technical Brief 2020). Shua-Dec, ¶¶44-45. Then for element 1.6 he alleges that LMR is determining using the vulnerabilities and network location, citing pages 14-15 of EX2067 (SideScanning Technical Brief 2020). *Id.* Even considering page 13—which is where the LMR language he quotes appears to come from—the Technical Brief's discussion of LMR does not say that LMR is determined using the vulnerabilities that Mr. Shua cited when discussing element 1.4. The Technical Brief describes looking for user credentials

(e.g., passwords and keys) as part of LMR detection, but these are not detected vulnerabilities that Mr. Shua relied on for element 1.4. EX2067, 13-14. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] The cited Technical Brief also describes “Vulnerability Scanning” and “Lateral Movement Risk” as separate data collection steps of the data-plane path. EX1067, 10. Mr. Shua further cites page 16 of EX1067, but this does not discuss what is used to determine LMR. Moreover, the type of risk discussed here (e.g., considering a vulnerability in a web service to be “[h]igh risk” if it is “internet-facing”) is not the LMR that Orca identifies as the “risk” for purposes of element 1.6 in the Patent Owner Response. EX2067, 16; POR, 36.

75. I note further inconsistencies for element 1.7 (prioritizing “by the determined risk”) and element 1.8 (reporting vulnerabilities as alerts prioritized “according to the determined risk”). These elements refer to the determined “risk” of element 1.6, and as I explained above, Orca maps this to LMR. POR, 36. Yet for element 1.7, Mr. Shua does not explain how the platform prioritizes the determined vulnerabilities by the LMR. Shua-Dec, ¶¶51-52. Instead, he cites the risk determination discussed on page 16 of EX2067. *Id.* As I explained above, the cited discussion does not mention prioritizing vulnerabilities by LMR. EX2067,

16. Mr. Shua's discussion of element 1.8 has a similar inconsistency. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] But again, this does not show prioritization according to the determined "risk" that Orca relied on for element 1.6 (i.e., LMR).

B. Orca's Evidence

1. Known Techniques/Features

76. Throughout his discussion of secondary considerations, Dr. Kaeli cites various documents referencing features that this evidence supports nonobvious. I have reviewed these citations, and in my opinion, all referenced aspects of Orca's platform were known prior to January 2019. Since many of these features are mentioned repeatedly throughout Dr. Kaeli's declaration, I discuss them below and will then cite back to this discussion to avoid unnecessary repetition.

Agentless

77. Agentless workload scanning was known. I discussed Veselov's disclosure of this technique and various other examples at length in my prior declaration, so I will not repeat that full discussion here. EX1002, ¶¶54-58, 62, 84-89. In brief, this technique was well known in the field before January 2019. Analogous agentless techniques were known for different types of assets (e.g.,

VMs and containers)—including offline assets—and for the data structures representing the state of those assets (e.g., VM images, snapshots, container images). *E.g.*, EX1007 (Veselov), Abstract, 1:22-49, 2:27-35; EX1040 (Ammons), EX1040 (Ammons), Abstract, ¶43; EX1048 (Price), ¶¶14, 32; EX1056 (Chari), ¶94; EX1065 (Thakar), 2:29-33; EX1073 (Mirage), 1; EX1078 (Hufsmith), ¶149; EX1129 (Michael), 2:28-48, 7:47-62 (scanning snapshot/image checkpoint allows malware assessment “without having an anti-malware agent running inside the virtual machine”). It was also known that the results of these types of agentless scanning were comparable to those of agent-based scans. *E.g.*, EX1129 (Michael), 2:28-48 (“Scanning the exposed physical hard disks is similar to, and more effective than, scanning the file-system data from within the virtual machine if an anti-malware agent were to be deployed inside the operating system.”), 7:47-62. Veselov itself makes this clear when it provides an overview of its process—including alternative approaches to performing the analysis step (using a duplicate assessment VM or analyzing the snapshot as a data file)—and explains that “[i]n any case, the snapshot is 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.” EX1007 (Veselov), 2:42-4:4.

Cloud APIs

78. Using cloud APIs to facilitate agentless workload scanning was also known. As I explained above, Veselov broadly describes the scanning service interacting with the cloud provider via APIs. EX1007, 5:1-9. This was nothing new since API-mediated interaction had long been a core mechanism for communication between computer components. Likewise, API-based resource discovery—including virtual-disk identification/location—was a known part of snapshot creation. *See, e.g.*, EX1020 (Sancheti), 21:9-43 (determining virtual disks for snapshots by “discover[ing] and identif[ing] the virtual disks ... associated with the virtual machine, for example by using APIs to query the virtual machine host”); EX1053 (Pahlad), ¶¶79-94. Cloud-API-mediated interactions were also a fundamental aspect of cloud systems, including cloud-security platforms.

