

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

WIZ, INC.,
Petitioner,

v.

ORCA SECURITY LTD.,
Patent Owner.

IPR2024-00863

PATENT NO. 11,663,031

DECLARATION OF ANGELOS STAVROU

Table of Contents

I. Background	1
II. Materials Considered	2
III. Level of Ordinary Skill in the Art.....	4
IV. Legal Principles.....	5
V. Claim Construction	5
A. “location”	5
B. “(i) taking at least one snapshot, and (ii) requesting taking at least one snapshot of the virtual machine at rest	6
C. Analyze/Analyzing “the at least one snapshot”	10
VI. Ground 1.....	11
A. Orca’s Embodiment-Mixing Argument	11
B. The Veselov-Price Combination	16
C. The Combination Satisfies the Claims	25
1. Elements 1.2-1.4/9.2-9.4/16.2-16.4.....	26
2. Elements 1.5-1.6/9.6-9.7/16.6-16.7.....	34
3. Element 1.7/9.8/16.8	36
4. Claims 4/12	37
VII. Ground 2.....	38
A. The POSA’s Reasons to Combine	39
B. The Limitations of Claims 2 and 10.....	44
VIII. Ground 3.....	46
IX. Secondary Considerations.....	48
A. No Presumption of Nexus	50
B. Orca’s Evidence	54
1. Known Techniques/Features.....	54
2. Long-Felt Need	65
3. Skepticism	68
4. Praise	71

5. Copying	76
6. Commercial Success	88
X. Concluding Statements	91

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 1-16 of U.S. Patent No. 11,663,031 (“’031 Patent” or EX1001) in IPR2024-00863. 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 ’031 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)
1403	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1404	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1405	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1406	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1408	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1413	[REDACTED] [PROTECTIVE ORDER MATERIAL]
1414	[REDACTED] [PROTECTIVE ORDER MATERIAL]

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

III. Level of Ordinary Skill in the Art

7. I understand that Dr. Kaeli provides a slightly different description of the POSA compared to the one I set forth in my original declaration. *Compare*

Kaeli-Dec, ¶30, with EX1002, ¶22. The differences in these descriptions of the POSA do not affect the specific descriptions of the POSA's understanding that I discussed in my previous declaration or in this declaration, including the POSA's familiarity with the relevant technology, the POSA's approach to combining teachings, and the POSA's understanding of the references. Therefore, I will continue to apply the level of ordinary skill that I applied in my original declaration, but my opinions would not change if were to apply Dr. Kaeli's version.

IV. Legal Principles

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

V. Claim Construction

A. "location"

9. I understand that Dr. Kaeli does not dispute the POSA's understanding of the term "location" that I provided in my previous declaration. Kaeli-Dec, ¶¶44-56. Additionally, Dr. Kaeli takes a very broad view of what a "location" encompasses. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] EX1121 (CreateSnapshot).¹

B. “(i) taking at least one snapshot, and (ii) requesting taking at least one snapshot of the virtual machine at rest

10. Dr. Kaeli argues that the “at rest” requirement recited in option (ii) also applies to option (i). Kaeli-Dec, ¶47. While my previous declaration explained why the claims were obvious even with an “at rest” requirement—meaning the claims are still obviousness under Dr. Kaeli’s interpretation—I

¹ I understand that EX1121 is a copy of archived AWS documentation obtained from

https://web.archive.org/web/20141226192222/http://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_CreateSnapshot.html and introduced during Dr.

Kaeli’s deposition. I have compared it with the webpage at that URL and confirm that it is an accurate representation.

disagree with his interpretation of this claim element. As I explain below, Dr.

Kaeli's interpretation is not how a POSA would have understood the claims' plain language read in light of the specification.

11. First, that is not how a POSA would have understood that plain language of the claims. The claim language introduces two options for obtaining the "at least one snapshot," reciting performing "at least one of: (i) taking at least one snapshot, and (ii) requesting taking at least one snapshot of the virtual machine at rest...." This language distinguishes the two options with numbering and punctuation. Each of the two options also independently recites the "at least one snapshot" rather than, for example, reciting at least one of "taking and requesting at least one snapshot...." The separate wherein clause then goes on to explain that the at least one snapshot represents "a copy of the virtual disks of the virtual machine at a point in time." The claim language thus distinguishes two options, with each option providing a separate, coherent description of what the at least one snapshots obtaining via that option would be. The plain language of the claims does not include the "at rest" requirement for option (i).

12. I disagree with Dr. Kaeli's argument about the POSA's understanding of this claim language. He states that a POSA would ask "snapshot of what?" and that only the language in option (ii) ("of the virtual machine at rest") provides that answer. Kaeli-Dec, ¶48. Thus, he argues, if we do not read "of the virtual

machine at rest” into option (i), “taking at least one snapshot” would not inform a POA about the object of the snapshot. *Id.* This is incorrect because the “at least one snapshot” of option (i) is coherent without reading in the language from option (ii), and the wherein clause explains that the one snapshot represents “a copy of the virtual disks of the virtual machine at a point in time.”

13. Second, I disagree with Dr. Kaeli’s argument based on the specification. He argues that his interpretation is supported by the specification because it supposedly makes “consistent” statements that “snapshots are of VMs ‘at rest’ whether taken or requested.” Kaeli-Dec, ¶49. He identifies several passages supposedly supporting that assertion, including descriptions of what at-rest data/machines include and descriptions of securing/analyzing assets (rather than the taking/requesting of the snapshot in particular). *Id.* However, the *analysis* step is recited separately in element 1.6 and does not limit the VM’s state when the snapshot is taken per option (i) of element 1.5. And while Dr. Kaeli identifies an “embodiment” involving a snapshot of a VM at rest (EX1001, 8:42), he omits other descriptions of taking a snapshot when the VM is *not* at rest. The specification illustrates its techniques with two VMs: VM 119-1 (which is an active machine) and VM 119-2 (which is “mostly” inactive). EX1001, 4:5-19. The inactive analysis step of element 1.6 tracks VM 119-2. *Id.*, 5:12-23. As a “mostly” inactive VM, there is a higher probability that VM 119-2 will be inactive

at a given time, but there is no expectation that it will be inactive when the snapshot is taken, so it may be in the “active” state or the “at-rest” state at the point. In fact, the specification expressly indicates that VM 119-2 does *not* need to be at rest when the snapshot is taken. In particular, the specification describes capturing the state of VM 119-2 when it “is **about to** transition into an inactive state or when the same VM 119-2 is at rest.” EX1001, 6:1-9 (emphasis added). In other words, even though the analysis for VM 119-2 is performed while the VM is inactive, the snapshot can be captured when the VM is “about to” transition into an inactive state (meaning that the VM is in the “active” state) or at rest. My interpretation of the claim language that I discussed above is therefore consistent with the POSA’s understanding of the relevant techniques described in the specification, while Dr. Kaeli’s interpretation is an unnatural parsing of the claim language that excludes one of the options described in the specification.

14. Third, I disagree with Dr. Kaeli’s additional argument involving the prosecution history. He points to Orca’s statements during prosecution that a references did not teach taking a snapshot of a VM at rest. Kaeli-Dec, ¶50. However, Dr. Kaeli appears to have overlooked that Orca was addressing an older version of the claims that *did* include an “at rest” for both options. As I explained in my previous declaration (EX1002, ¶¶75-76, this older version of the claims recited performing “at least one of taking a snapshot of the virtual machine at rest,

or requesting a snapshot of the virtual machine at rest.” EX1004 (File History), 123-27 (annotation omitted). After these arguments did not secure allowance, Orca submitted an amendment removing the “at rest” requirement from the first option (though this change was not identified with annotations in the listing of the amended claims). *Compare* EX1004, 123-27 *with id.*, 53-61. In an additional change that *was* marked with annotations, Orca added “wherein during detection of the vulnerabilities by analyzing the at least one snapshot, the virtual machine is inactive.” *Id.*, 53-61. Orca then exclusively argued for nonobviousness based on this wherein clause, and the Examiner allowed the claims. *Id.*, 34-35, 59-61. I therefore disagree with Dr. Kaeli’s argument that Orca’s argument addressing a materially different version of the relevant limitation supports his interpretation of the current claims. To the contrary, this aspect of the prosecution history is consistent with my interpretation since it shows Orca *had* an “at rest” requirement in the first option, then expressly removed it when it did not secure allowance.

