

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

WIZ, INC.,
Petitioner,

v.

ORCA SECURITY LTD.,
Patent Owner.

IPR2024-01190

PATENT NO. 11,740,926

DECLARATION OF ANGELOS STAVROU

Table of Contents

I. Background	2
II. Materials Considered	3
III. Level of Ordinary Skill in the Art.....	6
IV. Legal Principles.....	6
V. Claim Construction	6
A. “location”	6
B. Analyze/Analyzing “the at least one snapshot”	7
VI. Ground 1.....	7
A. Dr. Kaeli’s Embodiment-Mixing Argument.....	7
B. The Veselov-Mohanty Combination.....	12
C. The Combination Satisfies the Claims.....	27
1. Elements 1.4/14.4/15.4.....	27
2. Elements 1.5-1.6/14.5-14.6/15.5-15.6.....	38
3. Elements 1.7/14.7/15.7.....	38
4. Claim 13	40
VII. Ground 2.....	49
VIII. Grounds 3-4.....	54
IX. Grounds 5-6.....	56
X. Secondary Considerations.....	59
A. No Presumption of Nexus	61
B. Orca’s Evidence	64
1. Known Techniques/Features	65
2. Long-Felt Need	76
3. Skepticism	80
4. Praise	82
5. Copying	88
6. Commercial Success	101
XI. Concluding Statements	104

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 1-15 of U.S. Patent No. 11,740,926 (“’926 Patent” or EX1001) in IPR2024-01190. 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, ¶¶29-30, *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*, §X.

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, ¶¶43-48.

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, ¶¶44-48. 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. Moreover, as I explain in more detail below (*infra*, §VI.A), Veselov makes perfectly clear that its description of the system components, the preceding steps (including the various ways Veselov creates snapshots), and the subsequent steps (including the processing/reporting of detected vulnerabilities as alerts) apply regardless of which option for the analysis step is used.

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, ¶78. 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, Veselov is describing features common to all embodiments and then simply not repeating them 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—was 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, ¶¶79-85, 162. 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, ¶80), 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, ¶80.

13. Dr. Kaeli also incorrectly portrays direct snapshot analysis as 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, ¶¶78-79. 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, ¶162), 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 exists.” Kaeli-Dec, ¶162. 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

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.

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 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*, §X.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, ¶87. 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, ¶¶87, 93, 107 (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, ¶¶112-16. 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, ¶¶112-16. 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, ¶88; *see also id.*, ¶92. 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., threat determination, assessing risk associated with those threats, and prioritizing/reporting based on the risk (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,

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/>.

¶111. A POSA therefore would have recognized that the teachings of Mohanty I discussed 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, ¶88. 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.*, ¶¶89-92. 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.*, ¶93. However, at no point did I discuss installing agents on the target asset as part of the combination. EX1002, ¶¶105-17. Rather, I explained the POSA's motivations and reasonable expectation of success in applying Mohanty's teachings— e.g., threat determination, assessing risk associated with those threats, and prioritizing/reporting based on the risk—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, ¶111. Instead, Dr. Kaeli focuses on the *process* for collecting data (or the component that does the collecting), rather than the data itself. Kaeli-Dec, ¶94. 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, ¶¶94-95; *see also id.*, ¶111. 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, 110-11; 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, ¶96 (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 threat/risk determination and prioritization/prioritized alerts to snapshot-based assessments, nor does it address the benefits of the combination that I explained. EX1002, ¶¶112-16. 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, ¶97. 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.*, ¶¶97-98. 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.*, ¶98. 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, ¶98 n.3. Even if this were true, it would not impact the combination I discussed or the claim-

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 the data stored on virtual disks such as files, applications, configurations, etc. 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”

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

compared to Figs. 3A/B. Kaeli-Dec, ¶98. Veselov discusses the string/pattern comparison as an example of how data in the snapshot may be initially identified when performing direct snapshot analysis, not the subsequent assessment tasks that detect potential security problems. 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. EX1007, 18:32-53 (identifying “files, software packages, application versions, and the like”). Beyond the software/versions—which encompass the startup services I previously discussed—numerous types of information are stored as “files,” which were well known to include the types of data I discussed as part of the combination (e.g., unencrypted sensitive data,

personally identifiable information, passwords, etc.). *See id.*, 3:20-40, 7:16-22, 9:42-67; EX1002, ¶¶161-67.

24. I disagree with Dr. Kaeli's arguments addressing the background exhibits that I cited in my original declaration. Kaeli-Dec, ¶¶99-105. 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, ¶¶109-17. 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, ¶99 (citing EX1014, EX1034, EX1045, EX1046, EX1075, EX1077, EX1063; EX1008); *see also* Kaeli-Dec, ¶¶102-04 (dismissing EX1077, EX1094, and EX1095 because they do not mention snapshots). He similarly dismisses EX1025 (NIST Incident Handling). Kaeli-Dec, ¶99. 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, ¶100. 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

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/AWSE>

of which is a boot volume) along with metadata such as launch parameters.

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 dismissal of EX1056 (Chari). I cited this exhibit in the Ground 1 rationale section to illustrate the well-known and predictable nature of detecting sensitive data in both VMs and in stored disk data, and associated prioritization. EX1002, ¶¶109, 113, 117; EX1056, ¶¶38-39, 64-66.

C2/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.

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>.

Dr. Kaeli dismisses Chari because he claims it “relies on interactions with the target VM to label and sanitize sensitive data but does not ever detect sensitive data or cyber threats.” Kaeli-Dec, ¶101. This is incorrect. For example, Chari describes various types of sensitive data that present security threats (EX1056, ¶38), and it describes labeling and sanitizing such sensitive data in a snapshot (*id.*, ¶94). It would have been perfectly clear to a POSA that Chari’s description of labeling/sanitizing data in a snapshot indicates finding the data in the snapshot in order to add the label to that data and then sanitize it.

27. I also disagree with Dr. Kaeli’s argument that Mohanty’s techniques “work against Veselov’s efforts to reduce the complexity of security assessments.” Kaeli-Dec, ¶¶106-07 (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, threat determination and context-sensitive risk assessment/prioritization were commonplace, and a POSA would not have viewed the applied teachings as overly complex. In fact, Veselov encourages performing multiple assessment tasks. *Id.*, 2:35-41.