79. The particular API calls emphasized by Dr. Kaeli were also known.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

9

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

80. Indeed, cloud-API-based agentless workload scanning is squarely taught by Veselov itself. Mr. Shua states that Veselov is fundamentally different from his approach with respect to cloud APIs, resource discovery and snapshot analysis (Shua-Dec, ¶¶152-59), but that is not how a POSA would have understood Veselov. Rather, a POSA would have viewed Veselov as encompassing cloud-API-based resource discovery and snapshot creation/analysis that are comparable to the aspects of Orca's product that Mr. Shua says satisfy the claims. *Id.*, ¶¶31-46. For example, while Veselov is not limited to AWS, a POSA would have understood that Veselov encompasses that context since Amazon was the largest cloud provider (and Veselov is an Amazon patent). In any event, Veselov broadly describes cloud-API-mediated interactions and creating snapshots (not just obtaining previously generated snapshots). EX1007, 5:1-9, 9:14-41. The request initiating the scan may identify a particular VM, but even in that case, Veselov teaches having the scanning service identify the data to include in the snapshot. *Id.*, 9:25-28. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Regardless, Veselov's initial requests from the user are not limited to those that

identify the VM. For example, the request may simply identify the user. *Id.*, 6:34-38. In either case, a POSA would have understood Veselov as suggesting standard resource-discovery techniques so that it can create the snapshot(s). [REDACTED]

[REDACTED]

[REDACTED]

Read-Only Access and Block-Level Analysis

81. Dr. Kaeli cites documents mentioning Orca’s read-only access and block-level analysis, but these were known. Read-only access is nothing new, and it reflects the well-known security concept referred to as the principle of least privilege (i.e., giving users/systems the minimal access/permissions needed to perform specific functions). For example, Veselov describes read-only mounting to avoid inadvertent modifications. EX1007, 18:1-29. Veselov also describes “block-level” analysis. EX1007, 2:27-35 (technique suitable for assessment of “block-level storage devices” in data center), 3:27-40 (snapshot may include “block-level” image), 11:50-55, 16:21-26, 18:21-29. This, too, is nothing special since disk snapshots typically include block-level representations of the disk’s data.

Prioritized Alerts

82. Dr. Kaeli points to documents mentioning various aspects of alerts, including prioritized alerts, context-sensitive alerts, and surfacing actionable alerts.

As I previously explained, however, such alerts were commonplace long before 2019 (EX1002, ¶¶64-70), and they were a known aspect of agentless workload scanning (not just agent-based scanning). *E.g.*, EX1048 (Price), ¶¶17 (results provided in various ways, including “alerts”), 41-45 (discussing assessment of image and informing user of detected vulnerability along with priority level and recommended resolution/action); EX1078 (Hufsmith), ¶¶149 (teachings applicable to container images and other workload resources), 153-55, 159 (teachings described for “containers, images, and files,” and “similar approaches are contemplated for virtual machines, orchestration configuration files, serverless configuration files, and the like”), 180-183, 196-97. It was also known to prioritize based on a wide variety of contextual factors. *E.g.*, EX1044 (Giakouminakis), 6:9-58, 8:3-15 (considering “intrinsic value of compromised data due to an attack on the asset”), 8:47-9:3, 9:19-29 (considering risk level of connected assets and “the level of connection between the entities and the asset”), 11:3-10 (considering “function and the type of data stored”), 11:51-67; EX1063 (Hibbert), ¶¶80, 100-13 (considering application use); EX1075 (Mohanty), 5:16-58 (considering “workload context”), 5:59-6:44 (prioritization score based on vulnerability/threat/workload scores and other factors); EX1078 (Hufsmith), ¶¶34 (considering execution environment and other factors), 111, 117 (contextual metric used to measure exploitability of security problem based on asset’s network accessibility), 121, 140

(application usage); EX1085 (Hovor), ¶¶48 (considering “location of the asset in the network topology”), 75-77 (considering scores of connected assets).