C. Analyze/Analyzing “the at least one snapshot”

15. I understand that Dr. Kaeli argues that this limitation does not encompass analyzing a duplicate VM created from the snapshot. Kaeli-Dec, ¶¶52-56. 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. Orca's Embodiment-Mixing Argument

16. 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 them." Kaeli-Dec, ¶95. 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 itself later on to minimize unnecessary redundancy. Veselov is explicit about this, as I previously explained (EX1002, ¶91), and this would have made perfect sense to a POSA since the steps before the analysis (e.g., receiving scan requests, establishing interfaces, using cloud APIs to identify/locate virtual disks for the snapshots, and creating the snapshots) and after the analysis (e.g., processing/reporting the results as alerts) do not depend on whether the analysis step uses the snapshot to instantiate an assessment VM or reads the snapshot as a data file. *See, e.g.*, EX1007, 3:20-4:4, 8:52-10:1-36 (discussing Fig. 2 process, including alternative analysis options for step 208), Fig. 2. It is also consistent with the POSA's understanding of the relevant technology, since the POSA would have recognized that the well-known cloud-API usage that I discussed in my prior declaration—including the different equivalent ways to invoke the cloud APIs (web, console, programmatically, etc.) to achieve the same functionality—were equally applicable regardless of which of Veselov's analysis options were used. I discuss specific issues with Dr. Kaeli's embodiment-mixing arguments below.

17. 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 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, ¶¶95-99; *see also id.*, ¶102. 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. The 3A/B environments are “implemented by a computing resource service provider”—which Dr. Kaeli acknowledges (Kaeli-Dec, ¶95)—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, ¶95. The core premise of Dr. Kaeli’s contention is therefore incorrect.

18. Dr. Kaeli also incorrectly portrays direct snapshot analysis as limited to Figs. 5A/B. He states that my analysis “incorrectly blur[s]” the supposed line between the API-using environments and the direct snapshot analysis of Figs. 5A/B. Kaeli-Dec, ¶¶101-02. But Veselov draws no such line, as I explained above. Dr. Kaeli also argues that Figs. 5A/B are “the only example Veselov discloses that is not directed to assessments of duplicate VMs.” *Id.* This is incorrect. Veselov discusses 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).

19. 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, ¶95 n.3), 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).²

20. Dr. Kaeli also presents an overly narrow view of how APIs are used in Veselov. He claims that in the context of Fig. 1, "Veselov only discusses APIs in relation to a user device transmitting a request to the scanning service for a security assessment." Kaeli-Dec, ¶¶96-97. This is incorrect. While that is *one* way APIs are used, Veselov also explains that the "scanning service 110 may access the computing resource service provider 104, or any associated virtual computing environments and/or computing resources, via the API 120 or another API." EX1007, 5:1-22. In other words, Veselov broadly teaches that the scanning service interacts with the service provider and its resources via API calls. Indeed, that is the paradigm a POSA would have expected since cloud-API calls have long been the standard way that virtually every cloud-security platform—and cloud-based products in general—interact with cloud providers, as I discuss in more detail below. *Infra*, §IX.B.1.

² 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.

B. The Veselov-Price Combination

21. Dr. Kaeli presents several arguments against the combination of Veselov and Price. As I explain below, I disagree.

22. I disagree with Dr. Kaeli's argument that Veselov and Price have "unrelated goals and contradictory teachings." Kaeli-Dec, ¶104. This simply is not true. Veselov and Price both seek to improve the security of virtualized computer resources, and their teachings are closely aligned, as I explained in my previous declaration. EX1002, ¶¶92, 103. Dr. Kaeli states that the combination is contrary to Veselov's supposed "goal" of "keeping VMs online and avoiding downtime." Kaeli-Dec, ¶¶105-06. I disagree. Dr. Kaeli points to Veselov's description of a benefit of snapshot-based assessments in circumstances where the target asset is actively running at the time of the assessment. But a POSA would not have viewed this as limiting Veselov to that circumstance. Dr. Kaeli does not appear to dispute that a POSA would have expected many VMs to be frequently offline (Kaeli-Dec, ¶115), and Veselov says nothing about excluding inactive VMs from its scans. Nor would it make sense for Veselov to exclude suspended or powered-off VMs from its state-independent assessment since a POSA would have understood that the disk data Veselov evaluates during the analysis phase is present (and largely identical) regardless of the VM's state when the snapshot is taken. I explained the state-independent nature of Veselov's technique previously, and Dr.

Kaeli does not appear to dispute this. EX1002, ¶¶104, 107. Moreover, Veselov's straightforward, readily apparent applicability to offline assets does not interfere with benefits in circumstances when the target happens to be actively running. Its state-independent snapshot-based assessment conveys *both* benefits: it does not interfere with active VMs and it allows the scanning system to assess inactive VMs. Veselov also discusses the benefit of reducing data-center costs, which is applicable to both active and inactive VMs. EX1007, 2:42-3:19. In other words, regardless of whether the target is active, the scanning system can bring the snapshot to a different computing environment where the analysis process will not incur fees for data-center resource usage.

23. I also disagree with Dr. Kaeli's assertion that the combination is contrary to Veselov's comment regarding "complexity." Dr. Kaeli states that Veselov disparages systems that add complex installation and monitoring. Kaeli-Dec, ¶107. Dr. Kaeli appears to recognize that Veselov's comment is about installing/monitoring with *agents* on the target asset (not a comment on complexity generally). *Id.*; EX1007, 2:42-3:19. Yet he treats the combination as adding this type of complexity because it supposedly involves "incorporating Price's agent-based monitoring of its online VMs" and therefore "would have required installation of those agents" and associated slowdown/interference. Kaeli-Dec, ¶107. In other words, he treats the combination as requiring the wholesale

incorporation of Price into Veselov. This fundamentally misunderstands the combination, which does not involve incorporating the installation of agents on the target assets. EX1002, ¶¶104, 07. A POSA would have recognized that the actual combination—which includes performing offline scans of already-offline VMs—was not complex. Pet., 23-25. In fact, as I previously explained, it avoids the delays and costs associated with starting up offline assets. EX1002, ¶106.

24. I also disagree with Dr. Kaeli's assertion that a POSA would have had "no reason" to apply Price's teachings of assessing inactive VMs to Veselov because Veselov does not expressly mention inactive VMs. Kaeli-Dec, ¶108. But this "argument from silence" does not address how a POSA would have understood Veselov. For example, as I explained above, a POSA would have expected many target VMs to be at rest and would have known that Veselov's technique was state-independent. A POSA also and would have been familiar with agentless offline assessments, as I explained in my previous declaration. EX1002, ¶62; EX1065 (Thakar), 2:29-33; EX1073 (Ammons Mirage), 1. As to the POSA supposedly having no reason to perform inactive assessments in Veselov, I explained those reasons in detail. EX1002, ¶¶105-06.

25. I also disagree with Dr. Kaeli's arguments against the rationale for combining the references based on distinctions he makes between "images" and "snapshots." Kaeli-Dec, ¶¶109-13. He states that Price "does not mention

snapshots,” contrasts the term “image” and the term “snapshot,” and argues that they address “different classes of vulnerabilities.” *Id.*, ¶¶109-10. He concludes that applying Price’s teachings of analyzing images to Veselov “would not have been straightforward or routine,” that the references’ teachings are incompatible, and that the combination would have been unpredictable. *Id.*, ¶¶111-12. As I explain below, I disagree with each of these assertions.

26. In my opinion, Dr. Kaeli overstates the differences between images and snapshots and focuses on irrelevant distinctions that would not have impacted that straightforward application of image-analysis teachings to snapshot analysis. Before turning to the specific machine image resources that are used in cloud platforms like AWS, it is worth noting that the term “image” 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.

