

Filed on behalf of: Wiz, Inc.

By: Matthew A. Argenti (margenti@wsgr.com)
Michael T. Rosato (mrosato@wsgr.com)
Wesley E. Derryberry (wderryberry@wsgr.com)
Tasha M. Thomas (tthomas@wsgr.com)
Joseph M. Baillargeon (jbaillargeon@wsgr.com)
WILSON SONSINI GOODRICH & ROSATI
650 Page Mill Road
Palo Alto, CA 94304

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

WIZ, INC.,
Petitioner,

v.

ORCA SECURITY LTD.,
Patent Owner.

Case IPR2025-00441
Patent No. 11,868,798

**PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 11,868,798**

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	MANDATORY NOTICES UNDER 37 C.F.R. §42.8.....	1
III.	CERTIFICATIONS.....	3
IV.	IDENTIFICATION OF CHALLENGE; STATEMENT OF PRECISE RELIEF REQUESTED	3
V.	THE '798 PATENT.....	4
	A. Prosecution History	5
VI.	NO BASIS EXISTS FOR DENIAL UNDER 35 U.S.C. §325(D)	5
VII.	LEVEL OF ORDINARY SKILL	7
VIII.	CLAIM CONSTRUCTION	7
	A. “[Analyzing/Analyze] the Snapshot”	7
IX.	BACKGROUND	9
	A. Cloud Computing, Virtualization, and Snapshots.....	9
	B. Cyber Security	11
X.	PRIOR ART.....	13
	A. Veselov (U.S. Patent. No. 11,216,563; EX1007)	13
	B. Basavapatna (U.S. Pub. No. 2013/0191919, EX1008)	16
	C. Kapoor (US Patent No. 10,498,845, EX1111).....	17
	D. Czarny (U.S. Patent No. 9,749,349; EX1084)	18
	E. Chari (U.S. Patent Pub. 2015/0033221; EX1056)	19
	F. Roth (U.S. Patent No. 9,524,389; EX1116).....	19
XI.	GROUND 1: CLAIMS 1-5, 7-15, 17-22, 24-32, AND 34 WERE OBVIOUS OVER VESELOV AND BASAVAPATNA.....	20
	A. Reasons to Combine Veselov and Basavapatna.....	21
	B. Independent Claims 1 and 18.....	23
	1. Preambles	23
	2. Element 18.i	24

3.	Elements 1.1 and 18.1	25
4.	Elements 1.2 and 18.2	27
5.	Elements 1.3 and 18.3	28
6.	Elements 1.4 and 18.4	32
C.	Dependent Claims	34
1.	Claims 2 and 19	34
2.	Claims 3 and 20	36
3.	Claims 4 and 21	41
4.	Claims 5 and 22	44
5.	Claims 7 and 24	48
6.	Claims 8 and 25	51
7.	Claims 9 and 26	51
8.	Claims 10 and 27	52
9.	Claims 11 and 28	53
10.	Claims 12 and 29	53
11.	Claim 13 and 30	54
12.	Claims 14 and 31	55
13.	Claims 15 and 32	56
14.	Claims 17 and 34	57
XII.	GROUND 2: CLAIMS 6 AND 23 WERE OBVIOUS OVER VESELOV, BASAVAPATNA, AND KAPOOR	58
A.	Reasons to Combine Veselov, Basavapatna, and Kapoor.	58
B.	Claims 6 and 23	61
XIII.	GROUND 3: CLAIMS 11 AND 28 WERE OBVIOUS OVER VESELOV, BASAVAPATNA, AND ROTH	63
A.	Reasons to Combine Veselov, Basavapatna, and Roth	63
B.	Claims 11 and 28	65
XIV.	GROUND 4: CLAIMS 14 AND 31 WERE OBVIOUS OVER VESELOV, BASAVAPATNA, AND CZARNY	66

A.	Reasons to Combine Veselov, Basavapatna, and Czarny	66
B.	Claims 14 and 31	68
XV.	GROUND 5: CLAIM 13, 16, 30, AND 33 WERE OBVIOUS OVER VESELOV, BASAVAPATNA, AND CHARI	69
A.	Reasons to Combine Veselov, Basavapatna, and Chari	69
B.	Claims 13 and 30	71
C.	Claims 16 and 33	73
XVI.	CONCLUSION	74

LISTING OF CHALLENGED CLAIMS

1. A method for securing virtual cloud assets in a cloud computing environment against cyber threats, comprising:

[1.1] taking a snapshot of at least one virtual disk of a protected virtual cloud asset, wherein the virtual cloud asset is instantiated in the cloud computing environment;

[1.2] accessing the snapshot of the at least one virtual disk;

[1.3] analyzing the snapshot of the protected virtual cloud asset to detect potential cyber threats risking the protected virtual cloud asset;
and

[1.4] alerting detected potential cyber threats based on filtering or prioritizing using a determined risk of each detected threat.

2. The method of claim 1, further comprising: prioritizing each of the detected potential cyber threats based on their respective risk to the protected virtual cloud asset.

3. The method of claim 1, wherein analyzing the snapshot of the protected virtual cloud asset further comprises:

[3.1] parsing a copy of the snapshot; and

[3.2] scanning the parsed copy to detect the potential cyber threats, wherein the potential cyber threats include known and unknown vulnerabilities, and wherein the detection is based on a type of vulnerability.

4. The method of claim 3, wherein scanning the parsed copy further comprises any one of:

[4.1] checking configuration files of applications and operating system installed in the protected virtual cloud asset;

[4.2] verifying access times to files by the operating system installed in the operating system;

[4.3] analyzing system logs to deduce what applications and modules executed in the protected virtual cloud asset; or

[4.4] analyzing machine memory stored in the snapshot to deduce what applications and modules executed in the protected virtual cloud asset.

5. The method of claim 1, further comprising:

[5.1] instantiating an instance copy of the protected virtual cloud asset from the snapshot; and

[5.2] monitoring all activity performed by the instance copy of the protected virtual cloud asset.

6. The method of claim 3, wherein scanning the parsed copy further comprises any one of: reading process identification number (PID) files; and checking if the at least access times of the PID files match against process descriptors.

7. The method of claim 3, wherein scanning the parsed copy further comprises:

[7.1] matching at least one installed application list against a list of vulnerable applications;

[7.2] matching at least one application file against a database of files of vulnerable applications;

[7.3] matching sub-modules of an application against a database of sub-modules; or

[7.4] reading installation logs of package managers.

8. The method of claim 1, wherein the protected virtual cloud asset includes any one of: a virtual machine, a software container, or a micro-service.

9. The method of claim 1, wherein taking the snapshot is performed at a predefined schedule.

10. The method of claim 1, wherein taking the snapshot is performed based on receiving a request to take the snapshot.

11. The method of claim 1, wherein taking the snapshot is performed based on detection of a network event or an abnormal event.

12. The method of claim 1, further comprising mitigating a potential cyber threat posing a risk to the protected virtual cloud asset.

13. The method of claim 1, wherein analyzing the snapshot to detect potential cyber threats includes detecting cyber threats that do not represent vulnerabilities.

14. The method of claim 1, wherein analyzing the snapshot to detect potential cyber threat includes detecting a vulnerable application file by comparing a hash value of the vulnerable application file against a database of known application files.

15. The method of claim 1, wherein analyzing the snapshot to detect potential cyber threats includes looking for cyber threats inside containers stored on the virtual disk.

16. The method of claim 1, wherein analyzing the snapshot to detect potential cyber threat includes looking for sensitive data including at least one of personal identifiable information or private keys.

17. The method of claim 1, wherein the virtual cloud asset is a server or a storage system in a cloud computing platform.

18. A system for securing virtual cloud assets in a cloud computing environment against cyber threats, comprising:

[18.i] at least one hardware processor; and a storage device comprising computer readable instructions, to cause the at least one hardware processor to:

[18.1] take a snapshot of at least one virtual disk of a protected virtual cloud asset, wherein the virtual cloud asset is instantiated in the cloud computing environment;

[18.2] access the snapshot of the at least one virtual disk;

[18.3] analyze the snapshot of the protected virtual cloud asset to detect potential cyber threats risking the protected virtual cloud asset; and

[18.4] alert detected potential cyber threats based on filtering or prioritizing using a determined risk of each detected threat.

19. The system of claim 18, wherein the instructions further cause the at least one hardware processor to: prioritize each of the detected potential cyber threats based on their respective risk to the protected virtual cloud asset.

20. The system of claim 18, wherein analyzing the snapshot of the protected virtual cloud asset further comprises:

[20.1] parsing a copy of the snapshot; and

[20.2] scanning the parsed copy to detect the potential cyber threats, wherein the potential cyber threats include known and unknown vulnerabilities, and wherein the detection is based on a type of vulnerability.

21. The system of claim 20, wherein scanning the parsed copy further comprises any one of:

[21.1] checking configuration files of applications and operating system installed in the protected virtual cloud asset;

[21.2] verifying access times to files by the operating system installed in the operating system;

[21.3] analyzing system logs to deduce what applications and modules executed in the protected virtual cloud asset; or

[21.4] analyzing machine memory stored in the snapshot to deduce what applications and modules executed in the protected virtual cloud asset.

22. The system of claim 18, wherein the instructions further cause the at least one hardware processor to:

[22.1] instantiating an instance copy of the protected virtual cloud asset from the snapshot; and

[22.2] monitoring all activity performed by the instance copy of the protected virtual cloud asset.

23. The system of claim 20, wherein scanning the parsed copy further comprises any one of: reading process identification number (PID) files; and checking if the at least access times of the PID files match against process descriptors.

24. The system of claim 20, wherein scanning the parsed copy further comprises:

[24.1] matching at least one installed application list against a list of vulnerable applications;

[24.2] matching at least one application file against a database of files of vulnerable applications;

[24.3] matching sub-modules of an application against a database of sub-modules; or

[24.4] reading installation logs of package managers.

25. The system of claim 18, wherein the protected virtual cloud asset includes any one of: a virtual machine, a software container, or a micro-service.

26. The system of claim 18, wherein taking the snapshot is performed at a predefined schedule.

27. The system of claim 18, wherein taking the snapshot is performed based on receiving a request to take the snapshot.

28. The system of claim 18, wherein taking the snapshot is performed based on detection of a network event or an abnormal event.

29. The system of claim 18, further comprising mitigating a potential cyber threat posing a risk to the protected virtual cloud asset.

30. The system of claim 18, wherein analyzing the snapshot to detect potential cyber threats includes detecting cyber threats that do not represent vulnerabilities.

31. The system of claim 18, wherein analyzing the snapshot to detect potential cyber threat includes detecting a vulnerable application file by comparing a hash value of the vulnerable application file against a database of known application files.

32. The system of claim 18, wherein analyzing the snapshot to detect potential cyber threats includes looking for cyber threats inside containers stored on the virtual disk.

33. The system of claim 18, wherein analyzing the snapshot to detect potential cyber threat includes looking for sensitive data including at least one of personal identifiable information or private keys.

34. The system of claim 18, wherein the virtual cloud asset is a server or a storage system in a cloud computing platform.

I. INTRODUCTION

Petitioner Wiz, Inc. (“Wiz”) respectfully requests review of U.S. Patent No. 11,868,798 (“the ’798 patent”), currently assigned to Orca Security Ltd. (“Orca”). This petition demonstrates claims 1-34 are unpatentable.

The ’798 claims describe well-known techniques for securing virtual assets such as virtual machines (“VMs”) in a cloud computing environment. A “snapshot” of the asset’s virtual disk is taken, accessed, and analyzed to detect potential cyber threats. Alerts are then sent based on filtering or prioritizing using a determined risk of each threat.

This type of snapshot-based analysis was already well known, as demonstrated by the combination of Veselov and Basavapatna. Veselov discloses most aspects of the independent claims, though it does not expressly discuss risk-based filtering or prioritizing. However, these techniques were well known, as shown by Basavapatna. The dependent claims describe other well-known features.

Accordingly, Wiz respectfully requests institution.

II. MANDATORY NOTICES UNDER 37 C.F.R. §42.8

Real Party-in-Interest (37 C.F.R. §42.8(b)(1)): Petitioner Wiz is the real party-in-interest.

Related Matters (37 C.F.R. § 42.8(b)(2)): Wiz is involved in litigation involving patents related to the ’798 patent in *Orca Security Ltd. v. Wiz, Inc.*, No.

1-23-cv-00758 (DDE), filed and served on July 12, 2023. Wiz previously filed an IPR petition in IPR2024-00220 against U.S. Patent No. 11,431,735 (“the ’735 patent”), which is a related patent owned by Patent Owner that contain claims similar to those of the ’798 patent. IPR2024-00220, Paper 2. Like the current petition, the petition in IPR2024-00220 included a ground based on the combination of Veselov and Basavapatna. In response, Patent Owner disclaimed all challenged claims. IPR2024-00220, Paper 6. Wiz has also filed six petitions against patents that are involved in the abovementioned litigation: IPR2024-00863 against U.S. Patent No. 11,663,031, IPR2024-00864 against U.S. Patent No. 11,663,032, IPR2024-00865 against U.S. Patent No. 11,693,685, IPR2024-01109 against U.S. Patent No. 11,726,809, IPR2024-01190 against U.S. Patent No. 11,740,926, and IPR2024-01191 against U.S. Patent No. 11,775,326. The Board ordered institution in the first three of these cases on December 9, 2024, and in IPR2024-01109 on January 16, 2025. Institution decisions for the last two are expected shortly. Wiz also recently filed an IPR petition in IPR2025-00092 against U.S. Patent No. 11,516,231, which is related to the ’798 patent.

Lead and Back-Up Counsel (37 C.F.R. §42.8(b)(3)):

Lead Counsel: Matthew A. Argenti (Reg. No. 61,836)

Back-Up Counsel: Michael T. Rosato (Reg. No. 52,182); Wesley E. Derryberry (Reg. No. 71,594); Tasha M. Thomas (Reg. No. 73,207); Joseph M. Baillargeon (Reg. No. 79,685).

Service Information—37 C.F.R. §42.8(b)(4): Wiz consents to electronic service. Please direct all correspondence to lead and back-up counsel at the contact information below. A power of attorney accompanies this petition.

E-mail: margenti@wsgr.com; mrosato@wsgr.com; wderryberry@wsgr.com; tthomas@wsgr.com; jbaillargeon@wsgr.com

Post: WILSON SONSINI GOODRICH & ROSATI, 650 Page Mill Road,
Palo Alto, CA 94304

Tel.: 650-354-4154 Fax: 650-493-6811

III. CERTIFICATIONS

The '798 patent is available for IPR, and Wiz is not barred or estopped from requesting IPR on these grounds.

IV. IDENTIFICATION OF CHALLENGE; STATEMENT OF PRECISE RELIEF REQUESTED

Wiz seeks cancellation of the challenged claims for the reasons stated below, which are supported with exhibits, including the Declaration of Dr. Angelos Stavrou (EX1002). The claims are unpatentable under 35 U.S.C. §311 and AIA §6 based on at least the following grounds:

Ground	Claims	Basis
1	1-5, 7-15, 17-22, 24-32, 34	§103(a): obviousness over Veselov and Basavapatna.
2	6, 23	§103(a): obviousness over Veselov, Basavapatna, and Kapoor.
3	11, 28	§103(a): obviousness over Veselov, Basavapatna, and Roth.
4	14, 31	§103(a): obviousness over Veselov, Basavapatna, and Czarny.
5	13, 16, 30, 33	§103(a): obviousness over Veselov, Basavapatna, and Chari.