28. I disagree with Dr. Kaeli’s response to my explanation of motivation “to provide more accurate assessment of risks posed by particular threats.” EX1002, ¶113. Dr. Kaeli repeats his earlier argument that Mohanty is not

compatible with snapshots because they are point-in-time captures that become obsolete. Kaeli-Dec, ¶107. 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 determining threats based on stored data and measuring context-sensitive risk would be similarly useful to assessments of live VMs and snapshot-based assessments. Threats and 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, ¶107. I explained why these arguments are incorrect above.

29. I also disagree with Dr. Kaeli's response to my explanation of motivation "to obtain known benefits of prioritizing cyber threats and associated alerts." EX1002, ¶¶114-15; Kaeli-Dec, ¶¶108-09. 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, ¶108. 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, ¶¶108-17. 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, ¶109. 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 determine applicable cyber threats, 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, ¶111. 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 determining applicable cyber threats and measuring context-sensitive risk with Veselov's reasonably current and continuous snapshot-based approach.

30. 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, ¶¶110-12. 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.

31. 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

32. 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 1.4/14.4/15.4

33. I previously explained why the combination teaches or suggests all aspects of element 1.4. EX1002, ¶¶147-68. In response, Dr. Kaeli argues that I did not show this limitation was obvious. Kaeli-Dec, ¶¶115-40. I respectfully disagree.

34. As an initial matter, Dr. Kaeli portrays my analysis as presenting isolated arguments that Veselov teaches element 1.4 and that Mohanty teaches element 1.4. Kaeli-Dec, ¶¶118 ("Dr. Stavrou first argues that this limitation is met

by Veselov on its own.”), 130 (“Separately from Veselov, Dr. Stavrou also relies on disclosures of Mohanty.”). This framing is misleading since I explained how the *combination* satisfies element 1.4. EX1002, ¶148 (“As I explain below, the combination of Veselov and Mohanty teaches these elements.”).

35. In any event, I disagree with Dr. Kaeli’s arguments involving Veselov. Dr. Kaeli argues that Veselov does not broadly disclose detecting a wide variety of security problems. Kaeli-Dec, ¶120. In particular, he states that “Veselov’s disclosures only discuss ‘vulnerabilities’ and ‘host configuration assessments,’” which he claims do not encompasses cyber threats that based on the listed data types. *Id.*, ¶¶120-21. That is incorrect. Discussing known assessments, Veselov notes that service providers enable security platforms to identify “vulnerabilities, bad configurations, and the like.” EX1007, 1:41-50. Veselov then explains that its own system may perform “any suitable security assessment” and gives an open-ended list of examples including “Common Vulnerabilities and Exposures (CVEs), Center for Internet Security (CIS) benchmarks, ‘best practices’ packages, static or runtime behavior analysis, host configuration assessments, and the like.” *Id.*, 2:35-41. Veselov then confirms that its snapshot-based assessments identify the same types of security issues as assessments performed on target assets themselves. *Id.*, 3:41-4:4. I previously discussed the various types of known security issues on virtual assets detected by commonplace security-assessment

tasks like those referenced in Veselov, and these include cyber threats based on the recited types of data. EX1002, ¶¶49-53 (discussing various known risks), 54-55 (various types of issues detected, including sensitive data), 161-67 (discussing data types recited in element 1.4).

36. I also disagree with Dr. Kaeli’s argument that “cyber threats” as recited do not encompass any CVEs or other “vulnerabilities.” Though acknowledging that “vulnerabilities are a type of cyber threat,” he argues that the recited cyber threats—which are based on at least one of the recited data types—exclude vulnerabilities and CVEs. Kaeli-Dec, ¶¶122, 131. He bases this assertion on a paragraph in the specification describing certain “cyber threats that do not represent vulnerabilities. *Id.*, ¶123 (citing EX1001, 6:20-32). But nothing in that paragraph or the rest of the specification categorically excludes CVEs from the types of cyber threats recited in element 1.4, as Dr. Kaeli suggests. To the contrary, the specification uses “cyber threat” broadly and states that vulnerabilities in applications are types of threats. EX1001, 1:56-60 (discussing “security threat, such as vulnerability in the application”); *see also id.*, 4:45-47 (discussing vulnerabilities and “other” cyber threats), 6:40-43 (reporting “detected cyber threats (including vulnerabilities)”).

37. I also disagree with Dr. Kaeli’s response to my discussion of configuration assessments. He argues that element 1.4 does not recite “nonsecure

system configurations” as one of the data types. Kaeli-Dec, ¶126. But I explained why typical configuration assessments, as taught by the references, encompass determining cyber threats based on stored data (EX1002, ¶¶157-60) and then went on to explain specific recited data types that form the basis of the determined cyber threats (*id.*, ¶¶161-67). Dr. Kaeli argues that “[a]ssessments of system configurations do not include any of the specific types of data in the ’926 patent claims.” Kaeli-Dec, ¶127. This is incorrect since asset “configuration” covers a wide variety of data, including data types of element 1.4 that I discussed in my previous declaration. Dr. Kaeli points to additional types of configurations like “network port settings” that are mentioned in the evidence I cited to support my analysis. Kaeli-Dec, ¶127 (citing EX1008, ¶37). But whether *some* configuration data goes beyond what is recited does not address the fact that other types of configuration data are recited. EX1002, ¶157. Indeed, the very paragraph of EX1008 (Basavapatna) that Dr. Kaeli quotes demonstrates that broad nature of “configuration.” EX1008, ¶37 (asset configuration includes physical configuration, software, and software configurations). Turning to my explanation that “weak passwords and weak encryption schemes” are types of unsecure asset configurations routinely evaluated during configuration assessments (EX1002, ¶157), Dr. Kaeli claims the cited evidence shows otherwise. Kaeli-Dec, ¶127 (citing EX1008; EX1020). But EX1008 identifies weak passwords as one type of

security issue detected based on the previously discussed configuration. EX1008, ¶21 (discussing response to determination that “a password is too short”).