27. Dr. Kaeli appears to focus 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's portrays them as, and it would have been clear to a POSA that machine-image-analysis techniques were directly and readily applicable to snapshot/volume-image analysis. 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), whereas an AMI essentially just includes one or more snapshots (one of which is a boot volume) along with metadata such as a launch parameters. EX1131.³ In fact, an AMI can simply be created from a snapshot. EX1132.⁴ This type of close relationship between machine images and snapshots was common in cloud/virtualization platforms.

³ EX1131 is a correct copy of 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.

⁴ 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>.

28. In any event, it would have been clear to a POSA that the analytical techniques for machine images and container images were directly relevant and readily applicable to analytical techniques for snapshots since the underlying data being analyzed—e.g., files/applications stored in the image/snapshot—is equivalent. EX1007 (Veselov), 3:20-67, 9:42-67; EX1048 (Price), ¶42. The same is true of container images since these also include equivalent underlying data (e.g., files/applications) and were readily assessed via analogous techniques. *See, e.g.*, EX1078, ¶¶3-8, 31. A POSA considering the snapshot-analysis techniques taught by Veselov therefore would not have been driven by whether or not the terms “machine image” or “container image” were literally synonymous with “snapshot” or “volume image.” Nor would a POSA have viewed the differences between machine/container images and snapshots/volume images—e.g., the ability to directly boot a machine image—as undermining the applicability of teachings associated with machine/container-image analysis to snapshot analysis. Rather, a POSA would have recognized that the relevant data, the relevant analysis techniques, and the relevant security issues identified by such analysis were comparable.

29. Dr. Kaeli also overlooks the numerous similarities in the vulnerabilities and security threats that each reference addresses. He claims images and snapshots are used to address “different classes of vulnerabilities,” but

his explanation is that snapshots are used to detect malware present when the snapshot was taken whereas VM images are used to prevent the spread of malware. Kaeli-Dec, ¶110 (citing EX2023, 5:49-58; EX2024, Sec. 2.3). I disagree with his reasoning since, in both examples, he acknowledges the same class of vulnerability: malware. Furthermore, the presence of vulnerabilities and other security issues expose both snapshots and VM images to equivalent security threats. He appears to focus on how the detected malware is addressed after it is detected, rather than the vulnerabilities themselves, which does not make malware in snapshots/images “different classes of vulnerabilities.” Consistent with the POSA’s knowledge, the references also teach that their respective resources are used to address other analogous vulnerabilities/threats via analogous techniques. EX1007 (Veselov), 2:35-41 (discussing CVEs, benchmarks, best-practice packages, static or runtime analysis, configuration assessments, and the like); EX1048 (Price), ¶¶16-17 (scan addresses “policy compliance, vulnerability assessment, malware protection, and other security services”); *see also* EX1078 (Hufsmith), ¶¶33 (CVEs), 82 (static or dynamic analysis and configuration scanner), 147 (benchmarks), 149 (applicability to VM images), 159 (applicability to VMs).

30. I disagree with Dr. Kaeli’s assertion that the background exhibits I discussed in my previous declaration do not support my discussion of the

combination. Kaeli-Dec, ¶¶114-120. I discussed these, for example, to illustrate the POSA's understanding of the routine/predictable 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, ¶¶104-07. As an initial matter, Dr. Kaeli does not appear to be addressing the actual combination that I discussed. He begins this discussion by stating that "[c]ritically," the references do not show motivation to "shut down Veselov's target VMs to take a snapshot, when shutting down Veselov's target VMs would run directly counter to Veselov's goals." Kaeli-Dec, ¶¶114. As I explained above, however, the combination does not rely on shutting down active VMs, but rather performing inactive assessments of VMs that are already inactive or powered off. EX1002, ¶¶104-06.

31. I also disagree with Dr. Kaeli's arguments about these specific exhibits. Discussing EX1066-1067, he does not dispute the relevant point—that a POSA would have expected many VMs to already be offline—but claims that this does not support the combination because Veselov's focus is supposedly limited to active VMs. Kaeli-Dec, ¶¶115, 118. This is the same misunderstanding of Veselov that I discussed above. He focuses on irrelevant aspects of EX1066 (Zhou)—e.g., its discussion of offline patching and consideration of machine memory for patching purposes—but this does not address the relevant point (i.e.,

the POSA's understanding that many VMs in a typical cloud environment would often be offline). Kaeli-Dec, ¶116. He claims the discussion in EX1067 (Jeswani) of exemplary update techniques involving snapshots and images undermines the combination, but nothing in EX1067 supports his position that machine-image-analysis techniques are not applicable to snapshot/volume-image-analysis techniques. *Id.*, ¶117 (citing EX1067, 393-94, 398-99). He argues that EX1081 (Budko) does not mention snapshots, but he does not dispute the relevant point that it illustrates (i.e., the POSA's understanding keeping that having a potentially compromised VM remain offline during the assessment minimizes potential harm). EX1002, ¶105 (citing EX1081, ¶79). He says that EX1073 (Mirage) does not mention snapshots and that EX1015 (vSphere) "demonstrates the gap between snapshots and images" and does not suggest using snapshots for security assessments. Kaeli-Dec, ¶120. But neither of these piecemeal assertions addresses the well-known, predictable nature of offline assessments or the straightforward applicability of machine-image-analysis techniques to snapshot/volume-image-analysis techniques that I previously discussed and that would have been clear to a POSA.

32. Lastly, I disagree with Dr. Kaeli's assertion that a POSA would not have reasonably expected success in combining the references to arrive at the claims. Kaeli-Dec, ¶¶121-22. Dr. Kaeli states that Veselov does not indicate that

the process of taking a snapshot (or the subsequent analysis) is state-independent. *Id.*, ¶121. However, the state-independent nature of taking a snapshot in most cloud/virtualization platforms was widely known—e.g., the standard mechanism for creating snapshots in AWS generates snapshots regardless of whether the VM is active, suspended, or powered off (EX1121)—and Dr. Kaeli acknowledges the evidence showing “snapshots can be taken when a VM is powered off.” *Id.*, ¶120 (citing EX1015). The independent nature of snapshot analysis is readily apparent from Veselov’s description of analysis that—with the exception of one optional technique not applied in this combination that involves mirroring a VM’s current traffic—is performed separately from the target VM and without interaction with that VM. *E.g.*, EX1007, 3:44-67. Indeed, no modification of Veselov’s technique is even needed to arrive at the claims.

33. 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

34. 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 1.2-1.4/9.2-9.4/16.2-16.4

35. I previously explained why Veselov teaches or suggests all aspects of elements 1.2-1.4. In response, Dr. Kaeli makes several arguments attempting to show that Veselov does not satisfy these limitations. Kaeli-Dec, ¶¶124-43. Respectfully, however, Dr. Kaeli’s arguments do not accurately reflect how a POSA would have understood Veselov.

36. I disagree with his argument that Veselov does not teach using cloud APIs to identify virtual disks of a VM. Kaeli-Dec, ¶127. He argues that Veselov does not mention volume identifiers (*id.*, ¶127), but I explained multiple ways that a POSA would have understood Veselov’s process to include or suggest performing the recited identification. EX1002, ¶¶122-31. Even focusing exclusively on the word “identify” were the relevant inquiry—and I understand that in an obviousness inquiry, references are read from the perspective of a POSA for all that they fairly teach, not simply for verbatim use of claim terms—I explained that Veselov discloses having the scanning service “identify” the data elements to include in snapshots (which would encompass identifying virtual disks). EX1002, ¶128 (discussing snapshot tailoring); EX1007, 8:16-45.