83. That contextually prioritized, actionable alerts were known for agentless workload scanning is not surprising since the underlying data providing such context was known to be available to agent-based and agentless workload scanners alike. *E.g.*, EX1007 (Veselov), 3:32-35 (snapshot may include memory, device state, and “block-level image of the entire logical volume”), 4:1-4 (“In any case, the snapshot is 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.”), 9:42-67 (“The ‘state’ of the target resource represented by the snapshot data may include any properties of the target resource instantiated in the computing resource service provider 104 that are relevant to the security assessment.”), 15:50-59 (snapshot may include data describing VM’s “activity and connections”), 16:2-6 (snapshot may include traffic profile). And while Veselov does not provide many details of the reporting process itself, a POSA would have recognized that Veselov’s high-level description of reporting is consistent with these well-known, commonplace types of alerts. *See, e.g.*, EX1007, 2:35-41 (“The present disclosure contemplates implementation of any suitable security assessment....”), 3:20-4:4 (snapshot-based assessments produce same results as assessments performed on targets

themselves), 10:17-36 (results sent “in a user-readable format to a user interface to be displayed on the user device” and may identify remediation), 20:57-21:13 (result describes vulnerability).

Sensitive-Data Detection

84. Dr. Kaeli argues that certain exhibits are tied to sensitive-data detection, but as I previously explained, this was known prior to 2019. It was a feature of many commercially available cloud-security products. For example, Symantec product documentation described a solution that would “[d]iscover, monitor, and protect your sensitive corporate information.” EX1134, 1-7.¹⁰ It “can accurately detect all of the confidential data in your organization—whether it’s at rest, in motion, or in use.” *Id.* (discovering stored data “across all of your cloud, mobile, network, endpoint, and storage systems”); *see also id.*, 4 (scanning files/file systems, databases, etc.). It involves various detection techniques for structured and unstructured data, and it also “detects content by looking for matches on specific keywords, regular expressions or patterns, and file properties.”

¹⁰ EX1134 is a correct copy of documentation obtained from <https://docs.broadcom.com/docs/data-loss-prevention-solution-en>. It has a copyright date of 2015 and, based on my experience, its description of sensitive-data detection is consistent with the state of the art in that timeframe.

Id., 1-2 (“credit card number” example). It would “[n]otify users with an on-screen, pop-up window.” *Id.*, 3.

85. This example is not an outlier. Sensitive-data detection was known in a variety of platforms such as Forcepoint and McAfee Skyhigh Security Cloud. EX1135, 5 (“Forcepoint Data Discovery can be used to learn the location of sensitive data both on-premises and in supported cloud-based applications. It can be used to scan data on file servers, email servers, and databases.”); EX1136 (McAfee), 1 (“Gain complete visibility into data, context, and user behavior across all cloud services, users, and devices.”), 2 (“Applies unified policies to all cloud services across data at rest and in transit.”), 3 (“Leverages keywords, pre-defined alphanumeric patterns, regular expressions, file metadata, document fingerprints, and database fingerprints to identify sensitive data.”).¹¹ Lateral movement risk—

¹¹ EX1135 and EX1136 are correct copies of product documentation obtained from https://www.websense.com/content/support/library/data/v84/deploy/deploy_dlp.pdf and <https://media.erepublic.com/document/DS-Skyhigh-Cloud-Security-Platform.pdf>, respectively. They are dated 2017 and 2018, respectively, and based on my experience, their descriptions of sensitive-data detection are consistent with sensitive-data-detection capabilities in that timeframe.

which, as discussed in the Technical Brief Mr. Shua cites, involves user credentials that may allow an attacker to move to other assets (EX2067, 13-14)—was also well known and routinely evaluated. *See, e.g.*, EX1044 (Giakouminakis), 6:35-58 (considering impact of connections to other assets), 8:46-9:29 (similar), EX1045 (Canzanese), ¶¶34-35 (discussing prioritized alerts and high-priority situations involving interconnected assets where confidential-data loss allows attackers to “move laterally to other user endpoints in the enterprise network”); EX1141 (2018 Quarterly Threat Report), 4 (noting lateral-movement protection).¹² The evidence indicates nothing special about doing this in the context of snapshot-based assessments, which is not surprising since the relevant data for evaluating this issue—e.g., stored credentials and connected assets—was readily available in snapshot-based assessments. *E.g.*, EX1007, 9:55-10:16 (snapshot includes files, “configurations of network protocols and services; properties of connections to

¹² EX1141 is a correct copy of Rapid7’s Quarterly Threat Report obtained from https://www.rapid7.com/globalassets/_pdfs/research/rapid7-threat-report-2018q3.pdf. It is consistent with my understanding, based on my experience in the field, that lateral movement risk was a well-known and commonly assessed security issue well before 2019.

other computing resources; permission settings;” software/configurations, logs, etc.).