38. Dr. Kaeli claims I contended that Veselov can be used with “any security assessment” based on its broad teaching that its process can be used with any suitable assessment. Kaeli-Dec, ¶128. Veselov’s encouragement to use known assessment tasks is indeed broad, but for purposes of my declaration, I explained particular types of known assessment tasks that are “suitable” for Veselov’s snapshot-based approach, per Veselov’s teaching. EX1002, ¶¶109-11, 117. As to what supposedly makes the applied teachings unsuitable, Dr. Kaeli repeats his arguments that Veselov duplicate-VM approach and direct-snapshot-analysis approach are “entirely distinct embodiments” and that agent-based techniques are not suitable for snapshot-based assessments. I explained why each of these arguments is flawed above. *Supra*, §VI.A (discussing Veselov’s approaches), VI.B (Mohanty’s techniques consider same types of data as Veselov, and combination does not involve the incorporation of agents).

39. I disagree with Dr. Kaeli’s arguments about “external” threats. He dismisses Mohanty because it supposedly only considers external threats, not threats based on data stored on a virtual disk. Kaeli-Dec, ¶132. Dr. Kaeli is incorrect. While Mohanty describes examples in which threat information is

provided by intelligence systems,⁸ Mohanty expressly teaches determining whether threats are applicable to the asset by correlating threat information with the asset's stored data (including recited types of data). EX1002, ¶163; EX1075, 5:16-58 (discussing sensitive data), 7:30-58 (determining whether threat is applicable to asset by correlating threat information with "characteristics of the asset"), 10:27-34 (example remediation includes applying stronger encryption to sensitive data); *see also* EX1075, 7:59-8:18 (discussing example scenarios). In other words, Mohanty determines whether the type of data associated with the threat is stored on the asset. Thus, consistent with the claims, Mohanty teaches determining the existence of "potential cyber threats," each of which is "based on data stored on the virtual disk."

40. I disagree with Dr. Kaeli's further arguments involving the recited data types. He first argues that Veselov's description of the snapshot does not mention the recited data types (Kaeli-Dec, ¶129), but this assertion does not address my explanation of how a POSA would have understood the combined teachings, including Veselov's broad description of snapshot content and

⁸ The '926 specification similarly describes using such systems. EX1001, 4:34-44 (interfacing with "external intelligence systems" such as security systems "providing feeds on discovered threats").

assessment tasks (including CVE assessments) and Mohanty's further teachings of detecting particular types of cyber threats based on recited data types. EX1002, ¶¶161-67. Dr. Kaeli then dismisses Mohanty because it does not discuss analyzing snapshots and supposedly requires agents (Kaeli-Dec, ¶¶130-31), but I previously explained why Dr. Kaeli's arguments against the combination are incorrect.

Supra, §VI.B.

41. I disagree with Dr. Kaeli's arguments involving "unencrypted sensitive data." He claims Mohanty does not suggest determining sensitive data because it identifies that data "upon **creation** of the asset." Kaeli-Dec, ¶133. As an initial matter, this does not address the combination, which applies Mohanty's teachings in the context of Veselov's snapshot analysis (not wholesale incorporation of Mohanty's system). In other words, even if Mohanty were limited to initially detecting the presence of sensitive data when the asset is created and deployed, the combination applies Mohanty's teachings/suggestions to analysis of snapshots containing the same types of underlying data (including unencrypted sensitive data). And as I previously explained, determining threats that are based on stored sensitive data was commonplace and readily implemented in snapshot/image analysis. EX1002, ¶¶110-11, 164; EX1056 (Chari), ¶¶38-39, 60-61, 94; EX1078 (Hufsmith), ¶¶111, 181.

42. Regardless, in my opinion Dr. Kaeli is not reading Mohanty from the POSA's perspective to ascertain what it fairly suggests. In the sentence immediately after the one quoted by Dr. Kaeli, Mohanty explains that "[f]urther information can be written to tags"—not just dynamic tags—"as situations occur." EX1078, 4:4-9. I also discussed dynamic detection of suspicious data transfers. EX1002, ¶163. Dr. Kaeli argues that Mohanty does not describe using dynamic tags for sensitive data, that transfers are not "stored" data, and that suspicious data transfers "could" involve non-sensitive data (Kaeli-Dec, ¶¶138-39), but this does not consider what Mohanty suggests from the perspective of a POSA, who would have known that sensitive data stored on the system was a common target of such transfers. Dr. Kaeli claims I admitted that a POSA would be incapable of determining whether data is sensitive without further context. *Id.*, ¶139 (citing EX2172, 200:18-201:6). While *some* types of data may or may not be sensitive depending on a person's subjective opinion, what I actually said was that I could not answer the attorney's question about whether transferred data was generally "sensitive data" without further context about the type of data he was asking about. EX2172, 200:18-201:15. I then explained why Mohanty's teachings as a whole suggest determining cyber threats that are based on stored sensitive data. *Id.*, 201:7-204:9. Dr. Kaeli argues that Mohanty's discussion of recommending stronger encryption for some sensitive data is "beside the point" because Mohanty

also teaches marking sensitive data with tags. Kaeli-Dec, ¶137. But a POSA would have recognized that Mohanty's discussion of recommending stronger encryption in some cases suggests particularized detection/evaluation of the data in order to determine whether stronger encryption is warranted. EX1002, ¶163; EX1075, 10:27-34.

43. I disagree with Dr. Kaeli's arguments involving well-known best practices. Dr. Kaeli argues that my discussion of Veselov's disclosure of using best-practices packages (and my explanation that this typically included avoiding unencrypted or weakly encrypted data) is insufficient to satisfy the claims. Kaeli-Dec, ¶135. He further argues that my citation to background evidence covers many references that do not involve snapshots. *Id.* (citing EX1002, ¶162). But Dr. Kaeli does not appear to dispute the relevant facts that discussed here. Moreover, the passage of my declaration that he cites merely laid the groundwork for my subsequent discussion of weak encryption schemes and weak passwords. EX1002, ¶¶162-65. Dr. Kaeli does not appear to address my further discussion of weak passwords or personally identifiable information. *Id.*, ¶165.