37. I also disagree with his subsequent arguments about these snapshot-tailoring disclosures. Dr. Kaeli argues that this does not involve an API (Kaeli-Dec, ¶¶131, 134), but Veselov says that this is performed by the scanning service (EX1007, 8:16-45) and that the scanning service interacts with the provider/environment via APIs (*id.*, 5:1-15). Dr. Kaeli also argues that “data elements” are not virtual disks, (Kaeli-Dec, ¶¶131, 134), but this does not accurately reflect the POSA’s understanding of Veselov’s teachings. A virtual disk was typically the core component of a snapshot, and a POSA would have recognized that it made perfect sense for the scanning service to identify a virtual disk as a “data element” to include in the snapshot. EX1002, ¶128. For example, a VM may have multiple virtual disks, but the snapshot may not include all of these disks. Indeed, in AWS, an identification of a single volume/virtual disk is required when using the standard mechanism for creating snapshots. EX1121 (CreateSnapshot). Likewise, Veselov repeatedly refers to an exemplary snapshot as a “volume” image, and this volume refers to the virtual disk. EX1002, ¶122; EX1007, 9:42-55; 11:30-33. Regardless, identifying more granular “data elements” still suggests identifying the virtual disk in which those data elements reside. EX1002, ¶128. Dr. Kaeli also disputes motivation to modify Veselov to perform the claimed functions, but my analysis explained what a POSA would have understood Veselov to teach or suggest—i.e., a POSA’s understanding of

what the disclosed techniques involve/suggest—not modifications to Veselov’s broadly described process.

38. I also disagree with Dr. Kaeli’s embodiment-mixing arguments. He provides a condensed version of the arguments from the embodiment-mixing section of his declaration. Kaeli-Dec, ¶¶128-129. I disagree with these arguments for the reasons I discussed above. *Supra*, §VI.A.

39. Dr. Kaeli’s narrow portrayal of how APIs are used in Veselov is inaccurate. He states that Veselov’s APIs merely “provide interfaces to a user or otherwise convey user-specific data.” Kaeli-Dec, ¶130. As I noted above, however, Veselov does not merely describe using APIs for user-to-scanner communication. It broadly discloses using APIs for scanner-to-provider/environment communications. EX1002, ¶125; EX1007, 5:1-5, 5:59-60. A POSA would have understood this to include the scanner’s various communications for identifying the virtual disks that I discussed in my previous declaration. *E.g.*, EX1002, ¶125. Dr. Kaeli argues that I “handwave[]” over Veselov’s supposedly limited usage of “environment” and “web interface” APIs. Kaeli-Dec, ¶130. Again, however, Veselov is not so limited. Dr. Kaeli does not explain why he thinks those descriptors connote limited API usage, but regardless, 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. 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. Dr. Kaeli’s argument thus does not consider Veselov’s full disclosure as would have been understood by a POSA.

40. I disagree with Dr. Kaeli’s discussion of the evidence that I cited to illustrate the POSA’s understanding of routine techniques disclosed in Veselov. He mainly contends that this evidence does not show a motivation to modify Veselov. Kaeli-Dec, ¶132. However, as I previously discussed, the analysis does not rely on a modification of Veselov’s process. Rather, I explained what a POSA would have understood Veselov as teaching or suggesting. His footnote responds to specific exhibits by identifying individual aspects of the claims that are supposedly not disclosed (as if each exhibit is a reference being applied in an anticipation ground), but this does not meaningfully address my point. Kaeli-Dec, ¶132 n.6. As I previously explained, these exhibits corroborate my explanation, based on my experience in the field, that cloud platforms typically provided API-

based mechanisms for identifying and locating disks and that a POSA would have understood Veselov's description of its techniques as involving or suggesting these routine techniques. EX1002, ¶127. Dr. Kaeli argues that certain exhibits do not mention APIs, but cloud APIs are the fundamental underpinning of these operations in cloud environments and were routinely used to perform those operations, as I explain further below. *Infra*, §XI.B.1. He argues that EX1020 (Sancheti) does not show identifying virtual disks with APIs, but that is precisely what it discloses. EX1020, 21:9-41 (management component "discovers and identifies the virtual disks ... associated with the virtual machine, for example by using APIs to query the virtual machine host"). In any event, his suggestion that using cloud APIs to identify virtual disks was not routine is simply not accurate. *See infra*, §IX.B.1.

41. I disagree with Dr. Kaeli's argument that identifying a VM means that virtual disks are not also identified. He treats these two identifications as mutually exclusive, arguing that "the fact that Veselov's systems and methods work with and from the identification of VMs (*not virtual disks*)" refutes my explanation of why Veselov also teaches identifying virtual disks. Kaeli-Dec, ¶133. As an initial matter, Veselov is not limited to scans initiated by a request that identifies the target VM. While that is one option that I discussed, the request may also simply identify the user. EX1007, 6:25-38. Regardless, Veselov does not say that the

scanning service generates a snapshot simply by giving the cloud provider a VM identifier, and a POSA would have known that this is not how snapshot generation typically worked in cloud platforms. For example, in AWS, creating a snapshot via the standard API call requires an identifier of the volume (not the VM).

EX1121 (CreateSnapshot). In any event, Dr. Kaeli's treatment of VM identification as precluding subsequent identification of the VM's virtual disks is not consistent with Veselov's teachings of having the scanning service identify the VM's data elements, which I discussed above. EX1007, 8:16-45. Dr. Kaeli further argues that virtual disks and VMs can have different locations. Kaeli-Dec, ¶133. With respect to their general storage location, however, that is not typically the case 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). EX1133⁵ ("The volume and instance must be in the same Availability Zone.").

⁵ 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.

42. I disagree with Dr. Kaeli's arguments responding to the alternative UI-based mappings. EX1002, ¶¶129-31. First, he conflates this with the snapshot-tailoring argument. Kaeli-Dec, ¶134 (citing EX1007, 8:16-45). While they both involve Veselov's broad teachings of API-usage, they are separate arguments relying on different teachings in Veselov with respect to the disk identification itself. EX1002, ¶¶128 (snapshot-tailoring), ¶¶129-31 (UI-based mappings). He devotes most of his response to commenting on extraneous aspects of the citations I provided to show that the cited APIs are "cloud" APIs, but he does not even appear to dispute the relevant point. EX1002, ¶129 (citing EX1007, 2:22-3:19, 19:30-37, 19:60-64, 20:46-56). When he does turn to my explanation of how these cloud APIs are used to perform the recited operation, he repeats the embodiment-mixing argument and his overly narrow characterization of how APIs are used in Veselov, both of which I addressed above. Kaeli-Dec, ¶134; *supra*, §VI.A.

43. I also disagree with Dr. Kaeli's additional arguments regarding element 1.3. He argues that Veselov only "discusses queries" in the context of the snapshot storage service. Kaeli-Dec, ¶135. However, he again appears to be focusing on the verbatim use of the word "query" rather than reading Veselov from the POSA's perspective for all that it fairly teaches. He also argues that Veselov does not query disk locations to facilitate generating the snapshot because Veselov supposedly just obtains previously stored snapshots from the snapshot storage

service. Kaeli-Dec, ¶¶135-37. But Veselov makes clear that the snapshot storage service is optional and that scanning service may also generate a new snapshot. EX1007, 5:9-22, 9:9-25. Dr. Kaeli's only response to my explanation of how Veselov teaches/suggests using cloud APIs to query disk location as part of its snapshot-generating process is that Veselov supposedly does not use APIs to do this. Kaeli-Dec, ¶136. As I previously explained, however, Veselov broadly teaches that the scanning service communicates with the cloud provider and its environment/resources via APIs, which is consistent with standard practices at the time. EX1002, ¶¶134-35; EX1007, 5:1-5. Turning to the corroborating evidence, Dr. Kaeli again responds as if I was proposing a modification of Veselov rather than my actual argument, which explained the POSA's understanding of the well-known, routine way of doing what Veselov describes. EX1002, ¶¶134-35; Kaeli-Dec, ¶138.

44. I disagree with Dr. Kaeli's argument responding to my alternative mapping for element 1.3 involving querying a more general location. He acknowledges the discussion of the virtual network address but argues that Veselov does not mention querying that location. Kaeli-Dec, ¶139. Again, however, he focuses on verbatim use of the word query rather than responding to my explanation of what a POSA would have understood from Veselov's disclosure.