86. Sensitive-data detection was also known in the context of agentless workload scanning, including snapshot scanning. For example, Chari describes identifying sensitive data in various resources in order “label” it, assign it a sensitivity level, and sanitize it. EX1056, ¶¶38-42. Chari specifically describes using this technique on snapshots, explaining that “the virtual machine image manager labels and sanitizes the snapshot of the specific instance of the virtual machine image at 916 and stores the labeled and sanitized snapshot of the specific instance of the virtual.” *Id.*, ¶94.

87. In sum, the features described in the secondary-considerations evidence cited by Dr. Kaeli were known, both individually and as a combination. In other words, the techniques that I describe above are not merely isolated features known in the limited context of agent-based scanning. Rather, as Veselov and the other agentless examples I discuss above make clear, these were known aspects of agentless workload scanning. *See, e.g.*, EX1007, 2:35-41, 4:1-4; EX1040 (Ammons), ¶¶28 (API), 35 (cloud), 45 (sensitive components); EX1078 (Hufsmith), ¶¶42 (API), 44-46 (cloud), 67 (cloud), 140 (cloud), 148 (API).

2. Long-Felt Need

88. For the reasons I discuss below, the evidence does not show that Orca's product or the claims solved a long-felt need.

89. As an initial matter, Dr. Kaeli's and Mr. Shua's description of the supposed problems with prior cloud-security approaches does not consider the full scope of known techniques. Dr. Kaeli states that prior cloud-security approaches had shortcomings related to limited visibility and limited coverage. Kaeli-Dec, ¶¶168-72; *see also* Shua-Dec, ¶¶85-93 (citing EX1001, 1:60-2:17; EX2082, 0:10-3:00; EX2083, 1; EX2099; EX2100, 2; EX2102, 1; EX2103, 1; EX2104, 2; EX2133, 3-4). But he bases this assertion on the premise that the only known approaches were traffic inspection/monitoring, network scanning, and agents installed on the target. Kaeli-Dec, ¶¶169-70; Shua-Dec, ¶86. In other words, he argues, traffic inspection/monitoring and network scanning did not have sufficient visibility, whereas agents, though having good visibility, had coverage issues because they were cumbersome to install. Based on that narrow characterization of what was known, Dr. Kaeli argues that agentless workload scanning solved this supposed need by providing agent-like visibility without the coverage limitations of agents. Kaeli-Dec, ¶¶173-74 (citing EX2076, 7, 11, 31; EX2138, 1; EX2135, 1, 3); *see also* Shua-Dec, ¶¶8-10, 86-88. As I explained above, however, this argument overlooks that agentless workload scanning (including such scanning for

powered-off assets) was already known. *Supra*, §IX.B.1; EX1002, ¶¶56-58, 62.

Indeed, Veselov itself explains that its snapshot-based scanning—despite the analysis being performed separately from the target asset without using agents installed in on the target (EX1002, ¶¶149-53), a process that does not depend on whether the asset is actively running, suspended, or powered off (*id.*, ¶33)—provides agent-like visibility. EX1007, 3:20-4:4. Dr. Kaeli and Mr. Shua appear to be focusing exclusively on what had been widely commercialized rather than what was known in the field. I have been advised that, for purposes of considering nexus, the question of whether features were known is not limited to whether features were known to be part of commercialized products.

90. Dr. Kaeli and Mr. Shua mention other aspects like API usage and targeted vulnerability alerts, but these were also well known, as I explained above. *Supra*, §IX.B.1; Kaeli-Dec, ¶¶173-75 (citing EX2135, 3; EX2076, 31 [REDACTED]); Shua-Dec, ¶¶93-94 (citing EX2076, 7; EX2103, 1). They were also known as in combination with agentless workload scanning, which is not surprising since cloud APIs are the fundamental communication mechanism of cloud products (including cloud-security platforms), and targeted/prioritized vulnerability alerts were commonplace. *Supra*, §IX.B.1; EX1002, ¶¶42-47 (discussing well-known cloud computing and cloud APIs), 64-70 (discussing well-known alerting techniques).