44. I also disagree with Dr. Kaeli's arguments about startup services. Dr. Kaeli argues that I did not show that Veselov's or Mohanty's CVE assessment detects "specific CVEs" associated with startup services. Kaeli-Dec, ¶¶124-25. But he does not dispute my explanation that a POSA would have understood that

typical CVE assessments routinely looked for common startup services. Nor does he dispute my explanation of why a POSA would have expected the types of assets described in Veselov to include the types of startup services that are identified by such assessments. EX1002, ¶¶166-67. Dr. Kaeli argues that the exemplary CVEs I discussed to illustrate that point do not “reference ‘services to be run at startup,’ provide any indication of what services are run on startup, and do not explain what types of security assessments will detect the particular vulnerabilities.” Kaeli-Dec, ¶124. But CVE definitions describe the security flaw, not the typical use case for a given piece of software or the assessments that will look for that CVE. Regardless, Dr. Kaeli’s assertion does not address my point that these CVEs are illustrative examples of typical startup services on the types of systems described in Veselov. EX1002, ¶¶166-67. In other words, Dr. Kaeli’s analysis does not address how the POSA—who would have been familiar with typical virtual assets and CVE assessments—would have understood the references’ teachings. Dr. Kaeli also repeats his assertion that CVEs based on startup services stored on the virtual disk cannot be the recited cyber threats because they are vulnerabilities. Kaeli-Dec, ¶134. But CVEs and cyber threats are not mutually exclusive, as I explained above. Dr. Kaeli argues that my startup-services discussion does not address “snapshot” analysis (*id.*), but the combination expressly involves snapshot analysis

(EX1002, ¶¶147-48) and then went on to discuss startup services in that context (*id.*, ¶¶166-67).

45. Lastly, I disagree with Dr. Kaeli's arguments involving Wiz's DSPM solution. He argues that it was not obvious or well known to detect data-based cyber threats in the cloud by analyzing snapshots because Wiz's marketing for its DSPM solution states "securing cloud data is a difficult problem" and another documents notes the recent emergence of the DSPM product category. Kaeli-Dec, ¶136 (citing EX2154, 3; EX2155). But data-based threat detection did not originate with the DSPM label, and the evidence shows that DSPM covers far more than just detecting data. *Infra*, §X.B.5. Additional evidence confirms that detecting data-based cyber threats was well known before 2019 (including evidence discussing sensitive-data detection in snapshots), and none of Orca's evidence indicates that detecting such threats via snapshot analysis presented a meaningfully different technical challenge than detecting such threats via other known mechanisms. *Infra*, §X.B.1. That is not surprising since Veselov's snapshot analysis process is just a different way of examining the same types of data collected and evaluated by other assessment approach. *Supra*, §VI.B.

2. Elements 1.5-1.6/14.5-14.6/15.5-15.6

46. Dr. Kaeli repeats his argument that Mohanty only considers “external” threats, not threats that are based on stored data. Kaeli-Dec, ¶¶141-47. I previously explained why that argument is incorrect. *Supra*, §VI.C.1.

3. Elements 1.7/14.7/15.7

47. In my previous declaration, explained why the combination teaches or suggests element 1.7, which recites “reporting at least some of the determined plurality of potential cyber threats as alerts prioritized according to their associated risks.” I disagree with Dr. Kaeli’s arguments for this element.

48. I disagree with Dr. Kaeli’s arguments addressing Veselov. Dr. Kaeli argues that Veselov merely delivers results and does not suggest “alerts.” Kaeli-Dec, ¶¶149-52. But this focuses on verbatim use of the word “alert” rather than considering what Veselov conveys to a POSA. Dr. Kaeli argues that I incorrectly equated reporting results with reporting alerts and that these terms are not synonymous. Kaeli-Dec, ¶¶150-51. However, I did not simply equate the term result with the term alert, nor is Veselov limited to a description of reporting results generally. Rather, I discussed Veselov’s broad description of 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, ¶151 (alerts “inform stakeholders of a situation that needs attention”). Dr. Kaeli’s assertion that Veselov does not teach reporting alerts 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.

49. I also disagree with Dr. Kaeli’s arguments addressing Mohanty. He argues that I did not provide a reason to combine Veselov and Mohanty with respect to this claim element (Kaeli-Dec, ¶153), but the reasons to combine that I discussed squarely address prioritized alerts. EX1002, ¶¶109-17. He repeats his argument that Mohanty only detects “external” threats rather than threats based on stored data (Kaeli-Dec, ¶¶155-56), which I addressed above. *Supra*, §VI.C.1. He repeats his argument that Mohanty uses agents assessing real-time data, which does not apply to analysis of past snapshots (Kaeli-Dec, ¶156), which I also addressed above. *Supra*, §VI.B.

50. I also disagree with Dr. Kaeli’s secondary-considerations arguments. He argues that the evidence indicates nonobviousness of the claimed alerts in

combination with other elements. Kaeli-Dec, ¶157. But as I discussed further below, the evidence shows that prioritized alerts as recited were well known, including in the context of agentless scanning. *Infra*, §§X.B.1. Nor does Orca's evidence indicate need, skepticism, praise, copying, or success tied to the combination of this well-known technique with other claim elements. *Infra*, §§X.B.2-6. That is not surprising since snapshot-based assessments were known to consider the same types of data as a basis of threat determination, consider the same types of contextual information, and determine/report the same types of security issues presenting context-sensitive risk levels. *Supra*, §VI.B; *infra*, §X.B.1.

4. Claim 13

51. I previously explained why the combination of Veselov and Mohanty teaches or suggests all aspects of claim 13. EX1002, ¶¶226-35. 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 of the respective protected virtual cloud asset.” Dr. Kaeli makes several arguments in response (Kaeli-Dec, ¶¶161-68), but I respectfully disagree with his assessment.

52. Addressing the “cloud” aspect, he argues that “computing resource service provider” does not typically refer to a cloud provider. Kaeli-Dec, ¶162.

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).

53. I also disagree with Dr. Kaeli’s argument that Figures 5A/B of Veselov do not involve the service provider or APIs. Kaeli-Dec, ¶162. 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, ¶162)—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, ¶162. I explained this in my previous declaration (EX1002, ¶227), and Dr. Kaeli does not appear to respond to that analysis.

54. 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, ¶162 (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, ¶162), but I did not provide an inherency argument, and I understand that this is not a requirement for obviousness.