EX1002, ¶136. He also states that the VM and its virtual disks "may" not have the

same location, apparently agreeing that they sometimes would have the same location. Kaeli-Dec, ¶139. In any event, as I explained above, a POSA would have understood that they typically are stored in the same general location. *See, e.g.*, EX1133 (“The volume and instance must be in the same Availability Zone.”).

45. Dr. Kaeli’s arguments regarding element 1.4 largely follow the same pattern, and I disagree for similar reasons. He incorrectly treats my analysis as proposing a modification of Veselov in view of the state-of-the-art evidence and repeats his argument involving the general location from element 1.3. Kaeli-Dec, ¶¶140-42.

2. Elements 1.5-1.6/9.6-9.7/16.6-16.7

46. I disagree with Dr. Kaeli’s opinion that Veselov and Price do not teach limitations 1.5 and 1.6.

47. First, he disputes my assessment of how Veselov teaches “taking at least one snapshot” per option (i) of element 1.5 based on his interpretation of the claim as reading in the “at rest” requirement from option (ii) into option (i). Kaeli-Dec, ¶146. As I explained above, that is not how a POSA would have understood the claim’s plain language read in light of the specification. *Supra*, §V.B. Even if it were correct, I previously explained that at-rest taking of snapshots was well-known and routinely performed and a POSA would have understood that in Veselov’s state-independent technique, while VMs would be active in some

circumstances, in other circumstances the VM would be at rest when the snapshot is taken. EX1002, ¶¶33, 149. Dr. Kaeli responds by repeating his characterization of Veselov as limited to active VMs. Kaeli-Dec, ¶146. As I explained above, that is not how a POSA would have understood Veselov. *Supra*, §VI.B.

48. In any event, I presented additional analysis explaining why Price’s teachings provide an affirmative reason to take the snapshot when the VM is at rest. EX1002, ¶¶150-54. In response, Dr. Kaeli argues that taking snapshots at rest “was counter to conventional wisdom.” Kaeli-Dec, ¶147 (citing EX2117, 1 [REDACTED]). As I previously explained, however, this is wrong since the state-independent nature of taking a snapshot in most cloud/virtualization platforms was widely known. *Supra*, §VI.B; EX1002, ¶33; EX1121 (CreateSnapshot).

Moreover, neither of the exhibits cited by Dr. Kaeli actually indicate that taking snapshots at rest was unconventional. At best, they contrast scans by agents installed on the target (which run when the target runs) with the at-rest protection naturally provided by agentless scanning, which says nothing about conventional approaches to taking snapshots. EX2117, 1 [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]. Dr. Kaeli also refers to his previous arguments against the combination (Kaeli-Dec, ¶148), which I disagree with as explained above. *Supra*, §VI.B.

49. For element 1.6, Dr. Kaeli that Veselov’s snapshot-as-a-data-file analysis option is snapshot analysis, and for the “inactive” aspect he repeats his arguments against the combination. Kaeli-Dec, ¶¶149-51. As I explained above, I disagree with those arguments. *Supra*, §VI.B.

3. Element 1.7/9.8/16.8

50. I disagree with Dr. Kaeli’s argument that Veselov does not teach reporting the detected vulnerabilities as “alerts.” Dr. Kaeli argues that Veselov “does not mention” alerts. Kaeli-Dec, ¶155. 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, ¶¶156-57. Respectfully, however, Dr. Kaeli misstates my testimony and inaccurately characterizes Veselov.

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 description of reporting identified vulnerabilities in a way that brings specific problems (and identified remediations) to the user's attention. EX1002, ¶¶172-73. 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, ¶157 (alerts "notify stakeholders of situations that require 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.

4. Claims 4/12

51. Dependent claims 4 and 12 further recite "wherein the identification of the location of the virtual disks of the virtual machine includes a virtual address of at least one of the virtual disks." I explained two ways that Veselov teaches this limitation. EX1002, ¶¶179-82. In response, Dr. Kaeli repeats previous arguments

that he made for the independent claims. Kaeli-Dec, ¶¶159-63. With respect to the first mapping that I discussed (involving the VM’s virtual address), he repeats his assertion that a VM’s location/virtual address is not a location of virtual disks associated with that VM. *Id.*, ¶161. As I previously explained, that is not how a POSA would have understood Veselov or the claims. EX1002, ¶180; *see also supra*, §VI.C.1 (addressing similar argument for elements 1.2-1.4). With respect to the second mapping that I discussed (involving a more specific virtual address of the disk), he again refers to his previous arguments. Kaeli-Dec, ¶162. He also states that I did not identify disclosure of “obtaining” such an address. *Id.*

However, my declaration explained why Veselov teaches or suggests receiving this location. EX1002, ¶181. Dr. Kaeli also says that I did not explain how this location information would be obtained in a system that takes snapshots of VMs at rest (Kaeli-Dec, ¶162), but the claims do not require this. *Supra*, §V.B. In any event, disk virtual addresses are readily obtained in typical cloud environments regardless of the associated VM’s state. I therefore disagree with Dr. Kaeli’s arguments regarding claims 4 and 12.

VII. Ground 2

52. In my previous declaration, I explained how the combination of Veselov, Price, and Hufsmith teaches the limitations of claims 2 and 10, which require that the previously recited reporting include “indicating priority levels

associated with the detected vulnerabilities” (claim 2) or “providing indications of the detected vulnerabilities based on priority levels associated with the detected vulnerabilities” (claim 10). As I explain below, I disagree with Dr. Kaeli’s arguments.

A. The POSA’s Reasons to Combine

53. First, I disagree with his argument that a POSA would not have viewed Hufsmith’s teachings as applicable to snapshot analysis. Dr. Kaeli argues that Hufsmith’s teachings are not applicable because container images and snapshots are different (and because VM images and snapshots are different). Kaeli-Dec, ¶¶168-69. However, as I have previously explained, containers and VMs are closely related resources, and a POSA would have recognized that the relevant underlying data in container images and VM images—e.g., files, applications, configurations, etc.—and VM snapshots are analogous, as are the relevant techniques for detecting security issues such as vulnerabilities based on that data, making the aspects of Hufsmith that I discussed in my previous declaration readily applicable to the snapshot-based analysis taught by Veselov. EX1002, ¶¶201-04, 207; EX1007, 2:46-54, 3:20-67, 9:42-67 (snapshot includes data such as “file system and file information” and “software packages installed”); EX1078, ¶¶31 (examining image layers to obtain file/package information, then submitting “files/packages and other resources to various vulnerability scanning

engines (including Veracode™ and others enumerated below) for vulnerability assessment”), 149, 154, 159; *supra*, §VI.B. Dr. Kaeli also inaccurately minimizes Hufsmith’s broad teachings of applicability to VMs and VM images. For example, he states that when Hufsmith discusses applicability to other forms of “encapsulated functionality,” it does not elaborate on what those might be. Kaeli-Dec, ¶169 (citing EX1078, ¶149). This is simply not true, as Hufsmith—in that exact same sentence—gives several examples including “virtual machine images, machine images, unikernels, and AWS Lambda™ functions.” EX1078, ¶149; *see also id.*, ¶159 (“This approach leverages IDE abilities for issues specific to containers, images, and files, in some embodiments (though similar approaches are contemplated for virtual machines, orchestration configuration files, serverless configuration files, and the like).”).

54. Dr. Kaeli also focuses on extraneous aspects of container images that do not undermine the combination. For example, he emphasizes that container images have “layers” (Kaeli-Dec, ¶169), but these layers are essentially just a series of deltas relative to the previous layer. The type of analysis I discussed in Hufsmith ultimately analyzes the files/applications stored in these layers to identify and prioritize detected security issues, which a POSA would have recognized as relevant and readily applicable to snapshot analysis. *See, e.g.*, EX1078, ¶¶3-8, 31.