91. Orca’s own evidence actually contradicts Dr. Kaeli’s assertion that the field had long felt the alleged need and that Orca provided the alleged comprehensive solution. Mr. Shua states that “many people in the industry disagreed that network scanners and agent-based solutions were inadequate for cloud security.” Shua-Dec, ¶90. Mr. Shua also cites a discussion of Qualys’s product as describing limitations of agents (*id.*, ¶91), but that same document states that Qualys’s platform “resolves this conundrum.” EX2102, 1-2. As to the 2024 article that Dr. Kaeli and Mr. Shua cite repeatedly, it ultimately emphasizes the limitations of agentless workload scanning and the industry’s movement towards agents. For example, it observes that “[c]ontainers are the infrastructure of the cloud, powering the vast majority of cloud workloads” (EX2076, 10) but says that “[a]gentless scanning doesn't fit the containerized world well.” *Id.*, 12-13. It also states that, among other limitations, agentless solutions are “severely limited in what they can show,” yield higher false-positive rates, and do not support “a surprising number of workload types.” *Id.*, 10, 12-13. Ultimately, the author states that the market is “evolving to see the pitfalls of a pure agentless approach.” *Id.*, 23. That is why, in a portion of EX2076 that is obscured in the copy Orca

submitted but is clear in the unobscured version (EX1125),¹³ the author states that “[c]ompanies that once championed agentless solutions are now all building agents.” EX1125, 3; EX2076, 2.

92. I therefore disagree that the claims solved a long-felt need.

3. Skepticism

93. I disagree with Mr. Shua and Dr. Kaeli that the evidence shows that Orca overcame skepticism tied to the claims’ merits.

94. First, Dr. Kaeli and Mr. Shua focus primarily on the supposed skepticism of agentless workload scanning. Kaeli-Dec, ¶¶177-181 (citing EX2076, 7-8; EX2103; EX2106, 1; EX2135, 2; [REDACTED]); Shua-Dec, ¶¶95-97. As I previously discussed, however, agentless workload scanning (including cloud-API-based agentless workload scanning) was known. *Supra*, §IX.B.1.

95. They also cite evidence that does not involve overcoming skepticism in the first place. Dr. Kaeli points to a competitor’s article saying that agentless solutions fall short, but based on my experience in the field, this type of criticism of competing products is consistent with standard marketing tactics rather than

¹³ EX1125 is a correct copy of a webpage obtained from <https://softwareanalyst.substack.com/p/redefining-cnapp-a-complete-guide>.

being truly indicative of industry skepticism. Kaeli-Dec, ¶178; EX2106, 1. As to the 2024 retrospective, the author’s comment involving some who were “skeptics” says that these people were “right to be skeptics” (EX2076, 7), and as I previously discussed, the author discusses limitations of agentless approaches and the industry’s movement away from pure agentless approaches. *Id.*, 10, 12-13, 23; EX1125 (unobscured), 3. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

96. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

97. I therefore disagree that Orca overcame skepticism tied to the claims solved a long-felt need.

4. Praise

98. I disagree with Mr. Shua's and Dr. Kaeli's arguments that awards, Wiz documents, and other documents demonstrate industry praise tied to the claims. Kaeli-Dec, ¶¶182-86; Shua-Dec, ¶¶133-50.

99. As an initial matter, I note that much of the cited evidence appears to be Orca's own marketing material and other statements by Orca's partners/investors. I have been advised that self-praise and other forms of self-interested praise are generally not considered to be relevant to the praise inquiry for purposes of secondary considerations. That makes sense because, for example, an award from an investor is more likely to be driven by the investor's self-interest than an award from a neutral third party. As to the specific exhibits, many appear to be Orca's own press releases. EX2121-2122; EX2124-2125; EX2128-2130; EX2132. Other cited evidence quote Orca's statements or statements from Orca's investors/partners. For example, EX2117 acknowledges that Panorama Technologies is "an Orca Security partner," meaning their description of Orca is essentially self-marketing. That understanding is consistent with the rest of the document, which, based on my experience in the field, appears to be a form of advertising rather than an objective assessment of Orca's technology. *E.g.*, EX2117, 3 ("By choosing Orca with the support of Panorama Technologies, you choose a safe and secure digital future. Together, we will meet the challenges of

cybersecurity, ensuring the integrity and protection of what you value most every step of the way.”).

100. Mr. Shua touts other alleged praise that was made by Orca’s own partners and investors. Shua-Dec, ¶¶135-36,141, 146-48; EX2118 (Tag Security Report); EX2129-2130 (discussing Redpoint); EX2131 (Notable Capital); EX2145 (award sponsored by YL Ventures). I have been advised that Mr. Shua acknowledged relationships with Redpoint, Notable Capital, YL Ventures, and TAG Security. Orca’s relationship with TAG Security is further suggested by EX1115, which I understand is a copy of a webpage from Orca’s website that was introduced during Mr. Shua’s deposition.¹⁴ It indicates that the TAG Security Report submitted as EX2118 is part of Orca’s marketing material. The other relationships are also indicated by the submitted documents. For example, Mr. Shua points to Orca being named to Redpoint’s “InfraRed 100 List” (EX2129-2130), but another exhibit shows that Redpoint had invested in Orca. EX2113, 1 (“Orca Security’s C round also includes substantial participation from venture capital firms CapitalG (Alphabet’s independent growth fund), Redpoint Ventures,

¹⁴ I understand that EX1115 was obtained from <https://orca.security/lp/tag-cyber-cloud-security-vendor-landscape-tco-analysis/>. I have compared it with the webpage at that URL and confirm that it is an accurate representation.