55. Regardless, Dr. Kaeli is incorrect in his assertion that a POSA would not have viewed Veselov as using APIs to obtain snapshots. Kaeli-Dec, ¶162. 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, ¶227. 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, ¶228), and Dr. Kaeli does not appear to dispute to that explanation.

56. 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 of the respective protected virtual cloud asset,” 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, ¶231. 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.*, ¶¶232-33. Dr. Kaeli primarily responds to the first argument and provides only a brief comment on my analysis for the second. Kaeli-Dec, ¶¶163-64. I address each in turn below.

57. 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, ¶163. 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, ¶163. 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, ¶¶223, 231. Dr. Kaeli also argues that I did not explain how the combination “would *benefit* from such information.” Kaeli-Dec, ¶163. 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, ¶231.

58. I also disagree with Dr. Kaeli's argument that my analysis is "belie[d]" by Veselov's system use of VM identification. Kaeli-Dec, ¶164. In particular, Dr. Kaeli argues that Veselov simply identifies the VM and "has no need" to locate the virtual disk. *Id.*; *see also id.*, ¶165 (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 dvia 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, ¶165), 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). EX1133⁹ (“The volume and instance must be in the same Availability Zone.”).

59. 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, ¶166. But I cited this in

⁹ 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.

support of my explanation of how the POSA's familiarity with the relevant technology—in this case, cloud-management-console querying being a common aspect of the type of VM management discussed in Veselov—informs the POSA's understanding of Veselov's teachings. EX1002, ¶¶230-31. Dr. Kaeli argues that EX1031 (NIST IT Asset Mgmt) only refers to networked devices and software, not virtual disks of VMs, (Kaeli-Dec, ¶166), 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”).

60. I also disagree with Dr. Kaeli's arguments addressing the background exhibits from my discussion of claim 12 (locating the snapshot includes determining a virtual disk allocated to the respective asset). EX1002, ¶¶222-24 (citing, *inter alia*, EX1053 (Prahlad), ¶¶77-79, 81-82, 87-88, 92, Fig. 7; EX1020 (Sancheti), 21:9-41; EX1051 (Workstation Pro), 47, 125). Dr. Kaeli argues that the 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, ¶¶167-68. 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.

61. 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, ¶168 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, 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, 445-46. Dr. Kaeli similarly argues that EX1053 (Prahlad) and EX1020 (Sancheti) do not

address determining virtual-disk locations. However, EX1053 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 similarly discusses “discover[ing] and identif[ying] the virtual disks ... associated with the virtual machine.” EX1020, 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.

62. It is therefore still my opinion that claim 13 was obvious over Veselov and Mohanty.

VII. Ground 2

63. In Ground 2, I discussed the further combination with Ranum as applied to certain limitations of the independent claims (elements 1.4 and 1.7) and dependent claims 2 and 4. EX1002, ¶236. Dr. Kaeli argues that the POSA would not have been motivated to make the combination with a reasonable expectation of success (Kaeli-Dec, ¶¶171-72), that the combination does not satisfy element 1.4

(*id.*, ¶¶173-79), and that the combination does not satisfy claims 2 and 4 (*id.*, ¶¶180-82). I disagree.

64. First, I disagree with Dr. Kaeli's arguments against the reasons to combine. He argues that Ranum involves active scanning and passive network monitoring, which are "incompatible" with snapshot analysis "because snapshots lack active processes." Kaeli-Dec, ¶172 ("[T]here is no mechanism to communicate packets or other messages to a snapshot."). This is essentially the same argument Dr. Kaeli made against the combination of Ground 1, as Dr. Kaeli treats the combination as if it involves wholesale incorporation of Ranum's scanning components and their process of collecting data for evaluation. That is not what I proposed, nor is it how a POSA would have approached the references' combined teachings. Among other aspects of the combination, I explained Ranum's known approach of identifying potential security threats by detecting unexpected changes in stored files, including "computing cryptographic hashes of files contained on the asset and comparing those hashes to hashes of files contained on a baseline system (e.g., the state of the asset from an earlier scan)." EX1002, ¶239; EX1093 (Ranum), 17:37-48, 22:41-50. I further explained that this type of comparison to detected unexpected changes was well-known and predictable in the context of snapshot/image analysis. EX1002, ¶239; EX1040 (Ammons), ¶¶23, 40, 45-49 (discussing "drift detection" involving image analysis to detect non-

compliant changes, including malware); EX1072 (Derbeko), Abstract, 4:1-22 (comparing successive snapshots); EX1097 (Loureiro), ¶¶44-45. The combination thus applies Ranum’s teachings in a known, predictable way in the context of snapshot analysis (e.g., detecting unexpected changes in the files stored in the snapshot), and I explained known benefits and the POSA’s reasonable expectation of success in doing so. EX1002, ¶¶240-45. Dr. Kaeli thus does not address the actual combination I discussed or my analysis of the POSA’s reasons for making that combination.

65. I also disagree with Dr. Kaeli’s arguments addressing element 1.4. He disputes the relevance of Ranum’s additional teachings involving determining cyber threats that are based on the data types I previously discussed in Ground 1. Kaeli-Dec, ¶¶173-76. But he relies on the same misapprehension of the combination that I discussed above, assuming that the combination requires wholesale incorporation of Ranum’s scanner components. *Id.* Dr. Kaeli also argues that Ranum does not consider the same types of data considered in snapshot analysis. Kaeli-Dec, ¶177. But again he focuses on the “way” the data is collected rather than what the data is. *Id.* (references inspect assets in “fundamentally different ways”). Dr. Kaeli argues that I admitted that the data considered by Veselov and Ranum “look different.” Kaeli-Dec, ¶¶177-78. As can be plainly

seen in the transcript Dr. Kaeli quotes, I explained that the way the references collect the data looks different but that the data itself was similar:

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

A. While they are -- on the face of it, it looks different, **they rely on similar information. That's important here. So we are not relying on different set of information.**

EX2172, 251:18-252:2 (emphasis added); *see also id.*, 252:3-18 (references consider “files”), 252:19-253:17 (Ranum considers data similar to information in snapshot). Dr. Kaeli also quotes my comment about the data not being “identical.” Kaeli-Dec, ¶¶177-78. As I explained in the quoted testimony, however, while there are *some* other types of data in network traffic, the data that is ultimately compared (e.g., files) is the same type of information included snapshots. EX2172, 251:18-252:18. Dr. Kaeli is thus focusing on extraneous factors rather than recognizing, as a POSA would have, the comparable nature of the data that is ultimately assessed and the straightforward, predictable nature of applying Ranum’s teachings to snapshot analysis.