V. THE '798 PATENT

The '798 patent issued from U.S. Application No. 17/821,345 (“the '345 Application”), filed August 22, 2022. EX1001, Face. The '345 application claims priority to Provisional Application No. 62/797,718, filed January 28, 2019. The '798 patent thus has an effective filing date no earlier than January 28, 2019, and is subject to AIA §102 and §103. *Id.*; EX1002, ¶20.

The '798 patent describes securing virtual assets in a cloud environment. EX1001, Abstract. The specification describes well-known snapshot-based analysis that includes taking a snapshot of a virtual disk(s), accessing/analyzing the snapshots to identify cyber threats, and alerting based on filtering or prioritizing. *Id.*, 5:3-13, 7:13-8:6, Fig. 2; EX1002, ¶¶71-74.

The '798 patent includes 34 claims. Claims 1 and 18 are independent. Claim 18 essentially mirrors the method of claim 1 but is directed to a system. The dependent claims add other conventional aspects of cybersecurity and cloud computing. EX1002, ¶¶75-77.

A. Prosecution History

During prosecution, the '345 application (i.e., the patent application that led to the '798 patent) never received a rejection under §102 or §103. The first office action rejected the claims based on nonstatutory double patenting over claims from the parent application (now the '735 patent) and five other related applications. EX1004, 135-37. The Applicant filed six terminal disclaimers to secure allowance over the related patents, including the '735 patent which has now had all but two dependent claims disclaimed. *Id.*, 104-06, 110-17, 122-23; *Wiz, Inc. v. Orca Security Ltd.*, IPR2024-00220, Paper 7 (May. 09, 2024), 3. The notice of allowance identified three references as the closest prior art but only broadly indicated the supposedly lacking teachings by listing most of the independent claim language without explanation. EX1004, 59-60; EX1002, ¶78.

VI. NO BASIS EXISTS FOR DENIAL UNDER 35 U.S.C. §325(d)

Under the two-part *Advanced Bionics* framework, §325(d) analysis considers several factors to determine:

- (1) whether the same or substantially the same art previously was presented to the Office or whether the same or substantially the same

arguments previously were presented to the Office; and (2) if either condition of [the] first part of the framework is satisfied, whether the petitioner has demonstrated that the Office erred in a manner material to the patentability of challenged claims.

Advanced Bionics, LLC v. Med-El Elektromedizinische Geräte GmbH, IPR2019-01469, Paper 6 at 8 (precedential); 35 U.S.C. §325(d). While Veselov was disclosed as one of many references across multiple information disclosure statements, it was never applied in a rejection or substantively discussed. EX1004, 126, 135-37, 170-71, 179-80, 195. Veselov was also never considered in combination with Basavapatna, Kapoor, Chari, Czarny, or Roth, since these references were not disclosed. The Office thus did not consider any of the grounds presented herein. The Office also lacked additional evidence discussed herein, including the declaration provided by Wiz’s expert, Dr. Stavrou.

Allowance of the claims also constituted material error under part two of the *Advanced Bionics* test. The ’345 application never received an art-based rejection, and no particular limitation was identified as a basis for allowance. *Supra*, §V.A. The reasons given for allowance simply list the majority of the claim limitations as supposedly not disclosed by the “closest” art. *See* EX1004, 58-60. By contrast, the present grounds teach all limitations of the challenged claims as a whole. *Infra*, §§XI-XV. The claims therefore should not have issued, and they would not have issued if the Examiner had considered the present grounds.

VII. LEVEL OF ORDINARY SKILL

For purposes of this petition, Wiz assumes a priority date of January 28, 2019. A POSA as of January 2019 would have held at least a bachelor's degree in computer science, computer engineering, electrical engineering, or a related field, and would also have 2-3 years of professional experience working with cyber security analysis and virtualization. Additional experience could compensate for less education and vice versa. Relevant work experience includes, for example, malware analysis, security analysis of cloud computing systems, and security analysis of VMs. EX1002, ¶¶21-22. Dr. Stavrou meets these requirements and is qualified to credibly opine on the state of the art and the POSA's perspective. *Id.*, ¶¶1-19; EX1003 (Stavrou CV).

VIII. CLAIM CONSTRUCTION

Claim terms are given their ordinary and customary meaning, consistent with the specification, as a POSA understood them. 37 CFR §42.100(b); *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312-13 (Fed. Cir. 2005) (en banc). This petition applies the ordinary and customary meaning of the claim terms. EX1002, ¶79.

A. “[Analyzing/Analyze] the Snapshot”

Each independent claim recites analyzing “the snapshot.”

The ordinary and customary meaning of this language encompasses direct analysis of the snapshot data (e.g., analyzing the snapshot as a data file without

instantiating an assessment VM). EX1002, ¶¶80-81. This understanding is confirmed by the specification. *See, e.g.*, EX1001, 5:19-20 (“The snapshot is parsed and analyzed by the security system 140 to detect vulnerabilities.”), 5:36-39 (direct or hash-based matching of application files); *see also id.*, 6:4-11 (analyzing page file), 6:34-37 (security system computes cryptographic hash of sensitive areas in virtual disk and checks for differences), 6:54-58 (analysis of logs “derived from the snapshot”); EX1002, ¶81. Veselov describes this approach. *Infra*, §X.A.

In the related litigation (*supra*, §II), Orca’s infringement allegations treat this limitation as encompassing analysis of a VM instantiated from a snapshot, which Veselov also describes (*infra*, §X.A). For example, Orca alleges that the accused product satisfies “analyzing the at least one snapshot,” as recited in claim 9 of related U.S. Patent No. 11,693,685 because it “‘analyzes [the] operating system, application layer, and data layer’ of virtual machines.” EX1006, 23, 57-58. For purposes of this IPR, Wiz also applies that interpretation. *See also* EX1002, ¶82.

Orca has since argued that this limitation does not encompass analyzing a VM instantiated from the snapshot, though Orca still maintains that the limitation encompasses indirect snapshot analysis. IPR2024-00864, Paper 6, 9-11. Orca has not explained the contradiction between its recent statements and its infringement allegations—and, regardless, Veselov teaches that disclosures presented in its discussion of indirect snapshot analysis (i.e., analysis of a duplicate VM instantiated

from the snapshot) apply to embodiments involving direct snapshot analysis (EX1002, ¶¶91, 133, 174, 178, 185, 213, 237, 256; EX1007, 4:19-29, 8:52-62, 10:37-41, 16:15-18, 16:56-67, 18:16-18, 18:56-59, 24:54-25:2)—so this petition discusses both of Veselov’s approaches. *Infra*, §XI.B.5.


Accordingly, the discussion below applies a construction of “[analyzing/analyze] the snapshot” encompassing both direct analysis of the snapshot data and analysis of a VM instantiated from the snapshot. EX1002, ¶83.

IX. BACKGROUND

A. Cloud Computing, Virtualization, and Snapshots

Cloud computing was well known long before 2019. EX1002, ¶¶23, 40-42; EX1015, 55-58, 62-66, 164-66, 118, 138, Figs. 8-2, 9-1; EX1021, 1, 18-19, 94-95; EX1022, 29. The physical infrastructure was often provided by data centers that included large collections of physical resources. EX1002, ¶44; EX1013, 229; EX1021, 19.

Cloud systems typically used a “virtualization” layer that abstracts the underlying resources to efficiently manage the operation of multiple applications across multiple physical servers. EX1002, ¶¶24, 43; EX1009, xxiii; EX1010, 2; EX1011, 35; EX1021, 19. Each physical server could emulate multiple virtualized computer systems (e.g., VMs), running their own operating system/applications:



EX1009, 505 (Fig. A-5); *see also* EX1002, ¶¶25-27; EX1009, xxiii, 5, 505; EX1010, 2; EX1013, 229. Virtualized resources were commonly referenced via various types of virtual or non-virtual locations, including more general locations (e.g., the resource’s computing environment, storage service, or directory) and more specific locations (e.g., an address or file path). EX1002, ¶¶28-31; EX1009, xxiv, 2, 22, 242, 246-57, 505, 514-15, Fig. A-5; EX1010, 3-4; EX1012, 9:9-25; EX1013, 229; EX1014, 22, Fig. 2.1; EX1015, 56, 124; EX1016, ii; EX1017, 1:16-35; EX1021, 8; EX1031, 1; EX1048, ¶¶21, 31; EX1054, 1:31-42; EX1074, 12; EX1080, 5:34-42.

As early as 2005, virtualized systems employed backup techniques involving “snapshots,” which often saved data from the VM’s memory and disks, including sensitive data and any system/application vulnerabilities, to allow reversion to a previous state. EX1002, ¶¶32-37; EX1009, 257; EX1015, 164; EX1018, 2-6; EX1019, Abstract; EX1020, Abstract, 21:42-22:58; EX1049, 940-41; EX1051, 77, 119, 297; EX1052, 203; EX1069, 18:23-32; EX1064, ¶¶23, 31. Snapshot

generation routinely involved determining a location to store the snapshot files for later access. EX1002, ¶¶38-39; EX1009, 32, 221, 257-60; EX1015, 56, 164-66; EX1071, 6:35-39; EX1072, 4:1-13. Furthermore, snapshot generation routinely involved preliminary steps such as identifying/locating virtual disks that would be part of the snapshot. EX1002, ¶¶45-47; EX1048, ¶¶21, 42; EX1051, 47, 119, 125; EX1052, 445-46; EX1053, ¶¶36, 87-92, Fig. 7; EX1020, 21:9-22:18, Fig. 4; EX1055, 13, 23, 32-33, 53-56, 68-69.

B. Cyber Security

Traditional security systems sought to improve security by identifying security risks including vulnerabilities and potential exposure of sensitive data. EX1002, ¶¶48-49, 51, 55, 58; EX1023, 1-2; EX1024, 1; EX1036, 12:36-45; EX1044, 3:50-54; EX1056, ¶¶38-39, 60-62, 64-67; EX1057, 1-2; EX1058, 1-4; EX1059, 1-3; EX1075, 3:43-67, 5:16-58; EX1079, 31-36. Security scans commonly searched for known risks—e.g., those in published lists such as the Common Vulnerabilities and Exposures (“CVEs”)—and unknown risks that might be indicated, for example, by behavioral abnormalities. EX1002, ¶¶50, 59-60; EX1023, 1-3; EX1025, 12, 21, 29; EX1026, 9, 24-25; EX1044, 5:28-32; EX1047, 1-5, 7; EX1068, 14:58-15:18; EX1076, 2:24-67; EX1077, 2:45-51, 8:1-30; EX1084, 56-4:24, 11:27-47. VMs were known to be subject to at least the same security risks as nonvirtualized machines. EX1002, ¶¶52-53; EX1018, 3-1, ES-2,

4-1, 4-3; EX1021, 45-46; EX1027, 8, 24-26; EX1028, 6-7, 171-72; EX1029, 1, 321.

A POSA would have been familiar with a variety of scanning techniques. Both agent-based techniques—in which a security application, or agent, scans the same computer in which it has been installed—and agentless techniques were well-known and commonly used. EX1002, ¶¶54, 56-57; EX1030, 116; EX1031, 22; EX1032, ¶¶27, 43; EX1033, 1-2; EX1034, 5643; EX1037, 10:5; EX1038, 133-34; EX1040, Abstract (passive monitoring of virtual server “without utilizing agents executing within the virtual server”), ¶43; EX1095, ¶¶5-6, 43. For example, virtual machine introspection (“VMI”) is an agentless technique that was commonly used in virtualized environments and known to have several advantages (e.g., increasing efficiency and keeping the security software isolated from the potentially compromised guest VM). EX1002, ¶57; EX1018, 3-3; EX1033, 1; EX1034, 5643-44; EX1035, 389; EX1036, 3:56-65; EX1037, 10:9-10; EX1038, 133-34.

For virtualized systems, security scans were often performed on a snapshot of the protected resource rather than directly on the resource itself. EX1002, ¶57; EX1014, 25, 55; EX1038, 134-35; EX1039, 9; EX1040, Abstract, ¶43. Another common technique involved “parsing” snapshots or other data (e.g., reformatting the data) before scanning it. EX1002, ¶¶58, 62-63; EX1007, 15:60-16:3; EX1035,

389; EX1039, 14; EX1041, 3:1-9; EX1042, 10:45-55, 12:54-60; EX1043, 12:10; EX1048, Abstract, ¶¶14, 29, 32, 42, 45, Figs. 3A-C, 4, 6, claim 17; EX1065, 2:29-33; EX1066, 2; EX1067, 390; EX1073, 1; EX1081, ¶79.

Security systems typically alerted users of the identified risks. EX1002, ¶64; EX1025, 21. Given the number of risks typically identified and their varying importance, alerts were commonly prioritized. EX1002, ¶¶61, 65-68; EX1008, Abstract; EX1025, 32; EX1042, 2:10-27; EX1044, 6:9-18, 11:51-67; EX1045, ¶7; EX1046, 11:52-12:14; EX1048, ¶43, Fig. 5; EX1063, ¶¶80, 100-113; EX1075, 1:5-22, 1:23-42, 3:9-42, 4:24-5:15, 5:16-58, 5:59-9:26, Figs. 2-4; EX1085, ¶¶48, 75-77. Cybersecurity systems also routinely provided remedial actions to mitigate identified risks (e.g., software patches or asset quarantining). EX1002, ¶¶69-70; EX1025, 21, 24; EX1032, ¶33; EX1046, 12:15-53; EX1095, ¶¶54-55; EX1067, 390; EX1070, Abstract, ¶53, claim 2; EX1048, ¶14, claim 17; EX1081, ¶79.


X. PRIOR ART

A. Veselov (U.S. Patent. No. 11,216,563; EX1007)

Veselov was filed May 19, 2017, and is therefore prior art under 35 U.S.C. §102(a)(2).

Veselov describes scanning techniques mirroring those of the '798 patent. EX1007, Abstract; EX1002, ¶84. A scanning service generates/requests a snapshot (or obtains/accesses an existing snapshot), accesses/analyzes the snapshot, and

reports assessment results. EX1007, Abstract, 2:35-41, 3:20-4:18, 5:9-22, 8:16-51, 9:9-10:36, 11:5-33, 16:47-53; EX1002, ¶85. An exemplary process is depicted below:



EX1007, Fig. 2; *see also id.*, 8:52-10:36 (describing Fig. 2), 16:15-17:9 (describing Fig. 4), 18:16-56 (describing Fig. 6), Figs. 4, 6. Veselov indicates that teachings

discussed for certain embodiments broadly apply to other embodiments. EX1002, ¶91.

The scanning service can directly analyze the snapshot and/or analyze an assessment VM instantiated from the snapshot. EX1007, Figs. 3A-B (assessment VM), 5A-B (direct analysis of snapshot); EX1002, ¶¶86-87. Both assessments identify security risks. EX1007, 1:41-49 (known security assessment services identify “vulnerabilities, bad configurations, and the like”), 2:35-41 (analysis of CVEs, benchmarks, best practices, configurations, etc.); EX1002, ¶86.

Veselov notes that its techniques are particularly useful in cloud computing environments such as data centers. EX1007, 2:22-3:19; EX1002, ¶87. It also applies its techniques to distributed, virtualized systems commonly used for cloud computing:




FIG. 1

EX1007, Fig. 1; *see also id.*, 1:5-50; *supra*, §IX.A. Scans are performed in the target resource’s computing environment or in a test environment. EX1007, Figs. 3A, 5A (test environment), 3B, 5B (analysis in target’s environment).