Dr. Kaeli argues that the analyses involve “different actions and techniques” that

supposedly would have led a POSA considering snapshot-based analysis to avoid looking at techniques involving container-image analysis. Kaeli-Dec, ¶169. He does not say what these supposedly different actions and techniques are, but he appears to be approaching the combination as if it is limited to wholesale incorporation of Hufsmith's embodiments. Kaeli-Dec, ¶169 (arguing that POSA "would not understand that Hufsmith's embodiments could simply be appended to virtual machine images"). That is not how I approached the combination—i.e., I did not focus on whether a POSA would robotically copy-and-paste Hufsmith's embodiment into Veselov—nor is that how a POSA would have evaluated what techniques are applicable to snapshot-based analysis when considering Hufsmith's teachings.

55. Dr. Kaeli also quotes an article—written by someone who states "I don't consider myself a security expert"—describing differences between containers and VMs. EX2027. In particular, Dr. Kaeli quotes its statement that "The fundamental difference between containers and VMs is where the scanning occurs, with the analysis and action taken. In the container world, a single scan and detection can account for hundreds to thousands of nodes being remediated." In my opinion, however, this document does not support Dr. Kaeli's argument against the combination. As an initial matter, the author does not discuss snapshot analysis. EX2027. Regardless, the author identifies the difference between

containers and VMs as “where” the scanning occurs (i.e., where the analysis is performed and where the actions are taken). *Id.*, 4. Read in context, the author plainly is not talking about what types of data in the container image are being assessed, what assessment techniques are performed on that data, or the prioritization of identified security problems. Rather, he is contrasting assessment of a VM that “doesn’t necessarily trigger a recursive action for other like VMs” with assessment of container images in a container-image registry where a system can determine, for example, that a base image with a vulnerability has various other derived images (presumably inheriting that vulnerability) that are deployed on numerous production nodes. *Id.* In other words, the author is saying that by analyzing container images in image registries (i.e., “where” the scanning occurs) a set of associated security problems in related containers/container images can be identified and responsive actions taken. This is another extraneous issue that does not address the closely analogous aspects of container images and snapshots (e.g., the underlying data, the assessments performed on that data, and associated prioritization) that would have led a POSA to combine the references’ teachings as I previously discussed.

56. I also disagree with Dr. Kaeli’s argument that Hufsmith does not teach the benefits of alert prioritization. Kaeli-Dec, ¶¶170-71. As with Veselov, Dr. Kaeli focuses on Hufsmith’s limited use of the word “alert” (*id.*, ¶170), but this

does not address how a POSA would have understood and applied the teachings that I discussed. He states that paragraph 6 of Hufsmith does not describe problems with overwhelming vulnerability reports and quotes the first sentence of that paragraph. *Id.* However, he omits the second sentence, which describes the problem I mentioned. EX1078, ¶6 (discussing problem due to “computational and cognitive load associated with appropriately processing surfaced vulnerabilities”). Dr. Kaeli argues that teachings associated with embodiments involving the IDE are not applicable to Veselov because they are supposedly only relevant during the development process. Kaeli-Dec, ¶171. As an initial matter, Hufsmith teaches that its assessments are not limited to development/pre-deployment scenarios. *E.g.*, EX1078, ¶108 (“In turn, the combined threat scores may be considered during application or container development and deployment.”). Regardless, Dr. Kaeli is incorrect regarding how a POSA would have understood the applicability of these teachings. In other words, a POSA would not have approached the combination as a matter of robotically incorporating the full embodiment into Veselov (i.e., literally using source-code annotations). Rather, a POSA would have recognized that the well-known principles discussed in those embodiments—e.g., classifying/prioritizing detected vulnerabilities and conveying these priority levels to the user in alerts—were applicable and useful to VM-snapshot-based assessments such as those described in Veselov. EX1002, ¶¶205-07. Indeed,

Hufsmith expressly teaches that these techniques are broadly applicable to other approaches, including those associated with VMs/images. EX1078, ¶¶149, 154, 159.

B. The Limitations of Claims 2 and 10

57. I discussed several ways that applying Hufsmith's teachings satisfies the additional limitations of claims 2 and 10, including its discussion of using weights/metrics to prioritize vulnerabilities and its discussion of displaying classifications/scores for reported vulnerabilities. EX1002, ¶¶209-13. I disagree with Dr. Kaeli's arguments that Hufsmith does not teach the additional limitations of claims 2 and 10.

58. Dr. Kaeli's discussion of Hufsmith's reporting of weights and metrics is inaccurate. He argues that the weight does not indicate the severity of the risk associated with a given issue, but rather is simply a way to "count" the number of vulnerabilities detected. Kaeli-Dec, ¶175. Respectfully, that is not how a POSA would have understood these weights. As I previously explained, Hufsmith describes metrics and weights associated with respective security issues, and Hufsmith teaches/suggests reporting this information internally and to the user. EX1002, ¶¶209-10; EX1078, ¶¶90-94, 120, 134. Weights are not simply a way of counting vulnerabilities, as Dr. Kaeli contends. Rather, a weight is associated with a particular vulnerability and essentially increases or decreases the importance

assigned to that vulnerability (e.g., based on contextual factors). EX1078, ¶¶90 (discussing “different weights corresponding to different vulnerabilities”), 94 (discussing “context adjusted weights associated with the detected potential vulnerabilities”). Dr. Kaeli appears to conflate Hufsmith’s statement that a risk score for a container image “may be based, for example, on a count of the number of potential vulnerabilities detected” with Hufsmith’s separate statement that “[s]ome embodiments may calculate a weighted sum of detected potential vulnerabilities, and some cases with different weights corresponding to different vulnerabilities or types of vulnerabilities in a taxonomy of vulnerabilities.” *Id.* As to the reporting of the metrics/weights, Dr. Kaeli does not appear to dispute that Hufsmith reports these to the user. EX1002, ¶210. Instead, he argues that I said metrics and weights are the same. Kaeli-Dec, ¶176. This is untrue. As I explained, both are indicative of the associated vulnerability’s priority level, and Veselov teaches reporting both. EX1002, ¶210.

59. I also disagree with Dr. Kaeli’s arguments involving the annotations discussed in the IDE embodiments. He states that these teachings are “inapplicable to Veselov, which only analyzes deployed VMs or snapshots where such source code would be unavailable for such annotating.” Kaeli-Dec, ¶177. As I explained above, however, the combination does not involve a wholesale incorporation of the embodiment into Veselov (i.e., literally performing source-code annotations), nor

is that how a POSA would have approached the combination. *Supra*, §VII.A; EX1002, ¶¶201-07. He also states that global weights are not associated with the vulnerabilities. Kaeli-Dec, ¶178. While my analysis focused on vulnerability-specific priority levels, global weights are also associated with vulnerabilities because they are used to adjust their risk levels even if they are not *uniquely* associated with a single vulnerability. *See* EX1078, ¶¶135-43.

VIII. Ground 3

60. In Ground 3, I explained how dependent claims 6 and 14 were obvious in view of the combination of Veselov, Price, Hufsmith, and Huseinović. EX1002, ¶215. These claims further recite “wherein the snapshot includes a page file of memory associated with the virtual machine, the page file being configured to allow deduction of one or more applications running on the virtual machine.” Dr. Kaeli briefly disputes the reasons to combine but, aside from referring back to his previous arguments from Grounds 1-2, he does not otherwise dispute that the combination satisfies the claim limitations. Kaeli-Dec, ¶¶179-80.