GGV, ICONIQ Capital, Lone Pine Capital, Stripes, Adams Street Partners, Willoughby Capital, and Harmony Partners.”). Mr. Shua points to Orca being named to a similar list by Notable Capital, but the document shows that Notable Capital was also an Orca investor. EX2131, 2 (Orca “in Notable Capital’s portfolio”). Likewise, Mr. Shua points to an award “[s]ponsored by YL Ventures” (EX2145), but other documents shows that YL Ventures was an Orca investor. EX2113, 3 (“YL Ventures has supported Orca Security since its seed round.”).

101. However, even if I did not consider the connection to Orca’s investors, these documents still would not favor nonobviousness in my opinion because they are not tied to recited and novel aspects of the claims. As I discussed above, in my opinion, Dr. Kaeli and Mr. Shua have not shown that the product embodies and is coextensive with the claims, so general company/product praise would not have a presumption of nexus. *Supra*, §IX.A. Regardless, to the extent particular aspects of Orca and its product are discussed in the cited exhibits at all, they were known, as I discussed above. *Supra*, §IX.B.1. I see no other evidence showing that the above awards/lists are due to the claims. Dr. Kaeli does not address the Redpoint, Notable Capital, or YL Ventures lists, and Mr. Shua does not provide specifics. Shua-Dec, ¶¶141, 146-50.

102. I also disagree that other awards listed by Dr. Kaeli and Mr. Shua show praise tied to recited and novel features of the claims. Kaeli-Dec, ¶182

(citing EX2119, 9-10; EX2120, 6-7; EX2123); Shua-Dec, ¶¶138-46 (citing EX2119-2124; EX2127-2128; 2143-2144). Dr. Kaeli says that he reviewed EX2119, EX2120, and EX2123, and he quotes a few passages, but he does not explain why he thinks this is praise of the claims. Regardless, the passages he quotes merely describe the unrecited and/or known features I previously discussed. Automatically assessing “every discovered asset” is not recited. EX2119, 9-10; *supra*, §IX.A. Deep visibility and accessing cloud configurations to assess workloads’ block-storage out-of-band was known. EX2120, 6-7; *supra*, §IX.B.1. The “agentless approach” was known. EX2123; *supra*, §IX.B.1. Mr. Shua cites additional exhibits but similarly gives no indication that they relate to recited and novel features of the claims. For example, his quote of the Forbes 2021 exhibit addresses the agentless aspect of Orca’s approach, which was known, and he gives no meaningful specifics for the Forbes 2023 exhibits. Shua-Dec, ¶¶139, 146; EX2120; EX2127-2128. He gives no meaningful specifics for the SINET16 or Gartner “Cool Vendor” listings. Shua-Dec, ¶¶140, 142; EX2143-2144; EX2121. For the SC Media CSPM award, he quotes a customer’s statement about Orca’s “agentless approach,” but again agentless workload scanning was known. Shua-Dec, ¶145; EX2123; EX2125; *supra*, §IX.B.1. EX2120 repeats Orca’s marketing describing the platform as detecting issues including lateral movement risk. *Compare* EX2120 (Forbes article), 7 *with* EX2121, 2. But LMR detection is not

praised, the claims do not recite LMR, assessing LMR was known, and the evidence does not indicate any special about detecting LMR in combination with other claim elements. *Supra*, §IX.B.1.

103.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

104.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

105. Lastly, with respect to the additional cited evidence cited by Dr. Kaeli (Kaeli-Dec, ¶185), in my opinion it does not show praise tied to claimed and novel features. Dr. Kaeli quotes the author's discussion of agentless workload scanning enabled by cloud APIs and his description of Orca's platform as gathering block-level data via the cloud provider's APIs and performing a read-only scan (EX2067, 7, 11-12, 31), but this approach was known, as Veselov itself shows. *Supra*, §IX.B.1. Dr. Kaeli also quotes the report that I understand was commissioned by Orca (as I discussed above). EX2118, 4-5. Those are not objective assessments for the reasons I discussed above, and even if they were, the quoted statements

address known features of agentless scanning such as out-of-band scanning of block-level data, and reporting actionable results. *Supra*, §IX.B.1.