Veselov describes both agent-based and agentless scans. Some embodiments use an agent installed on an assessment VM instantiated from the snapshot. *E.g.*, EX1007, 14:22-17 (agent example), 16:45-53 (same), claim 2 (same), Fig. 3B (depicting security agent 330); EX1002, ¶89. Other embodiments do not use an agent at all; for example, some perform agentless analysis of the assessment VM and others analyze the snapshot as a data file. EX1007, 3:61-67, 16:67-17:3, Figs. 3A (depicting assessment VM without agent), 5A-B (agentless analysis of snapshot as data file); *see also id.*, 4:53-54, 6:38-43, 9:17-25, 10:32-36, 14:13-22 (appropriately authorized scanning service “may perform any suitable security assessment” on assessment VM and its underlying resources); EX1002, ¶¶88, 90; *supra*, §IX.B (agentless security assessments well-known).

B. Basavapatna (U.S. Pub. No. 2013/0191919, EX1008)

Basavapatna was filed on January 19, 2012, and published July 25, 2013. It is therefore prior art under §§102(a)(1)-(2).

Basavapatna detects vulnerabilities and associated threats by comparing an asset’s hardware/software configurations with vulnerability/threat definitions. EX1008, ¶¶20-30, 33-49, 50-62, ¶¶70-71; EX1002, ¶92. It then determines a

vulnerability-centric risk metric for each detected vulnerability (or threat-centric risk metrics for threats associated with vulnerabilities). EX1008, ¶¶45, 72-79, 82-85, 92-97, 100, 104-06, 112-15, Figs. 3A-B; EX1002, ¶¶93-95. The process is performed for multiple assets, and the underlying risk metrics are combined to produce an aggregate risk metric for each asset. EX1008, ¶¶120-37, Fig. 5; EX1002, ¶96. Vulnerabilities, threats, and assets can then be filtered and/or ranked in user reports to facilitate prioritization of more critical security issues. EX1008, Abstract, ¶¶22-29, 138-40, Fig. 6; EX1002, ¶97.

C. Kapoor (US Patent No. 10,498,845, EX1111)

Kapoor is a publication of a patent that relates to collecting and reporting “data associated with a process...to detect anomalies in a network environment.” EX1111, Abstract; EX1002, ¶98.

Kapoor teaches a cyber security method to address difficult to detect anomalies in virtualized cloud environments, such as datacenters containing VMs. EX1111, 1:18-24, 4:27-34 (datacenters comprise cloud-based resources), 15:54-67 (discussing the use of VMs), 4:15-21 (modeling datacenter activity to identify anomalies).

Kapoor further teaches collecting and transmitting datacenter node information such as the currently running processes on the node. *Id.*, 6:41-7:67, 8:31-48; EX1002, ¶¶99-100. Kapoor teaches searching the PID file directory and

mapping the PIDs on that node to determine the running processes. EX1111, 7:21-23, *see also id.*, 7:63-67; EX1002, ¶100. Kapoor explains that the operating system on a node uses PIDs to “identify processes running at a given time.” EX1111, 37:35-37. The collected PID information is transmitted to a “platform” for analysis. *Id.*, 8:41-43, 8:64-65, Fig. 3A; EX1002, ¶100.

Kapoor teaches the platform detects anomalies by periodically taking new snapshots to build a model of the datacenter activity. EX1111, 23:37-39, 24:8-15, Fig. 18. This model describes how node processes behave over time and provides information such as “all active processes and PIDs in the last hour.” *Id.*, 24:15-21; EX1002, ¶101.

D. Czarny (U.S. Patent No. 9,749,349; EX1084)

Czarny was filed September 23, 2016, and issued August 29, 2017. It is therefore prior art under §§102(a)(1)-(2).

Czarny describes identifying vulnerabilities by matching the scanned system’s software against a database of known vulnerable applications. EX1084, Abstract, 3:37-4:24; EX1002, ¶¶102-03. Czarny discusses certain conventional approaches that perform this matching using data identifying the software applications (e.g., software name/version). EX1084, 4:25-33. It further describes matching using other types of application data such as “complete binary level files of the software products” or hashes thereof. *Id.*, 4:7-47, 5:15-6:24, 9:17-10:42,

11:16-12:10, Figs. 1-5, 7-8, 10; EX1002, ¶¶104-05. Vulnerability data in the database may be obtained from CVE repositories and other public vulnerability repositories. EX1084, 6:25-38. Results are sent to the user in a report that may include a determined vulnerability level of the asset. *Id.*, 7:31-67, 12:11-18, Figs. 6, 10; EX1002, ¶106

E. Chari (U.S. Patent Pub. 2015/0033221; EX1056)

Chari was filed July 24, 2013, and published January 29, 2015. It is therefore prior art under §102(a)(1)-(2).

Chari describes detecting sensitive data in VM images/snapshots in cloud environments. EX1056, Abstract, ¶¶42, 52-62, 94, Figs. 1-2, 9, 10A-B; EX1002, ¶107. Given the various types of sensitive data at risk in cloud systems, Chari analyzes images to identify, prioritize, and remediate sensitive data. EX1056, ¶¶30, 38-39, 42, 61, 64-66, 76-80, 83-87, Figs. 5-8. This process may be applied to instantiated VMs, disk images, or snapshots. *Id.*, ¶¶61, 92-95, Fig. 9, claim 9; EX1002, ¶107.

F. Roth (U.S. Patent No. 9,524,389; EX1116)

Roth describes event-triggered snapshots used for VM security analysis. EX1116, Title, Abstract. Roth notes common security concerns for VMs such as deviation from standard configurations and installation of unauthorized programs. *Id.*, 1:12-36, 2:47-62, 6:16-48, Figs. 1-4. Roth discusses triggers for taking

snapshots so that VM state can be captured for security analysis. *Id.*, Abstract, 1:36-44, 2:20-3:16, 7:1-21, claim 1; *see also id.*, 3:17-4:36 (snapshots captured for forensic analysis), 4:46-5:15 (snapshot contents). Users can configure various types of snapshot triggers such as time-based events, systems errors, unauthorized configurations, intrusions, and other detected security problems. *Id.*, 2:67-3:16, 3:50-67, 7:59-8:29, 17:35-43; *see also id.*, 6:58-67 (discussing VM termination in response to various events); EX1002, ¶108.

XI. GROUND 1: CLAIMS 1-5, 7-15, 17-22, 24-32, AND 34 WERE OBVIOUS OVER VESELOV AND BASAVAPATNA

The combination of Veselov and Basavapatna teaches, suggests, or makes obvious all limitations of claims 1-5, 7-15, 17-22, 24-32, and 34 as a whole. Veselov describes taking a snapshot of a virtual disk, accessing the snapshot, analyzing the snapshot to detect cyber threats, and reporting the results. EX1002, ¶¶109-10; *supra*, §X.A. Basavapatna further discusses reporting cyber threats based on a determined priority or filtering. EX1002, ¶111; *supra*, §X.B. A POSA had reason to combine Veselov and Basavapatna to, for example, perform snapshot-based security assessments and alerting, as taught by Veselov, while further applying Basavapatna's teachings of risk-based filtering and prioritizing of alerts. EX1002, ¶112; *infra*, §§XI.A-B.

A. Reasons to Combine Veselov and Basavapatna

The references are analogous art. Veselov and Basavapatna both assess virtual assets to identify and report security issues. *Supra*, §§X.A-B; EX1007, Abstract, 4:5-8, 8:46-50; EX1008, Abstract, ¶¶1, 12. Both references are thus in the same field as the '798 patent and are pertinent to issues addressed by the '798 patent. EX1001, 1:12-14; EX1002, ¶¶113-14.

Arriving at the claims involved using known elements according to their known functions to achieve predictable results. *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 416-17 (2007). *Supra*, §§IX.A-B. Locating/accessing/analyzing a snapshot of a virtual disk to identify threats and reporting/mitigating identified security problems, as taught by Veselov, were well-known, predictable techniques. *Supra*, §§IX.A-B, X.A; EX1002, ¶¶57, 64, 69-70, 115 (citing EX1014, 24-25; EX1034, 5643; EX1040, Abstract, ¶43, claim 6; EX1076, 2:24-67; EX1077, 2:45-51, 8:1-30; EX1078, Abstract, ¶¶8, 32-34, 104-05, 149). Determining threat risk, determining application usage, and reporting results based on a determined priority, as taught by Basavapatna, were also well-known, predictable techniques. *Supra*, §§IX.B, X.B; EX1002, ¶¶61, 65-68, 115 (citing EX1025, 32; EX1034, 5643; EX1036, 12:36-45; EX1044, 6:9-7:31, 11:51-67; EX1046, 11:52-12:14; EX1048, ¶43, Fig. 5; EX1078, ¶¶121, 140). The combination uses these techniques to predictably perform their known functions since the references

describe analogous security assessments, and applying Basavapatna's techniques in snapshot-based VM assessments would not meaningfully change their function or result. EX1007, 4:1-4, 9:55-67; EX1002, ¶116.

A POSA would also have been motivated to combine Veselov and Basavapatna for several reasons.

First, a POSA would have been motivated to provide more robust and accurate risk assessments. Basavapatna teaches that determining cyber threats and reporting them based on a determined priority or filtering provides a more comprehensive and accurate risk assessment for the system. EX1008, ¶¶72, 120, 138. It was also well known that VMs in cloud environments typically had the same types of security considerations as non-virtual assets. EX1002, ¶¶117-18. The combination would therefore improve security by assessing and communicating important risks in snapshot-based security assessments of VMs. *Id.*, ¶118.

Second, a POSA would have been motivated to obtain known benefits of prioritizing and filtering. Basavapatna recognizes the well-known fact that assets typically face many security problems of varying importance. EX1008, ¶72; EX1002, ¶119. Basavapatna also teaches well-known advantages of prioritizing more critical threats based on associated risk levels and filtering out the less critical. EX1008, ¶¶120, 138; EX1002, ¶¶61, 65-68, 119 (citing EX1045, ¶¶7-8;

EX1046, 11:52-12:14; EX1063, ¶¶187-89, 207-11, 215-16; EX1078, ¶¶6, 155, 165; 180, 183, 192). Given the similar data and security considerations in virtual and non-virtual systems discussed above, a POSA would have recognized that Basavapatna’s teachings would similarly benefit the type of snapshot-based assessments taught by Veselov. EX1002, ¶¶57, 120 (citing EX1025, 3, 32; EX1044, 6:9-18; EX1045, ¶¶7-8; EX1046, 11:52-12:14).

A POSA would have reasonably expected success in combining the references to arrive at the claims. As discussed above, snapshot-based analysis, threat detection, risk assessments, alerting, and risk-based prioritization/filtering were well known and routinely implemented. EX1002, ¶¶65-68, 121 (citing EX1025, 32; EX1078, ¶¶31-33, 36, 96, 120, 128, 180-83). A POSA would have been familiar with the applied techniques, and the combination would not have presented meaningful technical challenges or produced unexpected results. EX1002, ¶121.

B. Independent Claims 1 and 18

1. Preambles

To the extent the preambles are limiting, Veselov teaches them. Claim 1 recites a “method for securing virtual cloud assets in a cloud computing environment against cyber threats, comprising” and claim 18 recites a “system for securing virtual cloud assets in a cloud computing environment against cyber

threats, comprising.” Consistent with the ’798 patent, Veselov describes methods/systems for securing virtual assets against cyber threats. EX1007, Abstract, 2:22-41, 3:20-4:18, 7:1-5, 10:17-36, Figs. 1-7; EX1002, ¶¶122-24 (citing EX1001, 3:13-16, 4:43-45, 6:38-41, 7:11-14); *supra*, §X.A. The references further teach securing virtual “cloud” assets against cyber threats “in a cloud computing environment.” Veselov’s techniques are particularly useful in a “distributed computing environment (e.g., a **‘cloud’ computing environment**),” such as a data center. EX1007, 2:42-3:19 (emphasis added); *see also id.*, 2:22-41 (techniques particularly suited for assessments in data center), 12:21-36, 19:30-37, 19:60-64, 20:46-56 (environment may be “distributed and/or virtual computing environment”), Figs. 1, 7; EX1002, ¶125; *supra*, §IX.A (data centers were typical components of cloud computing environments). The secured assets are virtual “cloud” assets because they are hosted in a cloud environment. EX1002, ¶¶126-27.

2. Element 18.i

Veselov teaches this element, which recites “at least one hardware processor; and a storage device comprising computer readable instructions, to cause the at least one hardware processor to....” The system’s components, including the scanning service, can include one or more servers, each of which may have a processor. EX1007, 5:23-40; EX1002, ¶¶128-29. Moreover, Veselov’s system also includes computer “memory,” such as hard disks and

random-access memory that store instructions that are executed by the processor. EX1007, 20:38-45, 20:57-23:18, claim 1; EX1002, ¶130. These instructions executed by the processor perform the system’s operations, and those operations—in the context of the proposed combinations in all grounds—include all operations recited in claim 18 and its dependent claims. EX1007, 4:60-5:1; EX1002, ¶130.


3. Elements 1.1 and 18.1

Veselov teaches these elements, which recite “[taking/take] a snapshot of at least one virtual disk of a protected virtual cloud asset, wherein the virtual cloud asset is instantiated in the cloud computing environment.”

Veselov teaches “taking a snapshot of at least one virtual disk of a protected virtual cloud asset.” To perform Veselov’s assessment, the scanning system obtains/accesses “a state of the resource at a point in time (e.g., a ‘snapshot’).” EX1007, 3:20-27 (emphasis added); *see also id.*, 2:42-63, 4:53-54, 6:25-30; EX1002, ¶132. Veselov teaches several ways for the scanning service to obtain/access the snapshot, including taking the snapshot itself. EX1007, 8:62-9:25 (scanning service may “obtain the snapshot data by executing the snapshot itself”), 16:15-26, 18:15-21, Figs. 2 (step 206), 4 (step 404), 6 (step 604); EX1002, ¶132; *see also* EX1007, 3:22-27 (snapshots may be taken in “conjunction with, initiating the security assessment”); *infra*, §XI.B.5 (discussing snapshot-analysis techniques). A snapshot represents the state of a virtual asset such as a VM.

EX1007, Abstract, 9:14-16 (snapshot represents “state of the target resource”), 16:15-26 (snapshot “of the target VM”); EX1002, ¶¶131-32. For example, it can be a snapshot of a logical volume of a VM (i.e., a “virtual disk of a protected virtual cloud asset”). EX1007, Abstract, 3:20-40 (snapshot may be copy of “the state of any devices (virtual or physical) allocated to the resource, block-level image of the entire logical volume” or portions thereof), 8:52-62 (target resources may be “virtual computing resources or segment(s) thereof, such as a virtual machine, ... logical volume, physical or virtual storage device, etc.”), 11:28-33 (logical volume “also known as ... ‘virtual disk drive’”); *see also id.*, 2:22-35, 3:61-67, 6:7-13, 8:52-9:9, 11:5-67, 16:15-26, 18:16-26, Figs. 2-6; EX1002, ¶¶91, 133-34. Veselov further teaches that these VMs are “protected virtual cloud” assets since it teaches securing such resources in cloud computing environments. *E.g.*, EX1007, 2:42- 3:19; *see also supra*, §§IX.A (cloud computing), XI.B.1 (preamble); EX1002, ¶¶135-36.

Veselov also teaches “wherein the virtual cloud asset is instantiated in the cloud computing environment.” The previously discussed protected virtual cloud assets are instantiated in the host environment (i.e., cloud computing environment). EX1007, Abstract, 6:9-34, 8:52-60, 10:37-48; EX1002, ¶¶137-38. Instantiated VMs are shown, for example, in Fig. 1:



EX1007, Fig. 1 (annotation added); *see also id.*, Figs. 3A-B, 5A-B (depicting VM in host environment).