66. I disagree with Dr. Kaeli’s arguments addressing claims 2 and 4. He argues that Ranum does not consider data stored on a virtual disk (Kaeli-Dec, ¶181), but the combination applies Ranum’s teachings to snapshot-based

assessments, and snapshots capture data stored on virtual disks. EX1002, ¶¶148-53, 236-45; *see also id.*, ¶¶261 (“Thus, in this combination, the data being analyzed would be the data in the snapshot of the virtual disk.”), 264. Dr. Kaeli argues that Ranum’s cryptographic hashes “are not used as part of a snapshot analysis” (Kaeli-Dec, ¶181), but again, Dr. Kaeli is not addressing the combination, which applies those teachings to snapshot analysis. Dr. Kaeli points to an example in which the hash-comparison baseline is from a clean system shipped from a vendor. Kaeli-Dec, ¶181 (citing EX1093, 17:9-23, 17:49-54). But Dr. Kaeli’s citations to Ranum notably skip over the example I discussed in which the baseline comes from a previous scan of the target asset. EX1093, 17:37-48; *see also id.*, 17:62-18:1, 22:2-34; EX1002, ¶¶250, 256, 263. Dr. Kaeli claims I admitted that the comparison baseline is not the running system in the network and is not comparable to the security assessment of the virtual disk. Kaeli-Dec, ¶181. However, I made the *opposite* point. EX2172, 259:19-268:17. Regardless, Dr. Kaeli’s argument is again not addressing the teachings as applied in the combination. EX1002, ¶264 (“POSA had reason to combine Ranum’s teachings with the snapshot-based assessments taught by Veselov such that the baseline hashes and the current hashes are of at least one area of the virtual disk represented by the snapshot”). Dr. Kaeli repeats his argument that Ranum’s active scanners would not be used in a snapshot (Kaeli-Dec, ¶¶181-82), but I addressed this point

above and explained why Dr. Kaeli mistakenly treats the combination as involving wholesale incorporation of Ranum's scanner components.

67. Lastly, I understand that Orca claims that I attempted to evade questions during cross-examination and stated that "I don't want to give you any precise answer." POR, 64-65 (citing EX2172, 272:21-273:18). This is inaccurate. During this exchange, I was asked to admit that I did not believe Ranum taught other types of data recited in element 1.4 beyond those I addressed in my declaration. EX2172, 271:3-15. I began explaining why my discussion of exemplary disclosures in Ranum was not an admission that other data types were not disclosed, at which point I was cut off as Orca's counsel interjected "I mean, this could take an hour if -- if we have to go through -- because, I mean...." *Id.*, 271:16-273:8. In response, I stated that I did not want to be imprecise just for the sake of saving time. *Id.*, 273:9-18. The "any precise" language in the transcript is a transcription error. *Id.* (emphasis added). As is clear in the context of the full exchange, I stated that "I don't want to give you [an imprecise] answer in the sake of saving time." *See id.*, 271:3-273:18. I was similarly giving complete answers in the additional passages cited by Orca. EX2172, 273:9-276:2.

VIII. Grounds 3-4

68. In Grounds 3-4, I discussed the further combination with Seo (with or without Ranum) such that Ground 3 builds upon Ground 1, and Ground 4 builds

upon Ground 2. EX1002, ¶266. These combinations discuss Seo’s teachings of well-known techniques for identifying illegitimate files and the straightforward application of those teachings to snapshot analysis. *Id.* These grounds address claim 3, which recites the method of claim 1 “further comprising detecting the data stored on the virtual disk by determining added or changed files on the virtual disk without a corresponding installation process.” Dr. Kaeli disputes the reasons to combine but does not dispute that the combination, if made, teaches the additional limitations of claim 3. Kaeli-Dec, ¶¶161-63.

69. I disagree with Dr. Kaeli’s arguments against the combination. Dr. Kaeli distinguishes Seo from Veselov because Seo focuses on “*preventing attacks ... in an artificial intelligence context*” whereas Veselov detects vulnerabilities in a snapshot. Kaeli-Dec, ¶187. But both references address cybersecurity and vulnerability detection, a key purpose of which is of course to prevent attacks. EX1002, ¶268. As to Seo’s discussion of “artificial intelligence” (“AI”), the relevant teachings do not involve neural networks or similar models, but rather discuss a simple, well-known principle—files without installation records are more likely to be illegitimate—that is readily applicable to snapshot analysis. EX1002, ¶¶266, 269-70. Dr. Kaeli does not dispute my explanation that detecting potential threats by checking whether added/changed files have corresponding installation records was well known and easily applied to snapshot analysis. Kaeli-Dec,

¶¶187-91; EX1002, ¶¶269-70 (citing EX1098, Abstract, 1:64-66, 6:6-17; EX1099, ¶¶65, 82). Instead, he argues that I did not indicate why a POSA would have looked solely to Seo’s installation-record-check teachings that are supposedly “based on Seo’s artificial intelligence network.” Kaeli-Dec, ¶187; *see also id.*, ¶191 (similar arguments for Ground 4). But Dr. Kaeli’s premise is incorrect since Seo does not use AI to check for installation records. The AI aspect is part of an earlier step, and the subsequent installation-record check is just a straightforward check of “whether there is an installation history of an application.” EX1094, ¶27; *see also id.*, ¶¶125, 169. As to why a POSA would have applied that teaching, I provided a detailed explanation that Dr. Kaeli does not address. EX1002, ¶¶269-74. For example, I discussed Seo’s explanation of the well-known reason for the check (EX1094, ¶125), multiple motivating factors, and the POSA’s recognition of the technique’s applicability to security assessments like Veselov’s. EX1002, ¶¶269-74.