61. I disagree with Dr. Kaeli’s arguments against the combination. He argues that Huseinović “explicitly counsels *against* collecting offline data.” Kaeli-Dec, ¶179 (citing EX1049, 940). He does not explain his position beyond providing partial quotes of Huseinović’s introduction. *Id.* In any event, I disagree that Huseinović counsels against the combination that was actually proposed. As

an initial matter, the cited passage describes prior forensic investigations in the pre-virtualization era, where powering off a physical machine caused volatile memory such as RAM to be lost. EX1049, 940. After this preface, Huseinović explains how this can be avoided in virtualized contexts by taking snapshots. *Id.*, 940-42. In other words, for platforms in which snapshots capture machine-memory state in a page file (when the VM is actively running or suspended), taking a snapshot preserves this page file, which is then readily analyzed to determine things like which applications are being used. *Id.*

62. With that context in mind, a POSA would not have viewed Huseinović as counseling against the combination. First, as I discussed above, snapshots taken in accordance with option (i) of element 1.5 do not require the VM to be “at rest,” so taking the snapshot while the VM is powered off is not required. *Supra*, §V.B; EX1002, ¶¶221-22. With respect to the alternative mapping addressing snapshots that are “request[ed]” while the VM is at rest per option (ii) of element 1.5 (EX1002, ¶¶223-25), I still disagree that Huseinović counsels against the approach I discussed. My declaration explained specific circumstances in which a snapshot of a powered-off VM is requested (satisfying option (ii) of element 1.5) and the POSA nevertheless had reason to power up the VM in order to take the snapshot (e.g., so the page file can be included) before returning the VM to the powered-off state. EX1002, ¶¶223-25. Dr. Kaeli’s assertion does not

address that reasoning either, and Huseinović's disclosure does not counsel against that approach. To the contrary, that approach is precisely what Huseinović suggests doing in some circumstances.

IX. Secondary Considerations

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

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

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

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

66. As I explain in the following sections, it is my opinion that product/platform (which Dr. Kaeli refers to as SideScanning™) is not entitled to the presumption of nexus. *Infra*, §IX.A. Even if the product were given this presumption, it would still be my opinion that the evidence lacks nexus because it is merely directed to known features (some of which are also unrecited). Since much the evidence relates to the same handful of features, I provide a summary of those known features upfront (*infra*, §IX.B.1) so that I can avoid unnecessary

repetition when I address the evidence for each type of secondary consideration discussed in Dr. Kaeli's declaration (*infra*, §§IX.B.2-6).

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

A. No Presumption of Nexus

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

69. In my opinion, the platform is not coextensive with the claims. Neither Dr. Kaeli nor Mr. Shua contend that it is, so there does not appear to be any dispute on this point. In the interest of clarity, however, I have reviewed Orca's exhibits describing the product, and in my opinion, the product is not

coextensive with the claims since it contains multiple features that are both significant and unrecited. I discuss several of these below.

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

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

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

EX2067, 1; *see also id.*, 5 (platform “covers 100% of an organization’s assets” and

“automatically assesses the security state of every discovered asset”). Likewise,

the subsequent description identifies this is part of “one of the most valuable features” that the platform provides. *Id.*, 8. In contrast to the platform’s features, the claims identify and query the location of virtual disks of a single VM, meaning they capture only a tiny subset of this cloud estate. Moreover, the claims focus on a VM, whereas Orca emphasizes coverage of many other types of workloads beyond just VMs. *Id.*, 1 (discussing “VMs, 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 vulnerabilities detected by analyzing the snapshot.

71. Numerous aspects of the data plane path are also unrecited. The claims recite analyzing the at least one snapshot “to detect vulnerabilities.” Yet Orca states that the platform’s data plane path performs numerous assessment steps beyond just detecting vulnerabilities. In the technical brief, “Vulnerability Scanning” is one of seven listed steps:

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

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

EX2067, 10. [REDACTED]

[REDACTED]

[REDACTED]

72. Orca also emphasizes the significance of performing its scans while some assets are active, whereas the claims recite that “the virtual machine is inactive” during the detection of the vulnerabilities. For example, the technical brief emphasizes at the outset that the platform covers both “running” and stopped workloads without disrupting business operations in live environments. EX2067, 1. Orca similarly emphasizes that its platform has “no impact whatsoever on the running cloud environment.” *Id.*, 19.

73. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

74. 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, ¶182. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

B. Orca's Evidence

1. Known Techniques/Features

75. 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

76. 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

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

were also a fundamental aspect of cloud systems, including cloud-security platforms.

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

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

⁶ EX1122 is a correct copy of documentation obtained from https://www.e-spincorp.com/pdf/product/Rapid7/Nexpose_User_Guide.pdf and includes a 2015 date. EX1123 is a correct copy of documentation obtained from https://static.tenable.com/documentation/nessus_aws_installation_and_configuration.pdf and includes a 2014 date. EX1124 is a correct copy of documentation obtained from <https://qualys.my.site.com/discussions/s/article/000005982> and

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

79. Orca contrasts its secondary-considerations evidence with Veselov for limitations 1.2-1.4—arguing that Veselov “does not support these functionalities” (POR, 30-31)—but that is not how a POSA would have understood Veselov. Rather, a POSA would have viewed Veselov encompassing these standard resource-discovery and snapshot-creation mechanisms. 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

includes an April 2018 date. Based on my experience, the content of these documents is consistent with the state of the art as of the listed dates.

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. In the AWS context, volumeId is a required parameter in the typical mechanism for creating a snapshot (EX1121), so a POSA would have expected Veselov's data-identification process in AWS to include discovering the mandatory volumeId at a minimum. 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

80. 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

81. 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), 121, 140 (application usage); EX1085 (Hovor), ¶¶48 (considering “location of the asset in the network topology”), 75-77 (considering scores of connected assets).

82. That contextually prioritized, actionable alerts were known for agentless workload scanning is not surprising since the underlying data providing such context were 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

83. 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, 1-7.⁷ It “can accurately detect all of the confidential data in your organization—whether it’s at rest, in motion, or in

⁷ 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.

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.

84. This example is not an outlier. Sensitive-data detection was known in a variety of platforms such as Forcepoint and McAfee Skyhigh Security Cloud. EX1135, 5 (“Forcepoint Data Discovery can be used to learn the location of sensitive data both on-premises and in supported cloud-based applications. It can be used to scan data on file servers, email servers, and databases.”); EX1136 (McAfee), 1 (“Gain complete visibility into data, context, and user behavior across all cloud services, users, and devices.”), 2 (“Applies unified policies to all cloud services across data at rest and in transit.”), 3 (“Leverages keywords, pre-defined alphanumeric patterns, regular expressions, file metadata, document fingerprints, and database fingerprints to identify sensitive data.”).⁸

⁸ EX1135 and EX1136 are correct copies of product documentation obtained from

85. Sensitive-data detection was also known in the context of agentless workload scanning, including snapshot scanning. For example, Chari describes identifying sensitive data in various resources in order “label” it, assign it a sensitivity level, and sanitize it. EX1056, ¶¶38-42. Chari specific 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.

86. 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.

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.

2. Long-Felt Need

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

88. 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, especially for machines that were at rest or inactive. Kaeli-Dec, ¶¶183-89; *see also* Shua-Dec, ¶¶79-87 (citing EX1001, 1:57-2:28; 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, ¶¶184-85; Shua-Dec, ¶80. 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, ¶¶190-91 (citing EX2076, 7, 11, 31; EX2138, 1; EX2135, 1, 3); *see also* Shua-Dec, ¶¶. As I explained above, however, this argument overlooks that

agentless workload scanning (including such scanning for powered-off assets) was already known. *Supra*, ¶IX.B.1; EX1002, ¶¶56-58, 62. Indeed, Veselov itself explains that its snapshot-based scanning—despite the analysis being performed separately from the target asset without using agents installed in on the target (EX1002, ¶¶160-64), a process that does not depend on whether the asset is actively running, suspended, or powered off (*id.*, ¶¶33, 165-67)—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.

89. Dr. Kaeli and Mr. Shua mention other aspects like API usage and targeted vulnerability alerts, but these were also well known, as I explained above. *Supra*, §IX.B.1; Kaeli-Dec, ¶¶190-92 (citing EX2135, 3; EX2076, 31 [REDACTED]); Shua-Dec, ¶¶87-88 (citing EX2076, 7; EX2103, 1). They were also known as in combination with agentless workload scanning, which is not surprising since cloud APIs are the fundamental communication mechanism of cloud products (including cloud-security platforms), and targeted/prioritized vulnerability alerts were commonplace. *Supra*, §IX.B.1; EX1002, ¶¶42-47 (discussing well-known

cloud computing and cloud APIs), 64-70 (discussing well-known alerting techniques).