4. Elements 1.2 and 18.2

Veselov teaches these elements, which recite “[accessing/access] the snapshot of the at least one virtual disk.” For the previously discussed snapshot (*supra*, §XI.B.3), the scanning service accesses/obtains the snapshot as part of the snapshot-analysis described below. EX1007, 6:38-46 (scanning service determines virtualization layer and “transmit[s] a command to the virtualization layer 144 to provide snapshot data 146 (or to provide access to snapshot data 146)”, 9:9-41 (scanning service may obtain snapshot “by accessing a data store” or “obtain access to the snapshot data, such as at a storage location of the snapshot data”); *see also id.*, 3:20-27, 4:53-5:22, 6:25-49, 11:34-36, 22:60-23:18, Figs. 2, 4, 6

(depicting obtaining step), claim 1; *infra*, §XI.B.5; EX1002, ¶¶139-141 (citing EX1074, 9, Fig. 1-10).

5. Elements 1.3 and 18.3

Veselov teaches these elements, which recite “[analyzing/analyze] the snapshot of the protected virtual cloud asset to detect potential cyber threats risking the protected virtual cloud asset.”

First, Veselov teaches “[analyzing/analyze] the snapshot of the protected virtual cloud asset.” Veselov states that “the scanning service 110 may obtain and **analyze snapshot data** 146.” EX1007, 4:49-54 (emphasis added); *see also id.*, 8:16-17 (“The scanning service 110 performs the desired security assessment using the snapshot data 146.”), 8:17-51 (describing variety of analysis techniques), Figs. 1 (depicting scanning service 110 analyzing snapshot data 146 to produce assessment results 132), 2 (step 210), Figs. 3A-B (scanning service 310), 4 (step 414), 5A-B (scanning service 510), 6 (step 614); EX1002, ¶¶142-43.

Veselov discloses two general ways to perform the analysis, each of which satisfies the previously discussed construction. EX1007, 8:30-45 (scanning service may “reconstitute an active duplicate virtual machine instance or ... analyze the snapshot as a data file”), 10:1-16 (describing types of analyses); *supra*, §VIII.B (construction). First, the scanning service can instantiate a duplicate/assessment VM from the snapshot and analyze this VM. EX1007, 3:54-61, 12:1-20 (same);

EX1002, ¶¶144-45; *see also* EX1007, 3:44-54 (duplicate VM may be instantiated in original or test environment), 14:13-22, Figs. 3A (assessment VM 312A in original computing environment), 3B (assessment VM 312B in testing environment).

Second, the scanning service can mount the snapshot as a data file and directly analyze the snapshot data. EX1007, 3:61-67 (scanning service may mount logical volume as hard drive or transform the snapshot into a data file “then perform the security assessment on the data at the file system level”), 8:30-45 (scanning service may “analyze the snapshot as a data file”), 16:56-17:29 (service mounts volume image and reads files); EX1002, ¶146. Consistent with well-known security-assessment techniques, neither of these analysis techniques in Veselov uses an agent installed on the target VM, nor do they require an agent installed on the assessment VM. *Supra*, §X.A; EX1002, ¶¶56-57, 147 (citing EX1040, Abstract (monitoring virtual server “without utilizing agents executing within the virtual server”), ¶¶8-12, 23, 43-45, 48, 52; EX1073, 1; EX1095, ¶¶5-6).

Veselov also teaches analyzing the snapshot to detect “potential cyber threats risking the protected virtual cloud asset” in several ways. Veselov “contemplates implementation of any suitable security assessment.” EX1007, 2:35-41; *see also id.*, 1:41-49 (known services identify “vulnerabilities, bad configurations, and the like”), 4:1-18 (evaluation of vulnerabilities), 4:53-59 (identifying “security risk information”), 5:9-11, 7:1-4, 10:17-36, 14:41-15:3 (scan

configuration determines “tests or tasks to perform, types of vulnerabilities to evaluate, and the like”), 20:57-21:3, 23:19-24:4, claims 1, 16-18; EX1002, ¶148. Security assessments for VMs routinely included threat detection. EX1002, ¶148; *see also id.*, ¶¶49-53 (common security risks), 54-60 (common detection techniques). Veselov also discloses several examples that encompass detecting cyber threats to an asset.

First, Veselov teaches applying rule packages such as CVEs, which were known to identify cyber threats to the asset. EX1007, 2:22-54, 7:16-22; EX1002, ¶149 (citing EX1077, 8:1-15); *supra*, §§IX.B, XI.B.1; *see also* EX1001, claims 3, 20 (cyber threats include vulnerabilities). For example, many CVEs were associated with application features risking an asset’s security and common startup services. EX1002, ¶149 (citing EX1090; EX1091; EX1100; EX1101; EX1102; EX1103). The ’798 specification identifies these types of security problems as “cyber threats.” EX1001, 6:19-30 (listing, e.g., “risky application features” and “services running on startup”). The ’798 specification also indicates that vulnerabilities are a type of cyber threat. *Id.*, 4:43-45, 6:38-43, 7:11-14; EX1002, ¶149.

Second, Veselov teaches identifying nonsecure system configurations, which also encompass “potential cyber threats risking the protected virtual cloud asset.” Veselov discloses performing “host configuration assessments” and describes

various types of configuration data stored in the analyzed snapshot. EX1007, 2:35-31 (analysis techniques include “host configuration assessments), 9:42-67 (listing properties of snapshot, including configuration information), 21:53-22:6, claim 9; *see also id.*, 1:46-49 (known scanning services identify “vulnerabilities, bad configurations, and the like”), EX1002, ¶150. A POSA would have understood detecting nonsecure configurations to include, or at least suggest, determining the existence of potential cyber “threats” risking the asset since such configurations pose potential threats to system security. EX1002, ¶150. This is confirmed by the ’798 specification, which identifies several types of “cyber threats” based on nonsecure system configurations. EX1001, 6:19-30 (listing, e.g., weak passwords and weak encryption schemes); EX1002, ¶150.

Consistent with Veselov, Basavapatna similarly detects “potential cyber threats risking” an asset in multiple ways. Basavapatna performs a threat-centric assessment that involves detecting potential threats to the asset. EX1008, ¶¶31, 43-49, 76-81, 83; *see also id.*, ¶¶20, 16, 37-40; EX1002, ¶151; *supra*, §XI.A.

Basavapatna similarly teaches a vulnerability-centric assessment that also includes detecting “potential cyber threats risking” the asset. EX1008, ¶¶52 (threat vectors linked within vulnerability definition data), 99-104 (detects vulnerabilities based on vulnerability definition data); *see also id.*, ¶¶53, 60, 104 (vulnerability score estimates whether asset vulnerable to threat); EX1001, 4:43-51 (vulnerabilities are

type of cyber threat that include vulnerable software and non-secure configurations); EX1002, ¶¶152-53.

6. Elements 1.4 and 18.4

Veselov teaches these elements, which recite “[alerting/alert] detected potential cyber threats based on filtering or prioritizing using a determined risk of each detected threat.”

Veselov teaches alerting detected potential cyber threats. For example, Veselov describes sending results describing the detected cyber threats for display to the user. EX1007, 10:17-36 (results reporting in “user-readable format to a user interface to be displayed on the user device”), 18:53-56, 20:57-21:13, 23:19-47, Figs. 2, 4, 6 (depicting results-providing step); *see also id.*, 5:9-13, 8:19-21, claim 1; EX1002, ¶¶154-55 (citing EX1025, 21).

Although Veselov does not expressly state that its alerts are “based on filtering or prioritizing using a determined risk of each detected threat,” Basavapatna teaches this with its discussion of alerting based on filtering or prioritizing using several determined risk metrics. EX1002, ¶156. These risk metrics include threat-centric and vulnerability-centric risk metrics and various associated aggregate risk metrics discussed below. *Supra*, §XI.B.5.

First, Basavapatna teaches determining a threat-centric risk metric for each threat detected on the asset. EX1008, ¶¶73-98, Fig. 3A. The threat-centric risk

metric is calculated based on multiple factors that quantify risk associated with cyber threats detected on the asset. *Id.*, ¶¶83-85; EX1002, ¶157. For example, an exposure factor for the asset “estimates the risk that an asset will be affected by the threat.” EX1008, ¶82; EX1002, ¶157.

Second, Basavapatna similarly teaches determining a vulnerability-centric risk metric for each vulnerability on the asset. EX1008, ¶¶99-113, Fig. 3B. For example, this risk metric is based on a “standardized vulnerability score” and a “composite vulnerability score,” each of which is associated with a risk to the asset. *Id.*, ¶¶100-06, 112, Fig. 3B; *see also id.*, ¶¶Abstract, 113; *supra*, §XI.B.5 (discussing vulnerabilities as type of threat); EX1002, ¶158. A POSA would have understood these risk determinations as encompassing risk determinations of each cyber threat. *Supra*, §XI.B.5; EX1002, ¶158.

Third, Basavapatna teaches aggregate risk metrics. For example, Basavapatna teaches determining a per-vulnerability aggregate risk by calculating “an aggregate risk metric for the particular vulnerability from the risk metrics for the assets for the particular vulnerability.” EX1008, ¶¶123-128; EX1002, ¶159. The threat-centric risk metrics may similarly be used to determine a per-threat aggregate risk. EX1008, ¶122. Second, Basavapatna teaches aggregating the vulnerability-centric or threat-centric risk metrics for an asset to determine an

asset-level risk. EX1008, ¶¶130, 132, Fig. 5; *see also id.*, ¶¶133-35; EX1002, ¶159.

Lastly, Basavapatna teaches alerting “based on filtering or prioritizing using” the determined risk of each detected cyber threat. Basavapatna teaches allowing the user to “view assets, vulnerabilities, or threats, sorted by the aggregate risk metric for the asset, vulnerability, or threat” (i.e., prioritizing using the determined risk) and further filter them by listing “a top number, e.g., top ten, of the assets, vulnerabilities, or threats.” EX1008, ¶138; *see also id.*, ¶139; EX1002, ¶160. Moreover, Basavapatna teaches alerting based on filtering or prioritizing by alerting the user if an aggregate risk metric “rises above a specified threshold.” EX1008, ¶140; *see also id.*, ¶138 (discussing importance of “[r]anking assets, vulnerabilities, and/or threats” to allow user to quickly identify most dangerous risk); EX1002, ¶¶160-61.

Accordingly, claims 1 and 18 were obvious.

C. Dependent Claims

1. Claims 2 and 19

The combination teaches claims 2 and 19, which depend from claims 1 and 18, respectively, and further recite “[prioritizing/prioritize] each of the detected potential cyber threats based on their respective risk to the protected virtual cloud asset.”

As an initial matter, both references also teach providing alerts for “each of the detected potential cyber threats,” as recited. EX1007, 10:17-36 (assessment results describe one or more threats); EX1008, ¶¶14-16 (determining risk metrics for one or more threats), 30, 138, claim 27; *supra*, §XI.B.6; EX1002, ¶¶162-63.

Moreover, Basavapatna teaches determining a risk metric for each cyber threat “based on their respective risk to the” protected asset. As discussed above, Basavapatna discloses determining several different risk metrics that represent a risk to the asset. *Supra*, §XI.B.6; EX1008, ¶¶92, 112, 120-135; EX1002, ¶164. First, the threat-centric risk metric is also based on several factors that each represent a risk to the asset, including the exposure factor, the applicability score, and the countermeasure score. EX1008, ¶¶12, 16, 82, 92; EX1002, ¶165. Second, the vulnerability-centric risk metric is also based on several factors that each represent a risk to the asset, including a standardized vulnerability score, a composite vulnerability score, and a countermeasure component score. EX1008, ¶¶12-13, 106, 112; EX1002, ¶166. Further, the aforementioned aggregate risk metrics are derived from either the vulnerability-centric or the threat-centric risk metrics and as such are similarly risks to the asset. *Supra*, §XI.B.6; EX1008, ¶¶120-135; EX1002, ¶167.

Basavapatna also teaches “prioritizing each of the detected potential cyber threats” based on their associated risks to the protected asset. As discussed above,

Basavapatna teaches prioritizing each cyber threat by reporting/alerting the threats based on their determined risk metrics, which represents risk to the asset as discussed above. *Supra*, §XI.B.6; EX1008, ¶¶138-40; EX1002, ¶167. Moreover, as previously explained, the combination teaches using these techniques for protected “virtual cloud” assets. *Supra*, §§XI.A, XI.B.1. Therefore, the combination teaches “prioritizing each of the detected potential cyber threats based on their respective risk to the protected virtual cloud asset,” as recited. EX1002, ¶¶168-70.

Accordingly, claims 2 and 19 were obvious.

2. Claims 3 and 20

Claims 3 and 20 depend from claims 1 and 18, respectively, and further recite “wherein analyzing the snapshot of the protected virtual cloud asset further comprises.” Both claims then recite two additional limitations taught by Veselov.

a. Claims 3.1 and 20.1

Veselov teaches these elements, which recite “parsing a copy of the snapshot.”

First, Veselov teaches use of a “copy” of a snapshot. As previously discussed, Veselov describes multiple ways of obtaining, or obtaining access to the snapshot, including bringing a copy of the snapshot into a testing environment within or outside of the target VM’s environment. *Supra*, §§XI.B.4-5; *see, e.g.*,

EX1007, 9:16-25, Figs. 3A (snapshot copy brought to test environment 305), 5A (snapshot copy brought to test environment 505); EX1002, ¶¶171-73. Thus, Veselov teaches methods and systems for intaking and using a copy of a snapshot as a part of analyzing the snapshot.

Veselov also describes “parsing” this copy of the snapshot as part of the analysis in several ways. The scanning system may mount the logical volume “or transform the snapshot into a data file.” EX1007, 3:61-67. Following this transformation, the system may “perform the security assessment on the data at the file system level.” *Id.* Veselov further explains that the scanning service may **“transform the volume image 508 into a format in which the scanning service 510 can read the copied files and other data”** *Id.*, 17:4-8 (emphasis added); *see also id.*, 8:21-23 (processing may involving transforming snapshot data), 9:32-37 (adding metadata tags). EX1002, ¶¶91, 174. A POSA would have understood that these transformations involve “parsing” the copy of the snapshot, as recited. EX1002, ¶174. At a minimum, parsing the snapshot would have been obvious in view of these teachings since parsing was a well-understood, commonly used technique for interpreting unstructured data such as VM images. *Id.*, ¶¶175-76 (citing EX1033, 6-7; EX1034, 5643; EX1035, 387; EX1037, 10:2; EX1038, 134; EX1039, 3; EX1041, 3:1-9; EX1042, 10:45-55, 12:54-60; EX1043, 12:10; EX1007, 17:4-8); *supra*, §IX.B.

b. Claims 3.2 and 20.2

The combination teaches these elements, which recite “scanning the parsed copy to detect the potential cyber threats, wherein the potential cyber threats include known and unknown vulnerabilities, and wherein the detection is based on a type of vulnerability.”

First, after transforming/parsing the snapshot copy, the scanning service scans this parsed copy. EX1007, 3:62-67 (system “may then perform the security assessment on the data at the file system level”); *see also id.*, 18:16-44 (scanning service may perform security assessment by first “obtain[ing] a snapshot (i.e., volume image) of the virtual machine” and then “reading the files of the mounted volume image” and performing “string or other pattern comparisons” to identify files, software, and other data); EX1002, ¶¶91, 177-78. Additionally, as previously discussed, Veselov’s scanning—which includes this scanning of the parsed copy—detects the cyber threats. *Supra*, §XI.B.5; EX1007, 18:44-53; EX1002, ¶178. Veselov thus teaches, or at least suggests, scanning the parsed copy of the snapshot to detect the potential cyber threats, which was a common, well-understood technique for performing the type of analysis described in Veselov. EX1002, ¶¶58, 179 (citing EX1007, 18:26-44).