IX. Grounds 5-6

70. In Grounds 5-6, I explained how dependent claim 11 was obvious in view of the combination of Veselov, Mohanty, and Hutchins (with or without Ranum). EX1002, ¶¶281-90. Claim 11 further recites “wherein mitigating a potential cyber threat includes at least one of: blocking traffic from untrusted networks to the respective protected virtual cloud asset, halting operation of the

respective protected virtual cloud asset, or quarantining the respective protected virtual cloud asset.” In particular, I addressed Hutchins’ teaching of the commonplace mitigation technique of quarantining. Dr. Kaeli briefly disputes the reasons to combine but, aside from referring back to his previous arguments from Grounds 1 and 2, he does not otherwise dispute that the combination satisfies the claim limitations. Kaeli-Dec, ¶¶197, 199.

71. I disagree with Dr. Kaeli’s arguments against the combination. Referring to his arguments from Ground 1, he argues that Hutchins’ use of a snapshot to replicate a VM is “fundamentally different than Mohanty’s analysis, which focuses on VA scanners to continuously monitor target resources.” Kaeli-Dec, ¶195. As I discussed above, however, Dr. Kaeli mistakenly assumes that the combination involves incorporation of agents, and he focuses on extraneous differences between collecting data from snapshots and collecting data via agents while overlooking that the relevant data collected is equivalent. *Supra*, §VI.A. Veselov also shows that snapshot-based duplicate VM assessments and direct analysis of the underlying asset data are alternatives in the same general assessment process. EX1007, 3:20-4:4. In any event, Dr. Kaeli’s assertion does not address the actual combination, which applies Hutchins’ well-known quarantine teachings to the previous Veselov-Mohanty combination from Ground 1.

72. I also disagree with Dr. Kaeli's separate arguments addressing Veselov. He argues that using Hutchins' "resource-intensive" analysis techniques like virus scanning "would require additional resources beyond those provided for in Veselov." Kaeli-Dec, ¶196. 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, ¶¶283-87. 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 types, and virus scanning was commonplace. EX1007, 2:35-41. Dr. Kaeli argues that a POSA wouldn't have used virus scanning in Veselov's 5A/B approach and repeats the inaccurate characterizations of those embodiments that I previously addressed. Kaeli-Dec, ¶196; *supra*, §VI.C.3 (discussing embodiments' involvement of cloud provider and APIs). I disagree since virus scanning was similarly known in the context of snapshot/image-based assessments. *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 quarantining does not depend on which of Veselov's analysis approaches is used

X. Secondary Considerations

73. 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.

74. 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.

75. 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 X.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.

76. 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*, §X.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*, §X.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*, §§X.B.2-6).

77. 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

78. 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.

79. 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.

80. 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 alerts for threats detected by analyzing the snapshot.

81. Mr. Shua’s discussion of the data plane path also undercuts coextensiveness. The claims recite analyzing the at least one snapshot to determine the existence of a plurality of potential “cyber threats.” Orca’s technical brief lists seven steps in the data-plane path:

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. Yet Mr. Shua only identifies three of these as involving the determination of cyber threats. Shua-Dec, ¶45 (pointing to “Configuration Scanning,” “Weak Password Detection,” and “Sensitive Information Scanning”).

82. 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]

83. While I understand that this lack of coextensiveness means that the presumption of nexus has not been established, I also briefly address the additional requirement that the patentee prove that the product embodies the claims. Dr. Kaeli adopts Mr. Shua's testimony that the product supposedly practices the claims, but in my opinion, Mr. Shua has not shown this. Kaeli-Dec, ¶202. For example, Mr. Shua's position is inconsistent with Dr. Kaeli's arguments. Dr. Kaeli argues that the claims do not encompass vulnerabilities or CVEs. Kaeli-Dec, ¶¶120-24. [REDACTED]

[REDACTED]

B. Orca's Evidence

1. Known Techniques/Features

84. Throughout his discussion of secondary considerations, Dr. Kaeli cites various documents referencing aspects of Orca's platform and argues 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

85. 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, 85-89. 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

86. Using cloud APIs to facilitate agentless workload scanning was also known. Mr. Shua claims that agentless API-centered approaches were relatively new in 2020 and that “Orca ... introduced it.” Shua-Dec, ¶126. As I explained above, however, 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 (Prahlaḍ), ¶¶79-94. Cloud-API-mediated interactions were also a fundamental aspect of cloud systems, including cloud-security platforms.

87. 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]

88. 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

10

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

analysis (Shua-Dec, ¶¶147-53), 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 corresponding aspects of the claims. *Id.*, ¶¶31-45. As an initial matter, 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]

In other words, you cannot generate a snapshot in AWS via the standard API call by providing only a VM identifier. 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

89. 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

90. 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).

91. 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....”), 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

92. Dr. Kaeli argues that certain exhibits are tied to sensitive-data detection, but 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, 1-7.¹¹ 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.” *Id.*, 1-2 (“credit card number” example). It would “[n]otify users with an on-screen, pop-up window.” *Id.*, 3.

93. This example is not an outlier. Sensitive-data detection was known in a variety of platforms such as Forcepoint and McAfee Skyhigh Security Cloud.

¹¹ 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.

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.”).¹² Lateral movement risk—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

¹² 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.

(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).¹³ 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 other computing resources; permission settings;” software/configurations, logs, etc.).

94. 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

¹³ 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.

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.

95. 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

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

97. 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,

¶¶203-09; *see also* Shua-Dec, ¶¶81-90 (citing EX1001, 1:47-2:17; 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, ¶¶204-05; Shua-Dec, ¶82. 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, ¶¶210-11 (citing EX2076, 7, 11, 31; EX2138, 1; EX2135, 1, 3); *see also* Shua-Dec, ¶¶8-10, 82-84. As I explained above, however, this argument overlooks that agentless workload scanning (including such scanning for powered-off assets) was already known. *Supra*, ¶X.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, ¶¶149-55), 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.