90. 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, ¶84. Mr. Shua also cites a discussion of Qualys’s product as describing limitations of agents (*id.*, ¶85), 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),⁹ the author states that “[c]ompanies that once championed agentless solutions are now all building agents.” EX1125, 3; EX2076, 2.

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

3. Skepticism

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

93. First, Dr. Kaeli and Mr. Shua focus primarily on supposed skepticism of agentless workload scanning. Kaeli-Dec, ¶¶194-97 (citing EX2076, 7-8; EX2103; EX2106, 1; EX2135, 2; [REDACTED]); Shua-Dec, ¶¶89-91. As I previously discussed, however, agentless workload scanning (including cloud-API-based agentless workload scanning) was known. *Supra*, §IX.B.1.

94. 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

⁹ EX1125 is a correct copy of a webpage obtained from

<https://softwareanalyst.substack.com/p/redefining-cnapp-a-complete-guide>.

being truly indicative of industry skepticism. Kaeli-Dec, ¶195; 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]

95. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

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

4. Praise

97. 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, ¶¶198-202; Shua-Dec, ¶¶124-141.

98. 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.”).

99. Mr. Shua touts other alleged praise that was made by Orca’s own partners and investors. Shua-Dec, ¶¶133, 138-40; 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.¹⁰ It indicates that the TAG Security Report submitted as EX2118 is part of Orca’s marketing material. The other relationships are also indicated by the submitted documents. For example, Mr. Shua points to Orca being named to Redpoint’s “Infra100 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,

¹⁰ I understand that EX1115 was obtained from <https://orca.security/lp/tag-cyber-cloud-security-vendor-landscape-tco-analysis/>. I have compared it with the webpage at that URL and confirm that it is an accurate representation.

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.”).

100. However, even if I did not consider the connection to Orca’s investors, these documents still would not favor nonobviousness in my opinion because they are not tied to recited and novel aspects of the claims. As I discussed above, in my opinion, Dr. Kaeli and Mr. Shua have not shown that the product embodies and is coextensive with the claims, so general company/product praise would not have a presumption of nexus. *Supra*, §IX.A. Regardless, to the extent particular aspects of Orca and its product are discussed in the cited exhibits at all, they were known, as I discussed above. *Supra*, §IX.B.1. I see no other evidence showing that the above awards/lists are due to the claims. Dr. Kaeli does not address the Redpoint, Notable Capital, or YL Ventures lists, and Mr. Shua does not provide specifics. Shua-Dec, ¶¶133, 138-41.

101. 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, ¶198

(citing EX2119, 9-10; EX2120, 6-7; EX2123); Shua-Dec, ¶¶130-38 (citing EX2119-2124; EX2127-2128; 2143-2144). Dr. Kaeli says that he reviewed EX2119, EX2120, and EX2123, and he quotes a few passages, but he does not explain why he thinks this is praise of the claims. Regardless, the passages he quotes merely describe the unrecited and/or known features I previously discussed. Automatically assessing “every discovered asset” is not recited. EX2119, 9-10; *supra*, §IX.A. Deep visibility and accessing cloud configurations to assess workloads’ block-storage out-of-band was known. EX2120, 6-7; *supra*, §IX.B.1. The “agentless approach” was known. EX2123; *supra*, §IX.B.1. Mr. Shua cites additional exhibits but similarly gives no indication that they relate to recited and novel features of the claims. For example, his quote of the Forbes 2021 exhibit addresses the agentless aspect of Orca’s approach, which was known, and he gives no meaningful specifics for the Forbes 2023 exhibits. Shua-Dec, ¶¶131, 138; EX2120; EX2127-2128. He gives no meaningful specifics for the SINET16 or Gartner “Cool Vendor” listings. Shua-Dec, ¶¶132, 134; 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, ¶137; EX2123; EX2125; *supra*, §IX.B.1.

102. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

103. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

104. Lastly, with respect to the additional cited evidence cited by Dr. Kaeli (Kaeli-Dec, ¶201), in my opinion it does not show praise tied to claimed and novel features. Dr. Kaeli quotes the author's discussion of agentless workload scanning enabled by cloud APIs and his description of Orca's platform as gathering block-level data via the cloud provider's APIs and performing a read-only scan (EX2067, 7, 11-12, 31), but this approach was known, as Veselov itself shows. *Supra*, §IX.B.1. Dr. Kaeli also quotes the Orca partner and the report that I understand was commissioned by Orca (as I discussed above). EX2117, 1; EX2118, 4-5. Those are not objective assessments for the reasons I discussed above, and even if they were, the quoted statements address known features of agentless scanning such as at-rest protection, out-of-band scanning of block-level data, and reporting actionable results. *Supra*, §IX.B.1.

5. Copying

105. 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.

106. 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

107. [REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]

[REDACTED]

108. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

¹¹ [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

109. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

110. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

111. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

112. 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.

EX1137.¹² 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

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

[REDACTED]

[REDACTED]

[REDACTED] However, Mr. Shua indicates that he had already heard about Wiz’s agentless workload scanner in “early 2020.” Shua-Dec, ¶113. [REDACTED]

¹² 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]

[REDACTED]

[REDACTED]

114. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

115. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

116. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

117. I also disagree with Dr. Kaeli's argument that Orca and Wiz's products are substantially similar. Kaeli-Dec, ¶¶228-60. He argues that Wiz practices the claims, but for elements 1.2-1.4, he simply cites a general description of workload scanning and then points to [REDACTED] without any explanation of how this shows practicing any particular element. Kaeli-Dec, ¶¶245-48. [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, ¶231 (citing EX2026; EX2115).

118. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] Orca's own marketing material distinguishes the products. EX1126.¹³ [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

¹³ EX1126 is a correct copy of a webpage obtained from

<https://orca.security/resources/comparisons/wiz/>.

[REDACTED]

[REDACTED]

[REDACTED]

119. 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).

120.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

6. Commercial Success

121. 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.

122. 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, ¶261. That financing/revenue is related to the product generally does not itself indicate nexus to the claims since, as I previously explained, the requirements for the presumption of nexus have been shown. *Supra*, §IX.B.1. In any event, the few aspects of the product Mr. Shua discusses are the same ones that I previously explained were known and/or unrecited. *Supra*, §§IX.A, IX.B.1; Shua-Dec, ¶¶96-107.

123. 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, ¶¶96-107. 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, ¶¶98, 102; EX2109, 1; EX2111, 1. These statements does not indicate that the financing was obtained due to the claims, and regardless, API-based agentless workload scanning was known. *Supra*, §IX.B.1. Mr. Shua points to a quote from a "GGV managing partner" (EX2109, 3), but that individual does not appear to be a disinterested third-party since, as the document states, GGV is an Orca investor. *Id.*, 1

(financing round “led by GGV Capital”). The rest of EX2111 largely just quotes or paraphrases Mr. Shua. EX2111, 2-4. Mr. Shua cites EX2112, which does not connect the funding to the claims, and he quotes the general partner of CapitalG, who also does not appear to be a disinterested third party because, as the document states, CapitalG led this funding round. Shua-Dec, ¶¶103-04; 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, ¶107; EX2113, 1.

124. 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, ¶262. Moreover, Orca’s evidence explicitly attributes Wiz’s success to other factors that are not recited. For example, the 2024 retrospective places particular emphasis on Wiz’s graph, explaining that a “core factor in Wiz’s success is its graph-based architecture, from the actual database to the user experience.” EX2076, 21-23; *see also id.*, 23 (attack-path analysis), 24 (data-clustering techniques informed by attacker research), 25 (Wiz code). The article also emphasizes Wiz’s responsiveness to customer needs:

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

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

X. Concluding Statements

125. 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.

126. 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

Declaration of Angelos Stavrou

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: July 15, 2025

By: /ANGELOS STAVROU/

Angelos Stavrou, Ph.D.