Second, Veselov teaches that the cyber threats it addresses include “known vulnerabilities” as recited. As previously discussed, Veselov contemplates that its

methods will be used to implement “security assessments that are defined by rules packages, such as **Common Vulnerabilities and Exposures (CVEs)**, Center for Internet Security (CIS) benchmarks, ‘best practices’ packages, static or runtime behavior analysis, host configuration assessments, and the like.” EX1007, 2:35-41 (emphasis added); *supra*, §XI.B.5; *see also* EX1002, ¶¶180-81 (citing EX1023, 1-3; EX1025, 12 (typical scans routinely detected known vulnerabilities)). The ’798 patent acknowledges that CVEs are lists of known vulnerabilities, stating “[t]he information provided by the intelligence systems may detect certain **known vulnerabilities identified in, for example, a CVE database.**” EX1001, 4:36-42 (emphasis added); EX1002, ¶181.

Veselov also teaches, or at a minimum suggests, that the cyber threats it addresses include “unknown vulnerabilities,” as recited. Veselov describes analysis techniques that go beyond looking for specific, known vulnerabilities. It builds upon prior systems that “analyze the behavior of computing resources to identify vulnerabilities, bad configurations, and the like.” EX1007, 1:46-50. It further describes “static or runtime behavior analysis” and “analyzing behavior of the second virtual machine instance in response to the simulated network traffic.” *Id.*, 2:35-41, claim 5; *see also id.*, 16:6-14 (scanning service may generate “‘dummy’ traffic” and “perform destructive testing on the assessment virtual machine 312A-B to verify the system’s health”). As Dr. Stavrou explains, a POSA

would have recognized these types of analysis were well-known techniques for identifying unknown vulnerabilities. EX1002, ¶¶87, 182 (citing EX1025, 26-30); *see also supra*, §IX.B (identification of unknown vulnerabilities was routine part of security analysis). Thus, Veselov teaches or at least suggests that “the potential cyber threats include known and unknown vulnerabilities.”

This limitation was also obvious further in view of Basavapatna. Concerning known vulnerabilities, Basavapatna’s “standardized vulnerability score” is used to rank the potential severity of “a particular vulnerability in a plurality of **known vulnerabilities**.” EX1008, ¶13 (emphasis added). Concerning unknown vulnerabilities, Basavapatna discusses “passive countermeasures” that may “protect against **unknown vulnerabilities**.” *Id.*, ¶22 (emphasis added); *see also id.*, ¶17 (describing threats not associated with any known vulnerabilities). Given the well-known benefits of identifying both known and unknown vulnerabilities—e.g., identifying vulnerabilities that create risks to the computer system but have not yet been defined in rules packages—and the commonplace, routine nature of performing such scans, a POSA had good reason to identify both types of vulnerabilities when combining the teachings of Veselov and Basavapatna. EX1002, ¶¶50, 54, 183-84; *supra*, §§IX.B, XI.A.

Lastly, Veselov further teaches “wherein the detection is based on a type of vulnerability.” Veselov states that “[v]arious aspects of the security assessment(s)

may be configured by user input, non-limiting examples of such aspects including: selection of the assessment target(s); parameters of the security assessment(s), such as tests or tasks to perform, **types of vulnerabilities to evaluate**, and the like....”

EX1007, 14:41-46 (emphasis added); *see also id.*, 2:35-41 (describing assessments based on, e.g., rules packages and CVEs), 7:16-22; EX1002, ¶¶91, 185. Moreover, as Dr. Stavrou explains, a POSA would have understood that a security assessment based on rules packages such as CVEs were typically based on the types of vulnerabilities. EX1002, ¶¶50, 59, 61, 185-87 (citing EX1075, 3:43-54, 4:46-64; EX1078, ¶¶33, 84, 113); *supra*, §XI.B.5 (CVE-based analysis).

Accordingly, claims 3 and 20 were obvious.

3. Claims 4 and 21

Claims 4 and 21 depend from claims 3 and 20, respectively, and further recite “wherein scanning the parsed copy further comprises any one of.” Both claims then recite four limitations, any of which is sufficient to satisfy each claim. EX1002, ¶¶188-89.

a. Claims 4.1 and 21.1

These elements recite “checking configuration files of applications and operating system installed in the protected virtual cloud asset.” Veselov alone, or in combination with Basavapatna, teaches these additional limitations.

Veselov teaches checking configuration files of applications and operating system installed in the VM. The snapshot may include “any properties of the target resource ... that are relevant to the security assessment,” including “operating system information” as well as “software packages installed, and configurations thereof.” EX1007, 9:55-67; *see also id.*, Abstract, 1:22-34 (describing configurations of operating system and installed applications), 7:16-22; EX1002, ¶¶190-92. Veselov further teaches checking these configuration files when scanning the parsed copy of the snapshot. EX1007, 1:43-50, 2:35-41 (applying “host configuration assessments”), 7:16-22 (using configuration data in snapshot), 18:32-50 (image scan includes character-pattern analysis to identify file/application features), 21:61-22:6 (analyzing VM configuration), claim 9 (same); *supra*, §X.C.2.b; EX1002, ¶193.

This limitation is also taught further in view of Basavapatna. For example, Basavapatna teaches that the vulnerability-centric and threat-centric risk metrics are determined by analyzing asset configuration data that includes “the configuration of the software running on the asset.” EX1008, ¶37; *see also id.*, ¶¶53, 60, 79, 105 (determining whether asset has “configurations evidencing the vulnerability”); EX1002, ¶194. Basavapatna also teaches checking configuration files of an installed operating system. *Id.*, ¶¶37 (asset configurations include “specific versions of operating systems”), 44 (applicability based on, e.g.,

“particular operating systems” targeted by threat), 53, 60, 78, 80, 104 (identifying vulnerability where asset runs operating system with known vulnerability); *supra*, §XI.A; EX1002, ¶¶194-95.

b. Claims 4.4 and 21.4

The combination teaches these elements, which recite “analyzing machine memory stored in the snapshot to deduce what applications and modules executed in the protected virtual cloud asset.”

Veslov teaches that the previously discussed scanning of the parsed copy includes “analyzing machine memory stored in the snapshot,” as recited. For example, Veslov explains that the analyzed snapshot may be a copy of the VM’s allocated devices, an image of the entire volume, or the “state of memory.”

EX1007, 3:27-40; EX1002, ¶¶196-97; *supra*, §XI.B.4-5, XI.C.2.b, XI.C.3.a-b.

The combination of Veslov and Basavapatna further teaches analyzing this memory “to deduce what applications and modules executed in the protected virtual cloud asset.” Veslov explains that the snapshot (which may include the state of memory as discussed above) includes properties such as the “software packages installed” on the VM and that the scanning service may identify which assessment tasks to run “based on the installed software packages.” EX1007, 9:55-67, 18:44-50; *see also id.*, 7:15-22 (assessment includes determination of installed software); *supra*, §XI.C.3.a; EX1002, ¶198. Veslov also teaches detecting cyber

threats associated with such software packages, as explained above. *Supra*, §§XI.B.5, XI.C.3.a; EX1007, 2:35-41, 7:16-22, 14:41-65; EX1002, ¶198. Consistent with these teachings, Basavapatna teaches identifying “running” software in the scanned asset as part of its previously discussed risk assessments, which teaches or at least suggests deducing the applications and modules “executed” in the asset. EX1008, ¶¶14-16, 53, 60 (asset vulnerable when running particular version of software), 78 (threat applicable when asset running exploitable software), 78-81, 104, claims 22, 27; *see also id.*, ¶37; *supra*, §XI.A; EX1002, ¶¶68, 199-200 (citing EX1063, ¶¶80, 100-12).

Accordingly, claims 4 and 21 were obvious.

4. Claims 5 and 22

Claims 5 and 22 depend from claims 1 and 18, respectively, and further recite two additional limitations.

a. Claims 5.1 and 22.1

These elements recite “instantiating an instance copy of the protected virtual cloud asset from the snapshot.” Veselov teaches these additional limitations.


As an initial matter, Veselov teaches, or at least suggests, performing multiple types of scans involving the same snapshot. Veselov describes multiple types of scans (*supra*, §X.A) and suggests combining the various scanning techniques it describes. *E.g.*, EX1007, 24:54-25:2 (“Furthermore, it is to be

understood that particular features, structures, or characteristics described may be combined in various ways in one or more embodiments.”); EX1002, ¶¶201-203. Veselov also teaches or suggests performing multiple scans involving the same snapshot. EX1007, 4:1-18 (after scan, system may “store the snapshot in the account of the user so that the user can perform additional analysis on the snapshot”), 14:37-65 (“In various embodiments, an owner or administrator of the target computing resources to be assessed may configure one or more security assessments....”); EX1002, ¶203. Veselov thus teaches or suggests performing multiple types of analysis on the snapshot, including both the parsing-based techniques discussed above with respect to claims 3 and 20—from which claims 5 and 22 respectively depend—and the VM-instantiation techniques discussed below. EX1002, ¶203; *supra*, §XI.C.2 (claims 3 and 20). A POSA also had reason to perform additional scans using Veselov’s instantiation approach in some circumstances to increase the chances of detecting security issues (e.g., issues detected by monitoring live activity) including those discussed further in §XI.C.4.b. EX1007, 15:41-16:3; EX1002, ¶203.

Turning to elements 5.1 and 22.1, Veselov explains, “the scanning service may implement a **duplicate** virtual machine that is **instantiated** with the same configuration.” EX1007, Abstract (emphasis added). Veselov further teaches using the snapshot data “to produce a duplicate implementation of the target

resource that can be scanned by the scanning service,” including by “**launching a duplicate virtual machine instance.**” *Id.*, 10:1-16 (emphasis added); EX1002, ¶204.

Veselov illustrates an example of this approach in which assessment VM 312A is instantiated from the snapshot:



EX1007, Fig. 3A; *see also id.*, Fig. 3B (depicting analogous assessment VM 312B). The duplicate “assessment virtual machine 312A-B may be **instantiated** with the same configuration as the target instance of the original virtual machine 312 had at the time the volume image was captured.” *Id.*, 13:53-58 (emphasis

added); *see also id.*, 3:54-61 (service launching an active duplicate instance from snapshot data), 12:1-20 (same); EX1002, ¶¶205-06.

b. Claims 5.2 and 22.2

Veselov teaches these elements, which recite “monitoring all activity performed by the instance copy of the protected virtual cloud asset.”

In embodiments involving instantiated assessment VMs, the scanning service causes “processing activity to occur within the assessment virtual machine 312A-B” in order to assess “real-time dynamic or static data, network connections, communications to computing resources, and other computational activities of the assessment virtual machine 312A-B.” EX1007, 15:41-57. The scanning service also “intercept[s] traffic to the target virtual machine 312 and duplicate the intercepted traffic to the assessment virtual machine 312A-B in order to stimulate activity in the assessment virtual machine 312A-B.” *Id.*, 15:60-16:10; *see also id.*, claim 5 (reciting “analyzing behavior of the second virtual machine instance”). Veselov thus teaches, or at a minimum suggests, monitoring “all activity” performed by the instance, as recited. EX1002, ¶¶207-10.

Accordingly, claims 5 and 22 were obvious.

5. Claims 7 and 24

Claims 7 and 24 depend from claims 3 and 20, respectively, and further recite “wherein scanning the parsed copy further comprises...” Both claims then recite four limitations, any of which is sufficient to satisfy the claims.

a. Claims 7.1 and 24.1

The combination teaches these elements, which recite “matching at least one installed application list against a list of vulnerable applications.”

Veselov teaches, or at least suggests, that scanning the parsed copy further comprises “matching at least one installed application list against a list of vulnerable applications.” As discussed above, Veselov teaches the parsed snapshot data scanned as part of the security assessment includes “software packages installed.” EX1007, 9:55-67; *supra*, §XI.C.2. Veselov also teaches identifying “files, software packages, application versions, and the like” that are installed on the target VM and running assessment tasks such as CVE assessments “based on the installed software packages.” EX1007, 7:6-22, 18:32-53, 22:35-46, claim 12; *see also id.*, 2:35-41, 14:58-65; EX1002, ¶¶91, 211-13. This discussion of analyzing the set of installed applications at least suggests analyzing a “list” of installed applications. A POSA also would have understood the disclosed analysis techniques such as CVE assessments as encompassing, or at least suggesting, matching this list against a list of vulnerable applications since this was a routine

approach for performing such assessments. EX1007, 2:35-41, 7:16-22, 14:58-65; EX1002, ¶214 (citing EX1023, 1-3; EX1044, 4:22-35, 13:59-14:7; EX1076, 2:24-45; EX1077, 2:45-51, 8:1-30; EX1084, 4:25-33, 6:25-38); *see also* EX1001, 4:40-42.

While this limitation is taught or suggested by Veselov alone, it is further taught by Basavapatna. As discussed previously, Basavapatna's threat-centric and vulnerability-centric analysis includes comparing "applicability data" with asset configuration data that includes "software running on the asset." EX1008, ¶¶37, 53; *see also id.*, ¶¶60, 76-85 (threat-centric risk), 104-06 (vulnerability-centric risk), Figs. 3A-B; *supra*, §XI.B.5; EX1002, ¶215. This comparison at least suggests matching an installed application "list" against a "list" of vulnerable applications since it involves comparing a set of installed/running software against a set of known vulnerable software. *See, e.g.*, EX1008, ¶¶44, 53; EX1002, ¶215. Accordingly, Basavapatna teaches "matching at least one installed application list against a list of vulnerable applications," as recited. EX1008, ¶60; EX1002, ¶¶215-16; *see also supra*, §XI.A (reasons to combine).

b. Claims 7.2 and 24.2

The combination teaches these elements, which recite "matching at least one application file against a database of files of vulnerable applications." The previous discussion of matching a "list" of installed applications against a list of

vulnerable applications also shows why matching an application “file” against a database of files of vulnerable applications was obvious. *Supra*, §XI.C.5.a; EX1002, ¶217. Moreover, a POSA would have understood that the types of matching-based threat detection taught by the references routinely included matching an application “file” against a “database” of vulnerable applications. EX1002, ¶218 (citing EX1017, 7:58-64; EX1023, 3-4); EX1007, 2:35-41, 3:20-40, 7:16-22, 9:42-67, 17:9-29; EX1008, ¶¶44, 53, 60, 79-80. For example, CVE detection was commonly performed by comparing files’ binary data (or hashes thereof) against a database of vulnerable application files. EX1002, ¶218 (citing EX1047, 1-5, 7; EX1068, 14:58-15:18; EX1076, 2:42-67; EX1084, 3:56-4:47, 11:27-47). Thus, at a minimum, it would have been obvious for the scan to include “matching at least one application file against a database of files of vulnerable applications,” as recited, since this was a standard, straightforward way of performing the matching analysis disclosed in the references. EX1002, ¶¶218-19.

c. Claims 7.4 and 24.4

Veslov teaches these elements, which recite “reading installation logs of package managers.” Veslov teaches that the previously discussed CVE assessment may be based on a package manager’s list of installed applications. EX1007, 7:6-22; *supra*, §XI.C.5.a. A POSA would have understood this process to encompass, or at least suggest, reading the package managers’ installation logs

since that was a simple, common way of obtaining the set of installed applications that would be analyzed in the disclosed application-level CVE assessment.

EX1002, ¶¶220-22 (citing EX1117, 2; EX1118, 5).