5. Copying

106. I understand that Orca’s copying arguments are primarily based on argument that Wiz had access to Orca’s product—with the earlier alleged access in August 2020—and that Wiz’s product then became substantially similar to Orca’s product. I have been advised that access and substantial similarity is considered evidence of copying. In my opinion, however, the evidence does not support Orca’s “access and substantial similarity” story, nor does the evidence otherwise show copying with nexus to the claims.

107. Below, I first discuss my assessment of several Wiz documents. These documents go the heart of Orca’s copying story since they show that [REDACTED]

[REDACTED]

[REDACTED] I then turn to Orca’s evidence and explain why, in my opinion, the evidence does not show copying, and regardless, the allegedly copied features were known.

a. Development of Wiz’s Platform

108. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

109. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

15 [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

110. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

111.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

112.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

113. Moreover, the evidence indicates that agentless scanning itself was one aspect of Wiz's platform and that other features were important drivers of its success. For example, a press release submitted by Orca notes that Wiz was the fastest software company to achieve \$100M in annual recurring revenue and

emphasizes the “Wiz Security Graph” and the way it presents a variety of risks on the graph. EX2116, 3. I understand that Wiz has obtained dozens of patents, including many involving its graph-based approach. I have also reviewed an announcement of Wiz’s agreement to be acquired by Google for \$32 billion.

EX1137.¹⁶ In announcing the agreement, Google praised Wiz for delivering “an easy-to-use security platform,” being an “innovative leader,” and “continu[ing] to deliver new products with strong adoption, fueling rapid business growth, including over the last 12 months in which it has begun to deliver new categories of cybersecurity solutions.” *Id.* Thus, the evidence indicates that Wiz did not copy its agentless scanning approach from Orca (as I discuss in more detail below) and that, regardless, other features were key drivers of its success.

b. Orca’s Evidence

114. As an initial matter, Dr. Kaeli’s copying story appears to conflict with Mr. Shua’s version of events. [REDACTED]

[REDACTED]

[REDACTED]

¹⁶ EX1137 is a correct copy of the webpage obtained from <https://www.prnewswire.com/news-releases/google-announces-agreement-to-acquire-wiz-302404500.html>.

[REDACTED] However, Mr. Shua indicates that he had already heard about Wiz's agentless workload scanner in "early 2020." Shua-Dec, ¶119. [REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]

115.

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

116.

[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

117. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[illegible]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

118. I also disagree with Dr. Kaeli's argument that Orca and Wiz's products are substantially similar. Kaeli-Dec, ¶¶214-238. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] He points to high-level summaries of agentless scanning that refer to accessing the customer's environment via cloud APIs and the use of snapshots, but this has the same issue. Kaeli-Dec, ¶216 (citing EX2026; EX2115). He argues that Wiz's product is substantially similar to Orca's with respect to the independent claim elements. But he does not meaningfully address element 1.5 ("correlating the determined potential cyber vulnerabilities with a network location of the protected virtual cloud asset"). Kaeli-Dec, ¶¶232-35. He quotes a statement on Wiz's website mentioning lateral movement and a statement on another website mentioning Wiz's "Security Graph," but he gives no explanation of how this supposedly shows that companies' platforms are meaningfully similar. Kaeli-Dec, ¶235 (citing

EX2177; EX2178). For elements 1.6, 1.7, and 1.8, Dr. simply states that Wiz's platform performs risk-based prioritization of vulnerabilities to mitigate alert fatigue and reports issues as alerts. Kaeli-Dec, ¶¶236-38 (citing EX2165, 8, 14-15). These short, extremely high-level assertions do not explain how Wiz's platform is supposedly similar to Orca's.

119. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] Orca's own marketing material distinguishes the products. EX1126.¹⁷ [REDACTED]

[REDACTED]

¹⁷ EX1126 is a correct copy of a webpage obtained from <https://orca.security/resources/comparisons/wiz/>.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

120. Notably, Dr. Kaeli repeatedly cites the 2024 article for the author’s positive comments about Orca, but he does not acknowledge the author’s discussion of significant differences between Orca’s and Wiz’s products. *E.g.*, EX2076, 7 (“Wiz pushed agentless scanning to new limits, most importantly by introducing the intuitive graphing of asset relationships to security teams, and building their entire UX around that capability.”), 8 (Wiz “introduced an agent

when people started asking for it” and “improved their DSPM capabilities”), 11 (emphasizing Wiz’s graph), 13 (investment in agents), 17 (emphasizing graph and Wiz code), 21-22 (“Wiz differentiated itself by taking a graph-based approach to visibility, detection, high-fidelity risk prioritization alerts and an obsessive customer-driven approach, setting it apart from the competition.”), 22 (“A core factor in Wiz's success is its graph-based architecture, from the actual database to the user experience.”), 23 (“It is widely known and should be emphasized Wiz launched a much more effective agentless scanner” leveraging graph), 23 (attack-path analysis), 24 (data-clustering techniques informed by attacker research), 25 (Wiz code).

121. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

6. Commercial Success

122. As I discuss below, I disagree with Dr. Kaeli's and Mr. Shua's argument that Orca's product and Wiz's product achieved commercial success due to the claims.

123. In my opinion, the evidence does not show that Orca's startup financing and revenue were obtained due to the claims' merits (i.e., recited and novel features). Kaeli-Dec, ¶239. That financing/revenue is related to the product generally does not itself indicate nexus to the claims since, as I previously explained, the requirements for the presumption of nexus have been shown. *Supra*, §IX.B.1. In any event, the few aspects of the product Mr. Shua discusses are the same ones that I previously explained were known and/or unrecited. *Supra*, §§IX.A, IX.B.1; Shua-Dec, ¶¶102-13.

124. I also disagree with Mr. Shua's argument that the financing/revenue is tied to the claims. Aside from Orca's own press releases (EX2065; EX2108;

EX2110; EX2113; EX2140), Mr. Shua cites articles that either just describe Orca or that just repeat statements from Mr. Shua and his investors. Shua-Dec, ¶¶102-13. He quotes EX2109 and EX2111 as supposedly showing a third-party's recognition of a connection to the claims, but these merely say that Orca's API-based agentless workload scanning makes it stand out. Shua-Dec, ¶¶104, 108; EX2109, 1; EX2111, 1. These statements does not indicate that the financing was obtained due to the claims, and regardless, API-based agentless workload scanning was known. *Supra*, §IX.B.1. Mr. Shua points to a quote from a "GGV managing partner" (EX2109, 3), but that individual does not appear to be a disinterested third-party since, as the document states, GGV is an Orca investor. *Id.*, 1 (financing round "led by GGV Capital"). The rest of EX2111 largely just quotes or paraphrases Mr. Shua. EX2111, 2-4. Mr. Shua cites EX2112, which does not connect the funding to the claims, and he quotes the general partner of CapitalG, who also does not appear to be a disinterested third party because, as the document states, CapitalG led this funding round. Shua-Dec, ¶¶109-10; EX2112. Mr. Shua cites a quote in Orca's press release and states that this quote is from an "industry expert," but again the quote does not appear to be from a disinterested third party since it is from a partner of Adams Street, which the document says provided "substantial participation" in the funding round. Shua-Dec, ¶113; EX2113, 1.

125. I also disagree with Dr. Kaeli's argument that Wiz's commercial success is due to the claims. He does not explain this opinion beyond citing a few press releases without analysis. Kaeli-Dec, ¶240. Moreover, Orca's evidence explicitly attributes Wiz's success to other factors that are not recited. For example, the 2024 retrospective places particular emphasis on Wiz's graph, explaining that a "core factor in Wiz's success is its graph-based architecture, from the actual database to the user experience." EX2076, 21-23; *see also id.*, 23 (attack-path analysis), 24 (data-clustering techniques informed by attacker research), 25 (Wiz code). The article also emphasizes Wiz's responsiveness to customer needs:

Wiz's true strength is in their ability to come to market at exactly the right time to meet customers where they're at. Rather than sit on their agentless success, they introduced an agent when people started asking for it. They improved their DSPM capabilities as teams cared about it. Most recently, they announced Wiz Code just as security teams are catching onto the power of ASPM capabilities. Wiz is frightening to the market precisely because they're not just a history lesson stuck in what they invented, but continue to release just the right features at just the right time to appease the majority of the market.

EX2076, 7-8. Thus, in my opinion, the evidence does not show that Wiz's success has a nexus to the claims.

X. Concluding Statements

126. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required of me, I will appear for cross-examination remotely within the United States during the time allotted for such cross-examination at a time and place agreed by counsel.

127. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Dated: August 15, 2025

By: /Angelos Stavrou/

Angelos Stavrou, Ph.D.