98. Dr. Kaeli and Mr. Shua mention other aspects like sensitive-data detection, API usage, and targeted vulnerability alerts, but these were also well known, as I explained above. *Supra*, §X.B.1; Kaeli-Dec, ¶¶207, 209-12 (citing EX2152, 5; EX2135, 3; EX2076, 31; [REDACTED]); Shua-Dec, ¶¶90-91 (citing EX2076, 7; EX2103, 1). They were also known 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 sensitive-data detection and targeted/prioritized threat alerts were commonplace. *Supra*, §X.B.1; EX1002, ¶¶42-47 (discussing well-known cloud computing and cloud APIs), 55 (discussing routine use of static-analysis tools to detect sensitive data), 64-70 (discussing well-known alerting techniques).

99. 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, ¶87. Mr. Shua also cites a discussion of Qualys's product as describing limitations of agents (*id.*, ¶88), 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),¹⁴ the author states that “[c]ompanies that once championed agentless solutions are now all building agents.” EX1125, 3; EX2076, 2.

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

¹⁴ EX1125 is a correct copy of a webpage obtained from <https://softwareanalyst.substack.com/p/redefining-cnapp-a-complete-guide>.

3. Skepticism

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

102. First, Dr. Kaeli and Mr. Shua focus primarily on the supposed skepticism of agentless workload scanning. Kaeli-Dec, ¶¶214-18 (citing EX2076, 7-8; EX2103; EX2106, 1; EX2135, 2; [REDACTED]); Shua-Dec, ¶¶92-94. As I previously discussed, however, agentless workload scanning (including cloud-API-based agentless workload scanning) was known. *Supra*, §X.B.1.

103. 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 being truly indicative of industry skepticism. Kaeli-Dec, ¶215; 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]

104. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

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

4. Praise

106. 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, ¶¶220-24; Shua-Dec, ¶¶127-44.

107. 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.”).

108. Mr. Shua touts other alleged praise that was made by Orca’s own partners and investors. Shua-Dec, ¶¶129-30, 135, 140-42; 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.¹⁵ It indicates that the TAG Security Report submitted as EX2118 is part of Orca’s marketing material. The other

¹⁵ EX1125 is a correct copy of a webpage obtained from <https://softwareanalyst.substack.com/p/redefining-cnapp-a-complete-guide>.

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, 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.").

109. 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*, §X.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*, §X.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, ¶¶135, 140-44 (describing awards and concluding “I believe Orca received the above awards because of the unique inventions claimed in the ’926 Patent and used in SideScanning™”).

110. 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, ¶220 (citing EX2119, 9-10; EX2120, 6-7; EX2123); Shua-Dec, ¶¶132-40 (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*, §X.A. Deep visibility and accessing cloud configurations to assess workloads’ block-storage out-of-band was known. EX2120, 6-7; *supra*, §X.B.1. The “agentless approach” was known. EX2123; *supra*, §X.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, ¶¶133, 140;

EX2120; EX2127-2128. He gives no meaningful specifics for the SINET16 or Gartner “Cool Vendor” listings. Shua-Dec, ¶¶134, 136; 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, ¶139; EX2123; EX2125; *supra*, §X.B.1.

111. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

112. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

113. Lastly, with respect to the additional cited evidence cited by Dr. Kaeli (Kaeli-Dec, ¶223), 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 (EX2076, 7, 11-12, 31), but this approach was known, as Veselov itself shows. *Supra*, §X.B.1. Mr. Shua also quotes the report that I understand was commissioned by Orca (as I discussed above). Shua-Dec, ¶¶129-30; EX2118, 3-7. This is not an objective assessment for the reasons I discussed above, and even if it were, it

discusses known features of agentless scanning that I previously discussed. *Supra*, §X.B.1.

5. Copying

114. 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.

115. 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

116. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

117. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

16 [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

118. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

119.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

120.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

121. 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. EX1137.¹⁷ 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

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

[REDACTED]

[REDACTED]

¹⁷ 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] Kaeli-Dec, ¶226. However, Mr. Shua indicates that he had already heard about Wiz’s agentless workload scanner in “early 2020.” Shua-Dec, ¶116. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

123. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

124. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

125. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

126. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

127. I also disagree with Dr. Kaeli's argument that Orca and Wiz's products are substantially similar. Kaeli-Dec, ¶¶260-91. He argues that Wiz's platform is similar to Orca's with respect to elements 1.5-1.7 but does not explain how the aspects of Wiz's documents supposedly indicate this similarity. Kaeli-Dec, ¶¶288-91. For example, Mr. Shua quotes a general statement that Wiz prioritizes/stops "attack paths," provides "ongoing monitoring, alerting, and policy

enforcement,” and “build a single prioritized queue” of attack paths, but he does not explain how this supposedly shows that Wiz’s platform is similar to Orca’s with respect to risk determination or threat prioritization. *Id.*, ¶289 (citing EX2155, 3-4). Mr. Shua argues that Wiz provides ongoing alerting and reports issues as alerts. Shua-Dec, ¶¶90-91 (citing EX2156, 4; EX2157, 4). Again, however, he does not explain how this supposedly shows that Wiz’s platform is similar to Orca’s with respect to reporting cyber threats as alerts prioritized according to their associated risks, as recited. [REDACTED]

[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, ¶263 (citing EX2026; EX2115).

128. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] Orca's own marketing material distinguishes the

products. EX1126.¹⁸ [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

¹⁸ 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]

[REDACTED]

129. 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).

130. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

6. Commercial Success

131. 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.

132. 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, ¶292. 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*, §X.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*, §§X.A, X.B.1; Shua-Dec, ¶¶99-110.

133. 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, ¶¶99-110. 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, ¶¶101, 105; 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*, §X.B.1. Mr. Shua points to a quote from a "GGV managing partner" (EX2109, 2), 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, ¶¶106-07; 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, ¶110; EX2113, 3.

134. 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, ¶293. Likewise, he seems to imply that Wiz’s revenue growth in 2023 and 2024 was due its supposed copying of DSPM from Orca, but he provides no analysis supporting this suggestion, and as I previously discussed, the evidence does not support his premise that Wiz copied DSPM from Orca. *Id.*; *supra*, §X.B.5. Moreover, Orca’s evidence 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. I address additional evidence at the end of my discussion of Wiz's product development. *Supra*, §X.B.5.a. Thus, in my opinion, the evidence does not show that Wiz's success has a nexus to the claims.

XI. Concluding Statements

135. 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.

136. 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.