Accordingly, claims 7 and 24 were obvious.

6. Claims 8 and 25

Veselov teaches claims 8 and 25, which depend from claims 1 and 18, respectively, and further recite “wherein the protected virtual cloud asset includes any one of: a virtual machine, a software container, or a micro-service.” Veselov teaches that the protected virtual cloud asset includes a “virtual machine” and a “software container,” either of which satisfies the claims. Veselov describes performing a security assessment on a VM or instance thereof. EX1007, Abstract; *see also id.*, 3:30-32; *supra*, §XI.B.3; EX1002, ¶¶223-24. The target resource may also be a software container or a VM that includes software containers. EX1007, 1:34-40, 2:44-51; EX1002, ¶¶224-25.

Accordingly, claims 8 and 25 were obvious.

7. Claims 9 and 26

Veselov teaches claims 9 and 26, which depend from claims 1 and 18, respectively, and further recite “wherein taking the snapshot is performed at a predefined schedule.” Veselov explains that the user may authorize periodic assessments of a target resource, which cause the scanning system to “obtain and

analyze snapshots on a scheduled basis.” EX1007, 4:11-14; *see also id.*, 8:62-9:5, 14:41-53, 15:15-23, 21:23-32, 23:48-56, claim 4; EX1002, ¶226. Moreover, as discussed above, Veselov teaches taking a snapshot as part of obtaining a snapshot. EX1007, 8:62-9:25, 3:22-27; *supra*, §XI.B.3; EX1002, ¶¶226-27.

Accordingly, claims 9 and 26 were obvious.

8. Claims 10 and 27

Veselov teaches claims 10 and 27, which depend from claims 1 and 18, respectively, and further recite “wherein taking the snapshot is performed based on receiving a request to take the snapshot.” Veselov explains that the assessment “may be initiated by receipt of a request.” EX1007, 6:30-31; *see also id.*, 6:30-43, 14:41-52, Fig. 1. This teaches or at least suggests taking the snapshot based on receiving a request to take the snapshot since the requested assessment may include taking a snapshot. *Supra*, §XI.B.3; EX1002, ¶228. Additionally, as discussed above, Veselov teaches that the scanning service can obtain a snapshot “by requesting another service to capture the snapshot and provide it to the scanning service.” EX1007, 9:14-25; *see also id.*, 6:38-43, 7:24-31, 8:60-9:14; *supra*, §XI.B.3. Veselov thus also teaches having a separate service take the snapshot based on a request from the scanning service. EX1002, ¶¶228-29

Accordingly, claims 10 and 27 were obvious.

9. Claims 11 and 28

Veselov teaches claims 11 and 28, which depend from claims 1 and 18, respectively, and further recite “wherein taking the snapshot is performed based on detection of a network event or an abnormal event.” Veselov teaches that the user may identify “trigger events that cause the assessment to begin.” EX1007, 14:41-57. This teaches or suggests taking a snapshot based on detection of such events since Veselov also teaches that taking the snapshot may be a preliminary step in the assessment. *Supra*, §XI.B.3; EX1007, 9:14-25; EX1002, ¶230. A POSA would have understood these event-trigger teachings to encompass, or at least suggest, a network event or abnormal event since these were typical types of assessment triggers for event-triggered security assessments such as Veselov’s. EX1002, ¶¶230-31 (citing EX1001, 7:19-25; EX1036, 14:51-15:13 (scan triggered by events such as “changes in the external environment”); EX1064, Abstract, ¶¶31, 42-44, 63).

Accordingly, claims 11 and 28 were obvious.

10. Claims 12 and 29

Veselov, alone or in combination with Basavapatna, teaches claims 12 and 29, which depend from claims 1 and 18, respectively, and further recite “further comprising mitigating a potential cyber threat posing a risk to the protected virtual cloud asset.” Veselov teaches comparing assessment results to a remediation

framework, identifying remedial actions, and “providing the identified actions to the user.” EX1007, 10:17-36. Veselov thus teaches or suggests providing actions that address the identified security problems (e.g., fixing the problem or taking an action to prevent damage). EX1002, ¶¶232-33 (citing EX1040, ¶49, EX1070, Abstract; EX1056, ¶¶41, 60, 85). Basavapatna similarly teaches mitigating the identified cyber threats using various countermeasures, and it was commonplace for security systems to mitigate detected cyber threats posing a risk to the asset. EX1008, ¶¶20-25, 28-29; EX1002, ¶¶69-70, 234-35 (citing EX1025, 24).

Accordingly, claims 12 and 29 were obvious.

11. Claim 13 and 30

Claims 13 and 30 depend from claims 1 and 18, respectively, and further recite “wherein analyzing the snapshot to detect potential cyber threats includes detecting cyber threats that do not represent vulnerabilities.” As explained below, Veselov alone, or in combination with Basavapatna, teaches these limitations.

Veselov teaches analyzing the snapshot to detect “cyber threats that do not represent vulnerabilities.” As explained above, Veselov broadly teaches identifying various cyber threats. *Supra*, §XI.B.5; EX1007, 1:46-49 (discussing identification of “vulnerabilities, bad configurations, and the like”). This includes cyber threats associated with risky application features and services run on startup (e.g., with its teachings of snapshot-based CVE assessments), which the ’798

specification identifies as examples of cyber threats that do not represent vulnerabilities. *Supra*, §XI.B.5; EX1001, 6:19-30; EX1007, 15:41-59, 16:56-67; EX1002, ¶¶91, 236-37.

Basavapatna also teaches detecting “cyber threats that do not represent vulnerabilities,” consistent with the ’798 specification. EX1001, 6:19-30 (examples include weak passwords and authentication methods). Basavapatna teaches or at least suggests detecting weak passwords and weak authentication methods. EX1008, ¶21 (discussing security problems such as when “a browser setting is too loose, or a password is too short”); EX1002, ¶238. Basavapatna also discusses analyzing the types of system configurations well known to include authentication and password configurations. EX1008, ¶¶14-16, 34-37, 44, 49; EX1002, ¶¶238-39; *see also supra*, §XI.A.

Accordingly, claims 13 and 30 were obvious.

12. Claims 14 and 31

Claims 14 and 31 depend from claims 1 and 18, respectively, and further recite “wherein analyzing the snapshot to detect potential cyber threat includes detecting a vulnerable application file by comparing a hash value of the vulnerable application file against a database of known application files.” A POSA would have understood the types of matching-based security assessments discussed in Veselov and Basavapatna (e.g., CVE assessments) to include, or at least suggest,

the recited hash comparison since this type of matching was routinely performed by computing hashes of application files and comparing those hashes against a database associated with known vulnerable applications, such as a CVE repository. *Supra*, §XI.C.5; EX1002, ¶¶240-41 (citing EX1017, 7:58-60; EX1053, ¶162; EX1063, ¶¶75, 244-48; EX1068, 4:29-39; EX1076, 2:39-59). The additional limitations of claims 14 and 31 therefore would have been obvious given the references' discussion of matching-based assessments and the POSA's familiarity with typical ways of performing such matching. EX1002, ¶¶241-42.

Accordingly, claims 14 and 31 were obvious.

13. Claims 15 and 32

Veslov teaches claims 15 and 32, which depend from claims 1 and 18, respectively, and further recite “wherein analyzing the snapshot to detect potential cyber threats includes looking for cyber threats inside containers stored on the virtual disk.” As discussed above, Veslov teaches that the analyzed assets may be VMs or containers and that a VM itself may include multiple containers. *Supra*, §XI.C.6; EX1007, 1:34-40, 2:44-51. A POSA would have understood that such containers were typically stored on a virtual disk (e.g., a virtual disk of the associated VM), meaning the containers' data would be included in the previously discussed snapshot. *Supra*, §XI.B.3; EX1007, 1:34-40, 3:20-40; EX1002, ¶¶243-44 (citing EX1017, 17:28-18:12, 19:51-63, Fig. 3; EX1120, ¶¶41-44, Figs. 8-10;

EX1121, 15:1-44, Fig. 2A). Thus, regardless of whether the analysis involves directly analyzing snapshot data or indirectly analyzing snapshots via instantiated virtual assets (*supra*, §XI.B.5), Veselov teaches that the analysis includes “looking for cyber threats inside containers stored on the virtual disk.” EX1002, ¶244 (citing EX1007, 1:34-40, 3:61-67). At a minimum, this would have been obvious since identifying threats in containers was well-known and routine, including during security assessments of a VM that—as taught by Veselov—includes multiple containers. *Supra*, §X.A; EX1007, 1:34-40, 2:46-58; EX1002, ¶¶245-46 (citing EX1078, Abstract, ¶¶31-37).

Accordingly, claims 15 and 32 were obvious.

14. Claims 17 and 34

Claims 17 and 34 depend from claims 1 and 18, respectively, and further recite “wherein the virtual cloud asset is a server or a storage system in a cloud computing platform.” These limitations are recited in the alternative, meaning either is sufficient.

Veselov teaches both that the virtual cloud asset is a server and a storage system in a cloud computing platform. First, Veselov teaches that the virtual cloud asset may be a server. For example, Veselov teaches that the target VM may be a server (e.g., a web server). EX1007, 2:54-58, 9:5-8; *see also supra*, §XI.B.1 (target asset is virtual cloud asset); EX1002, ¶¶247-48. Veselov also teaches the

virtual cloud asset may be hosted on a server. EX1007, 5:23-50, 18:57-19:20, Figs. 1, 7. Second, Veselov teaches or suggests that the virtual cloud asset is a “storage system.” For example, Veselov explains that the target asset stores data and has storage capabilities such as virtual disk storage. *Id.*, 6:7-24, 11:21-33; *see also id.*, 2:22-35; EX1002, ¶249 (citing EX1012, 1:14-37; EX1026, 74).

Regardless of whether the asset is a server and/or a storage system, Veselov teaches or at least suggests that it is in a cloud computing platform, as previously discussed. *Supra*, §XI.B.1; EX1007, 2:64-3:19; EX1002, ¶¶250-51.

Accordingly, claims 17 and 34 were obvious.

XII. GROUND 2: CLAIMS 6 AND 23 WERE OBVIOUS OVER VESELOV, BASAVAPATNA, AND KAPOOR

This combination builds upon Ground 1 and further applies Kapoor’s teachings of reading PID files to determine the currently running processes. The combination does not impact the previous obviousness arguments for limitations beyond those discussed in §§XII.A-B. EX1002, ¶252.

A. Reasons to Combine Veselov, Basavapatna, and Kapoor.

The references are analogous art. Kapoor describes methods of performing a security assessment of a virtualized cloud environment. EX1111, 4:27-34 (datacenters comprise cloud-based resources), 15:54-67 (discussing the use of VMs); *supra*, §X.C. The references are thus in the same field as the ’798 patent

and reasonably pertinent to issues addressed by the '798 patent. EX1001, 1:12-14; EX1002, ¶¶253-54; *supra*, §XI.A.

Arriving at the claimed invention involved using known elements according to their known functions to achieve predictable results. *KSR*, 550 U.S. 398, 416-17 (2007). In addition to the known elements of Veselov and Basavapatna discussed above (*supra*, §XI.A), reading PID files to determine currently running processes, as taught by Kapoor, was well-known and predictable. EX1111, 7:21-24 (scanning PID files to evaluate running processes), 7:63-67, 37:35-37 (OSs use PIDs to identify running processes); EX1002, ¶255 (citing EX1112, 1-22; EX1113, 1-2). The further combination with Kapoor would merely involve using the well-known method of reading PID files to predictably perform its known function. EX1002, ¶255. These teachings are consistent with Ground 1 since Basavapatna generally teaches determining the running applications as part of its security assessment by using similar system configuration files, as discussed above. EX1008, ¶¶37, 60, 75-76, 78, 79, 82, 92, 104, 106, 112; EX1002, ¶255; *supra*, §XI.C.3.b.

A POSA would have also been motivated to combine Kapoor with Veselov, Basavapatna, and VMware SDK. As previously explained, Basavapatna teaches determining an asset's running applications as part of its risk assessment. *Supra*, §XII.D.5.b; EX1008, ¶¶37, 78-81, 83-85; *see also id.*, ¶¶60, 104-05. However, Basavapatna does not discuss what specific data is analyzed to determine the

running applications. Kapoor's techniques were a simple and effective way to make this determination using PID files. EX1111, 7:21-23, 7:63-67, 10:56-61, 36:5-26, 37:35-37; *see infra*, §XIII.B; EX1002, ¶256. A POSA would have recognized that determining running applications using PID files, as taught in Kapoor, was equally applicable and useful to the Veselov-Basavapatna combination. EX1002, ¶256 (citing EX1007, 1:22-40, 3:32-40, 10:55-67; EX1008, ¶37). A POSA therefore would have recognized the PID files discussed in Kapoor would have been present and readily analyzed in Veselov's snapshots. EX1002, ¶256 (citing EX1112, 4-5, 15). The combination is therefore a straightforward, routine implementation of Basavapatna's general teachings to determine if an application associated with a cyber threat is running. EX1002, ¶256.

A POSA would have had a reasonable expectation of success. The method of reading PID files to determine information on the running processes as discussed in Kapoor was well understood and routinely practiced. EX1002, ¶257 (citing EX1112, 1-22; EX1113, 1-2). As Dr. Stavrou explains, combining the references' teachings as discussed herein would not have presented meaningful technical challenges or produced unexpected results since a POSA also would have expected this technique to be readily accomplished in the context of the snapshot-based analysis. EX1002, ¶257 (citing EX1007, 3:61-67, 8:34-47, 9:28-32, 10:1-36,

11:47-67, 13:53-63, 16:56-17:9, 17:9-29 (scanning service reads OS, applications, and files/file system), 18:16-56).

B. Claims 6 and 23

Claims 6 and 23 depend from claims 3 and 20, respectively, and further recite “wherein scanning the parsed copy further comprises any one of: reading process identification number (PID) files; and checking if the at least access times of the PID files match against process descriptors.” These two limitations are recited in the alternative so either is sufficient. As explained below, the combination at least teaches “reading process identification number (PID) files,” as part of scanning the parsed copy.

As a preliminary matter, the combination of Veselov and Basavapatna teaches that “scanning the parsed copy” of a snapshot may include analyzing machine memory stored in the snapshot to deduce what applications are running, as previously discussed. *Supra*, §XII.C.3.b. As also discussed above, Veselov teaches identifying “files, software packages, application versions, and the like, for software installed on the target virtual machine 312.” EX1007, 18:16-44; *see also id.*, 9:55-67; *supra*, §§XI.B.5, XI.C.2.b; EX1002, ¶¶258-59. The combination of Veselov and Basavapatna thus teaches that the parsed copy of the snapshot may include the VM’s files/file system, and that the analysis of this parsed copy may

include analyzing machine memory in the snapshot to determine running applications. EX1002, ¶259.

While Basavapatna does not explicitly describe reading a system's PID files to determine whether a vulnerable application is running, Kapoor teaches or at least suggests determining the running applications in this manner. Kapoor teaches collecting information about a node—which may be a VM (EX1111 (Kapoor), 13:30-54, 14:20-44, 15:54-67)—including the processes running on it. EX1111, 6:41-7:67, 8:31-37, 14:14-19, Fig. 2. A POSA would have understood that many Unix-like systems (e.g., Linux) store a PID file for each running process. EX1002, ¶260 (citing EX1112, 1-2; EX1113, 2; EX1114, 3). Consistent with this understanding, Kapoor explains that PIDs “identify processes running at a given time.” EX1111, 37:28-37. Kapoor further teaches scanning within an OS's PID directory to read PID files to determine which processes are currently running. *Id.*, 7:10-31; *see also id.*, 7:63-67. It was therefore obvious for the scanning of the parsed copy of the snapshot, as taught by Veselov, to include reading PID files. *Supra*, §XII.A; EX1002, ¶¶261-62 (citing EX1111, 8:31-37, 14:17-19, claims 15, 36, Figs. 2, 3A-B; EX1112, 4-5, 15).

Claims 6 and 23 were therefore obvious.

XIII. GROUND 3: CLAIMS 11 AND 28 WERE OBVIOUS OVER VESELOV, BASAVAPATNA, AND ROTH

This combination builds upon Ground 1 and further applies Roth's teachings of event-triggered snapshots. The combination does not impact the previous obviousness arguments for limitations beyond those discussed in §§XIII.B. EX1002, ¶263.

A. Reasons to Combine Veselov, Basavapatna, and Roth

The references are analogous art. Roth describes snapshot techniques facilitating security analysis of virtual assets. EX1116, Abstract, claim 1; *supra*, §X.F. The references are thus in the same field as the '798 patent and reasonably pertinent to issues addressed by the '798 patent. EX1001, 1:12-14; EX1116, 1:28-44, 2:56-3:67, claim 1; EX1002, ¶¶264-65; *supra*, §XI.A.

Arriving at the claims involved using known elements according to their known functions to achieve predictable results. *KSR*, 550 U.S. at 416-17. In addition to the known elements discussed in Ground 1 (*supra*, §XI.A), Roth discusses known, predictable techniques for taking snapshots based on the detection of various events. EX1116, 2:67-3:67; EX1002, ¶266 (citing EX1119, 19:29-63). Consistent with Veselov's discussion of event-triggered snapshots (EX1007, 9:14-25, 14:41-53, 15:35-40, 23:48-56, claim 4), Roth further describes specific types of events that trigger the taking of snapshots for subsequent security analysis. EX1116, 2:42-47, 2:67-3:16. The combination thus uses Roth's known

techniques to achieve their known function (i.e., determining an appropriate moment to take a snapshot for subsequent security analysis) with predictable results. EX1002, ¶266.

A POSA would have been motivated to apply Roth’s teachings to efficiently provide snapshots with data relevant to security analysis. Veselov analyzes snapshots—which may be triggered by an event—to identify potential security problems (*supra*, §X.A; EX1007, 9:14-25, 14:41-53), and Basavapatna recognizes that various security problems have different priorities (*supra*, §XI.B). EX1002, ¶267. Roth further teaches taking snapshots when certain security-related events are detected (e.g., system intrusion and other security threats), which facilitates security analysis of the system’s state at times of particular concern. EX1116, 1:25-36, 2:42-3:16, 3:50-67; EX1002, ¶267. A POSA thus would have recognized that applying Roth’s teachings to the Veselov-Basavapatna combination would facilitate the efficient, targeted capture/analysis of snapshots following security-relevant events. EX1002, ¶267.

A POSA would have had a reasonable expectation of success. Event detection and snapshot generation, as taught by Roth, were routine and predictable in systems such as those addressed by Veselov and Basavapatna. *See* EX1007, 14:41-53, 15:35-40, 23:48-56, claim 4 (discussing event-triggered assessments); EX1002, ¶268 (citing EX1119, 19:29-63). The combination would have been

readily implemented since Veselov and Roth discuss analogous virtual assets/resources and snapshots. See EX1007, 9:14-17, 9:42-55; EX1116, 2:42-47, 3:64-67, 8:25-30, 13:37-14:58; *compare* EX1007, Fig. 7 with EX1116, Fig. 10. Moreover, Veselov already teaches event-triggered snapshots (EX1007, 9:14-25, 14:41-53; *supra*, §XI.C.9), so applying Roth’s teachings of triggering snapshots based on the detection of particular types of events would have been straightforward and predictable. EX1002, ¶268.

B. Claims 11 and 28

Roth teaches the additional limitation of claims 11 and 28, which recite “wherein taking the snapshot is performed based on detection of a network event or an abnormal event.” Like Veselov (*supra*, §XI.C.9), Roth teaches event-triggered snapshots that are used for security analysis. EX1116, Abstract, 2:20-41, 3:64-67, 7:1-21; EX1002, ¶¶269-70. For example, after an initial snapshot is taken, a VM may operate as normal “until the occurrence of the predetermined event (i.e., thereby triggering the event trigger), whereupon a second snapshot may be obtained.” EX1116, 3:50-67. Roth discusses several examples of detected events that trigger the taking of a snapshot. For example, the event may be an API request from a separate system or a system intrusion. *Id.*, 2:67-3:16. Each of these teaches or suggests a “network event” since such requests/intrusions typically occurred over a network. EX1116, 1:14-24, 8:30-47, 23:60-24:22, Figs. 2, 10;

EX1002, ¶¶40-44, 271 (citing EX1007, Fig. 1; EX1008, ¶59; EX1012, 10:16-25; EX1025, 16, 55-57; EX1026, 133; EX1033, 1-2). Other snapshot-triggering events include “security compromise of the virtual machine” or a “software or hardware error of the virtual machine.” EX1116, 2:67-3:16; *see also id.*, 17:35-43 (event may be execution error or detection of unauthorized software installation/modification). Each of these teaches or suggests an “abnormal event” since they indicate deviation from normal operation. EX1002, ¶¶272-73.

Claims 11 and 28 were therefore obvious.

XIV. GROUND 4: CLAIMS 14 AND 31 WERE OBVIOUS OVER VESELOV, BASAVAPATNA, AND CZARNY

This combination builds upon Ground 1 and further applies Czarny’s teachings of matching techniques such as cryptographic-hash comparison. The combination does not impact the previous obviousness arguments for limitations beyond those discussed in §XII.B. EX1002, ¶274.

A. Reasons to Combine Veselov, Basavapatna, and Czarny

The references are analogous art. Czarny describes detecting vulnerabilities in computer assets. EX1084, Abstract, 3:37-4:47, 9:17-10:42, Figs. 1, 7-10; *supra*, §D. The references are thus in the same field as the ’798 patent and reasonably pertinent to issues addressed by the ’798 patent. EX1001, 4:43-45; EX1002, ¶¶275-76; *supra*, §XI.A.

Arriving at the claims involved using known elements according to their known functions to achieve predictable results. EX1002, ¶277; *KSR*, 550 U.S. at 416-17. Like the elements of Veselov and Basavapatna discussed above, matching applications via hash comparison, as taught by Czarny, was well-known and predictable. EX1084, 3:37-4:47, 5:44-52, 7:31-67, 9:17-10:42, 11:16-12:18, Figs. 1, 7-10; *supra*, §§X.A, XI.C.12; EX1002, ¶277 (citing EX1047, 1-9; EX1053, ¶162; EX1068, 14:58-15:18, 4:29-39; EX1063, ¶¶75, 244-48; EX1076, 2:24-67). The combination merely involves using these well-known matching techniques to predictably perform their known functions. EX1002, ¶278. Czarny describes known matching techniques for performing the same types of vulnerability detection discussed in Veselov (e.g., CVE assessments) and Basavapatna. EX1007, 2:35-41, 7:4-22, 10:51-11:6; EX1008, ¶¶15, 51, 60; EX1084, 3:56-4:29, 6:25-38. The applied teachings would therefore perform their known functions of detecting vulnerable applications to achieve predictable results. EX1002, ¶279.

Additionally, a POSA would have been motivated to combine the previous references with Czarny to provide robust vulnerability detection and reporting. Czarny discusses advantages of hash comparisons. EX1084, 4:7-47 (provides “a more thorough, robust and flexible level of functionality” than matching based on software name/version), 5:44-6:24 (improves detection of custom/unofficial software). EX1002, ¶280. Hash-based matching was known to provide fast and

accurate matching against lists of known applications. EX1002, ¶280 (citing EX1076, 2:39-59, 6:40-58; EX1025, 23; EX1063, ¶¶244-48; EX1068, 4:29-39). A POSA therefore would have been motivated to use these common, effective matching techniques in order to assess risk and report results more accurately. EX1002, ¶280.

A POSA would have had a reasonable expectation of success. As explained above, matching via binary/hash comparison, as taught by Czarny, was a well-understood, predictable technique that was routinely used to perform the types of matching-based analysis described in Veselov and Basavapatna. EX1002, ¶281 (citing EX1063, ¶¶75, 244-48; EX1068, 4:29-39; EX1076, 2:24-67). Arriving at the claims by combining the references' teachings therefore would not have presented meaningful technical challenges or produced unexpected results. EX1002, ¶281.

B. Claims 14 and 31

The combination teaches claims 14 and 31, which recite “wherein analyzing the snapshot to detect potential cyber threat includes detecting a vulnerable application file by comparing a hash value of the vulnerable application file against a database of known application files.” Czarny detects vulnerabilities by matching an application file’s binary data against a database of known vulnerable applications (e.g., CVE database). EX1084, 6:25-38, 11:7-47, Figs. 1-5, 10; *supra*, §X.D.

Czarny describes several ways to perform this matching, including a hash-based comparison (i.e., comparing a hash value of the application file against a database of known vulnerable application files). EX1084, 4:7-24, 5:59-6:24, 9:28-64, claims 1, 6, Figs. 4-5; EX1002, ¶¶282-83. As explained above, this approach is consistent with Veselov’s teachings and conventional approaches for matching applications files. *Supra*, §XIV.A; EX1002, ¶¶284-85. The combination thus teaches that the snapshot analysis includes detecting a vulnerable application file via the recited hash-value comparison.

Claims 14 and 31 were therefore obvious.

XV. GROUND 5: CLAIM 13, 16, 30, AND 33 WERE OBVIOUS OVER VESELOV, BASAVAPATNA, AND CHARI

This combination builds upon Ground 1 and further applies Chari’s express disclosure of detecting sensitive data, including personally identifiable information (“PII”) and private keys. Applying these teachings does not impact the obviousness arguments for the limitations beyond those discussed below in §§XV.B-C. EX1002, ¶286.

A. Reasons to Combine Veselov, Basavapatna, and Chari

As an initial matter, the references are analogous art. Chari describes security assessments that include detecting sensitive data and remediating vulnerabilities in virtual images. EX1056, Abstract, ¶¶1-5, 14; *supra*, §X.E. The references are therefore in the same field as the ’798 patent and are reasonably

pertinent to issues addressed by the '798 patent. EX1001, 1:12-14; EX1002, ¶287; *supra*, §XI.A.

Arriving at the claims involved using known elements according to their known functions to achieve predictable results. *KSR*, 550 U.S. at 416-17. In addition to the known elements discussed above (*supra*, §XI.A), sensitive-data detection was well known and predictably used in security assessments such as those described in Veselov and Basavapatna. EX1056, ¶¶38-41, 64-66; EX1002, ¶¶34, 50, 54, 288. For example, Chari teaches detecting sensitive data in images/snapshots, which are the same or similar to the types of files analyzed in Veselov and Basavapatna. EX1056, ¶¶64, 73; EX1002, ¶288; *supra*, §XI.B.5. Detecting sensitive data, as taught by Chari, would therefore have its same predictable operation and function in the combination. EX1002, ¶288.

A POSA would have been motivated to apply Chari's teachings in order to provide a more complete identification of security risks. The presence of unsecured sensitive data was a common concern in security assessments. EX1002, ¶¶50, 54, 289. Veselov involves cloud systems, and Chari observes that residual sensitive data is of particular concern to cloud systems. EX1007, 2:27-3:19; EX1056, ¶¶59-62; EX1002, ¶289; *supra*, §XI.B.1. Chari mitigates this issue by scanning images/snapshots to detect and sanitize such data, which a POSA would have recognized was equally beneficial in the security assessments taught by

Veselov and Basavapatna. EX1056, ¶¶62-66; EX1002, ¶289. In other words, consistent with Veselov and Basavapatna's general teachings of detecting vulnerabilities based on the potential exposure of sensitive data (e.g., CVEs/CVSS), Chari expressly discloses detecting specific instances of sensitive data (e.g., social security numbers, account numbers, and encryption keys) in an image/snapshot. EX1056, Abstract, ¶¶14, 38, 76-80; EX1007, 2:35-41, 7:15-22; EX1008, ¶¶15, 21, 45, 55-56, 101, 106, 112; EX1002, ¶289.

A POSA would have had a reasonable expectation of success. Detecting sensitive data, as taught by Chari, was routinely and predictably performed as part of security assessments such as Veselov's and Hufsmith's. EX1002, ¶¶50, 54, 290. Applying these teachings is consistent with and readily implemented in combination with Veselov and Basavapatna. EX1007, 2:35-41, 14:13-22; EX1007, 2:35-41, 14:13-22; EX1008, ¶¶21, 45, 56; EX1002, ¶290. Thus, the combination would not have presented significant technical challenges or produced unexpected results. EX1002, ¶290.

B. Claims 13 and 30

Claims 13 and 30, which depend from claims 1 and 18, respectively, further recite "wherein analyzing the snapshot to detect potential cyber threats includes detecting cyber threats that do not represent vulnerabilities." The '798 patent identifies sensitive data not being encrypted, private keys, and system credentials

stored clearly as types of cyber threats that do not represent vulnerabilities.

EX1001, 6:19-26. As explained below Chari teaches detecting these types of cyber threats.

First, Chari teaches detecting sensitive data that is not encrypted. Chari teaches that a “labeler” may be used to detect and label sensitive data identified in the target VM. EX1056, Abstract, ¶¶37-38, 60, 76; *supra*, §XV.A. Chari also teaches or suggests that this data is not encrypted. EX1002, ¶¶291-92 (citing EX1056, ¶¶38 41, 85, claim 10).

Second, Chari teaches or suggests detecting private keys. Chari discloses that the detected sensitive data may include “cryptographic keys.” EX1056, ¶38; *see also id.*, ¶¶41, 60. A POSA would have understood this to encompass, or at least suggest, detecting private keys since private keys were a common type of cryptographic key stored by such systems. EX1002, ¶293 (citing EX1081, ¶¶14, 110-15, 117, 160-66).

Third, Chari teaches detecting system credentials stored clearly. For example, Chari teaches detecting “user credentials (i.e., passwords),” which is a type of system credential. EX1056, ¶60; *see also id.*, ¶61. As explained above, Chari explains or at least suggests detecting such data in an unencrypted form (i.e., system credentials “stored clearly”). *Id.*, ¶¶41, 85, claim 10; *see also id.*, ¶39, 77; EX1002, ¶294 (citing EX1018, 3-2 to 3-4; EX1025, 58; EX1081, ¶¶94-101).

For each of these examples of “detecting cyber threats that do not represent vulnerabilities,” it would have been obvious to detect such data as part of “analyzing the snapshot” in view of the references’ combined teachings. *Supra*, §XV.A. Claims 13 and 30 were therefore obvious. EX1002, ¶295.

C. Claims 16 and 33

Chari teaches the additional limitations of claims 16 and 33, which depend from claims 1 and 18, respectively, and further recite “wherein analyzing the snapshot to detect potential cyber threat includes looking for sensitive data including at least one of personal identifiable information or private keys.” As discussed above, Chari teaches detecting private keys, thus Chari teaches “looking for sensitive data including ... private keys” for the same reasons. *Supra*, §XV.B; EX1056, ¶¶38, 41, 60; EX1002, ¶296. Chari also teaches detecting sensitive data such as “credit card numbers, social security numbers, [and] bank account numbers,” which were well-known types of PII. EX1056, ¶38; *see also id*, ¶¶41, 60, 76; *supra*, §XV.B; EX1002, ¶266 (citing EX1105, 2). As previously explained, it would have been obvious to detect such data as part of “analyzing the snapshot” to detect cyber threats in view of the references’ combined teachings. *Supra*, §XV.A. EX1002, ¶¶296-99.

Claims 16 and 33 were therefore obvious.

XVI. CONCLUSION

The challenged claims are unpatentable. Wiz respectfully requests that IPR be instituted.

Dated: January 17, 2025

Respectfully submitted,

/Matthew A. Argenti /
Matthew A. Argenti, Lead Counsel
Reg. No. 61,836

CERTIFICATION UNDER 37 C.F.R. §42.24(d)

Under the provisions of 37 C.F.R. §42.24(d), the undersigned hereby certifies that the word count for the foregoing Petition for *Inter Partes* Review totals 13,930 which is less than the 14,000 allowed under 37 C.F.R. 42.24(a)(i). In accordance with 37 C.F.R. 42.24(a), this word count does not include any table of contents, table of authorities, mandatory notices under §42.8, certificate of service or word count, or appendix of exhibits or claim listing.

Dated: January 17, 2025

Respectfully submitted,

/Matthew A. Argenti /
Matthew A. Argenti, Lead Counsel
Reg. No. 61,836

PAYMENT OF FEES UNDER 37 C.F.R. §§42.15(A) AND 42.103

The required fees are submitted herewith. If any additional fees are due at any time during this proceeding, the Office is authorized to charge such fees to Deposit Account No. 23-2415.

LIST OF EXHIBITS

Exhibit No.	Description
1001	U.S. Patent No. 11,868,798 to Shua
1002	Declaration of Dr. Angelos Stavrou
1003	<i>Curriculum Vitae</i> of Dr. Angelos Stavrou
1004	File History of U.S. Patent Application No. 17/821,345 ('798 Pat.)
1005	Intentionally Left Blank
1006	Second Amended Complaint for Patent Infringement (Dkt. No. 15), <i>Orca security ltd. v. Wiz, Inc., No. 1:23cv-00758-GBW (DDE)</i>
1007	U.S. Patent No. 11,216,563 B1 to Veselov
1008	U.S. Publication No. 2013/0191919 A1, published July 25, 2013, to Basavapatna
1009	Wolf, "Virtualization: From the Desktop to the Enterprise" APRESS (2005)
1010	Waldspurger, "Memory Resource Management in VMware ESX Server," USENIX, Proceedings of the 5 th Symposium on Operating Systems Design and Implementation (Dec. 9-11, 2002)
1011	Goldberg, "Survey of Virtual Machine Research," COMPUTER, Vol. 7, No. 6 (June 1974)
1012	U.S. Patent No. 10,944,811 B2 to Doctor
1013	Wood, "Black-box and Gray-box Strategies for Virtual Machine Migration" NSDI '07: 4 th USENIX Symposium on Networked Systems Design & Implementation, USENIX
1014	Klemperer, "Efficient Hypervisor Based Malware Detection," Carnegie Mellon University (May 2015)

Exhibit No.	Description
1015	vSphere Web Services SDK Programming Guide, VMWare, Inc. (April 17, 2018)
1016	Souppaya, “Application Container Security Guide,” NIST Special Publication 800-190 (Sept. 2017) https://doi.org/10.6028/NIST.SP.800-190
1017	U.S. Patent No. 11,106,784 B2 to Rosendahl
1018	Scarfone, “Guide to Security for Full Virtualization Technologies: Recommendations of the National Inst. of Standards and Technology,” NIST Special Publication 800-125 (Jan. 2011)
1019	U.S. Patent No. 9,881,040 B2 to Rawat
1020	U.S. Patent No. 9,116,633 B2 to Sancheti
1021	“NIST Cloud Computing Standards Roadmap,” NIST Special Publication 500-291, Version 2 (July 2013).
1022	Roberts, “Public Cloud Service Definition,” VMware vCloud® Architecture Toolkit™ for Service Providers, VMware, Inc. (2018)
1023	ITL Bulletin, “Security Patches and the CVE Vulnerability Naming Scheme: Tools to Address Computer System Vulnerabilities,” NIST (Oct. 2002)
1024	ITL Bulletin, “Preventing and Handling Malware Incidents: How to Protect Information Technology Systems from Malicious Code and Software,” NIST (Dec. 2005)
1025	Cichonski, “Computer Security Incident Handling Guide: Recommendations of the National Institute of Standards and Technology,” NIST Special Publication 800-61, Rev. 2 (Aug. 2012) http://dx.doi.org/10.6028/NIST.SP.800-61r2
1026	ENISA Threat Landscape Report 2018, “15 top Cyberthreats and Trends,” V1.0 (Jan. 2019)

Exhibit No.	Description
1027	Verizon, “2019 Data Breach Investigations Report” (May 2019)
1028	Di Pietro, “Virtualization Technologies and Cloud Security: Advantages, Issues, and Perspectives,” Springer Nature Switzerland AG (2018) https://doi.org/10.1007/978-3-030-04834-1_9
1029	Singh, “Overview of Attacks on Cloud Computing,” Int’l. J. of Eng. And Innovative Technology (IJEIT), Vol. 1, Issue 4 (Apr. 2012)
1030	Bairwa, “Vulnerability Scanners: A Proactive Approach to Assess Web Application Security,” Int’l. J. on Computational Sciences & Applications (IJSCA), Vol. 4, No. 1 (Feb. 2014)
1031	Stone, “IT Asset Management,” NIST Special Publication 1800-5 (Sept. 2018) http://doi.org/10.6028/NIST.SP.1800-5
1032	EP Patent No. 3,360,071 B1 to Li
1033	Garfinkel, “A Virtual Machine Introspection Based Architecture for Intrusion Detection,” Computer Science Dept., Stanford University (2003)
1034	Tang, “VirtAV: An Agentless Runtime Antivirus System for Virtual Machines,” KSII Transactions on the Internet and Information Systems, Vol. 11, No. 11 (Nov. 2017) https://doi.org/10.3837/tiis.2017.11.026
1035	Lengyel, “Scalability, Fidelity and Stealth in the DRAKVUF Dynamic Malware Analysis System (2014) https://dl.acm.org/doi/pdf/10.1145/2664243.2664252
1036	U.S. Patent No. 8,613,080 B2 to Wysopal
1037	Bauman, “A Survey on Hypervisor-Based Monitoring: Approaches, Applications, and Evolutions,” ACM Computing Surveys, Vol. 48, No. 1, Article 10 (Aug. 2015)

Exhibit No.	Description
1038	Suneja, “Exploring VM Introspection: Techniques and Trade-offs,” ACM, VEE’15 (Mar. 2015) https://dl.acm.org/doi/abs/10.1145/2817817.2731196
1039	Payne, “Simplifying Virtual Machine Introspection Using LibVMI,” SANDIA Report-SAND2012-7818, Sandia National Labs. (Sep. 2012)
1040	U.S. Publication No. 2013/0111018 A1 to Ammons
1041	U.S. Patent No. 5,893,131 to Kornfeld
1042	U.S. Patent No. 7,975,305 B2 to Rubin
1043	Jiang, “Stealthy Malware Detection and Monitoring Through VMM-Based “Out-of-the-Box” Semantic View Reconstruction,” ACM Transactions on Information and System Security, Vol. 13, No. 2, Article 12 (Feb. 2010)
1044	U.S. Patent No. 9,141,805 B2 to Giakouminakis
1045	U.S. Publication No. 2019/0379700 A1 to Canzanese Jr.
1046	U.S. Patent No. 10,783,241 B2 to Crabtree
1047	Bourquin, “Binslayer: Accurate Comparison of Binary Executables,” 2nd ACM SIGPLAN Program Protection and Reverse Engineering Workshop (2013)
1048	U.S. Publication No. 2013/0247133 A1 to Price
1049	Huseinović, “Virtual machine memory forensics,” 2013 21st Telecommunications Forum Telfor (TELFOR) (2013)
1050	Intentionally Left Blank
1051	Using VMware Workstation Pro: VMware Workstation Pro 14.0, VMware, Inc. (December 17, 2017)
1052	VMware VirtualCenter User’s Manual version 1.4, VMware, Inc. (July 7, 2006)

Exhibit No.	Description
1053	U.S. Publication No.2010/0070725 A1 to Prahlad
1054	U.S. Patent No. 9,830,182 B2 to Ackley
1055	Virtual Disk Programming Guide: vSphere Storage APIs for Data Protection (VADP) 6.0 & Virtual Disk Development Kit (VDDK) 6.0, VMware, Inc. (March 25, 2016)
1056	U.S. Publication No. 2015/0033221 A1 to Chari
1057	Archived copy of “CWE-200: Information Exposure,” archived on January 26, 2019, by the Internet Archive’s “Wayback Machine” https://web.archive.org/web/20190126150012/https://cwe.mitre.org/data/definitions/200.html
1058	Archived copy of “CWE-311: Missing Encryption of Sensitive Data,” archived on January 17, 2019, by the Internet Archive’s “Wayback Machine” https://web.archive.org/web/20190117150049/https://cwe.mitre.org/data/definitions/311.html
1059	Archived copy of “CWE-359: Exposure of Private Information (‘Privacy Violation’),” archived on January 13, 2019, by the Internet Archive’s “Wayback Machine” https://web.archive.org/web/20190113052451/https://cwe.mitre.org/data/definitions/359.html
1060	Intentionally Left Blank
1061	Intentionally Left Blank
1062	Intentionally Left Blank
1063	U.S. Publication No. 2014/0245376 A1 to Hibbert
1064	U.S. Publication No. 2017/0180318 A1 to Lutas
1065	U.S. Patent No. 9,900,333 B2 to Thakar

Exhibit No.	Description
1066	Zhou, “Always Up-to-date—Scalable Offline Patching of VM Images in a Compute Cloud,” 26th annual computer security applications conference (2010)
1067	Jeswani, “ImageElves: Rapid and Reliable System Updates in the Cloud,” 2013 IEEE 33rd International Conference on Distributed Computing Systems (2013)
1068	U.S. Patent No. 8,984,478 B2 to Epstein
1069	U.S. Patent No. 11,061,682 B2 to Vorbach
1070	U.S. Publication No. 2013/0024940 A1 to Hutchins
1071	U.S. Patent No. 9,875,160 B2 to Bezbaruah
1072	U.S. Patent No. 10,536,471 B1 to Derbeko
1073	Ammons, “Virtual machine images as structured data: the Mirage image library,” 3rd USENIX conference on Hot topics in cloud computing (2011)
1074	VMware Infrastructure Architecture Overview White Paper, VMWare, Inc. (2006)
1075	U.S. Patent No. 9,692,778 B1 to Mohanty
1076	U.S. Patent No. 10,055,576 B2 to Milner
1077	U.S. Patent No. 9,824,216 B1 to Khalid
1078	U.S. Publication No. 2020/0097662 A1 to Hufsmith
1079	Carroll, “Challenges in Modern Digital Investigative Analysis,” published in Forensic Science and Forensic Evidence I, U.S. Att’y’s Bull. Vol. 65 No. 1 (January 2017)
1080	U.S. Patent No. 7,831,739 B2 to Lambeth
1081	U.S. Publication No. 2010/0169948 A1 to Budko

Exhibit No.	Description
1082	Intentionally Left Blank
1083	Intentionally Left Blank
1084	U.S. Patent No. 9,749,349 B2 to Czarny
1085	U.S. Publication No. 2017/0048266 A1 to Hovor
1086	Intentionally Left Blank
1087	Intentionally Left Blank
1088	Intentionally Left Blank
1089	Intentionally Left Blank
1090	Archived copy of “Vulnerability Details: CVE-2006-2923,” published by CVE Details and archived on July 29, 2014, by the Internet Archive’s “Wayback Machine” https://web.archive.org/web/20140729171850/https://www.cvedetails.com/cve/CVE-2006-2923/
1091	Archived copy of “Vulnerability Details: CVE-2007-0018,” published by CVE Details and archived on August 9, 2014, by the Internet Archive’s “Wayback Machine” https://web.archive.org/web/20140809144720/https://www.cvedetails.com/cve/CVE-2007-0018/
1092	Intentionally Left Blank
1093	Intentionally Left Blank
1094	Intentionally Left Blank
1095	U.S. Publication No. 2017/0372070 A1 to Burdett
1096	Intentionally Left Blank
1097	Intentionally Left Blank

Exhibit No.	Description
1098	Intentionally Left Blank
1099	Intentionally Left Blank
1100	Archived copy of “Vulnerability Details: CVE-2018-6954,” published by CVE Details and archived on December 11, 2018, by the Internet Archive’s “Wayback Machine” https://web.archive.org/web/20181211090742/http://www.cvedetails.com/cve/CVE-2018-6954
1101	Archived copy of “CVE-ID: CVE-2018-15687,” published by MITRE Corp. and archived on December 11, 2018, by the Internet Archive’s “Wayback Machine” https://web.archive.org/web/20181211082546/https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-15687
1102	Archived copy of “CVE-ID: CVE-2018-1312,” published by MITRE Corp. and archived on January 5, 2019, by the Internet Archive’s “Wayback Machine” https://web.archive.org/web/20190105194958/https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-1312
1103	Archived copy of “Vulnerability Details: CVE-2013-2765,” published by CVE Details and archived on July 20, 2014, by the Internet Archive’s “Wayback Machine” https://web.archive.org/web/20140720123338/https://www.cvedetails.com/cve/CVE-2013-2765
1104	Intentionally Left Blank
1105	Archived copy of “Defense Privacy, Civil Liberties, and Transparency Division, U.S. Department of Defense, FAQs,” archived on May 17, 2018, by the Internet Archive’s “Wayback Machine” at https://web.archive.org/web/20180517201830/http://dpcl.d.defense.gov/Privacy/About-the-Office/FAQs .
1106	Intentionally Left Blank
1107	Intentionally Left Blank

Exhibit No.	Description
1108	Intentionally Left Blank
1109	Intentionally Left Blank
1110	Intentionally Left Blank
1111	U.S. Patent No. 10,498,845 B2 to Kapoor
1112	Archived copy of “Linux Programmer’s Manual: proc–process information pseudo-filesytem,” published by Man7 and archived on December 15, 2018, by the Internet Archive’s “Wayback Machine” at https://web.archive.org/web/20181215072140/http://www.man7.org:80/linux/man-pages/man5/proc.5.html
1113	Archived copy of “PID Definition,” published by The Linux Information Project and archived on November 1, 2018, by the Internet Archive’s “Wayback Machine” at https://web.archive.org/web/20181101094536/https://www.linfo.org/pid.html
1114	Archived copy of “Unix/Linux–Processes Management,” published by TutorialsPoint and archived on December 17, 2018, by the Internet Archive’s “Wayback Machine” at https://web.archive.org/web/20181217170426/https://www.tutorialspoint.com/unix/unix-processes.htm
1115	Intentionally Left Blank
1116	U.S. Patent No. 9,524,389 B1 to Roth
1117	Archived copy of “20 Log Files that are located under /var/log Directory,” archived on October 13, 2018, by the Internet Archive’s “Wayback Machine” at https://web.archive.org/web/20181013220819/https://www.thegeekstuff.com/2011/08/linux-var-log-files/ .
1118	Archived copy of “Linux YUM Tutorial,” archived on November 10, 2017, by the Internet Archive’s “Wayback Machine” at

Exhibit No.	Description
	https://web.archive.org/web/20171110120546/https://linuxmoz.com/linux-yum-tutorial/ .
1119	U.S. Patent No. 11,093,629 B2 to Hiltgen
1120	U.S. Publication No. 2020/0026527 A1 to XUE
1121	U.S. Patent No. 10,579,355 B2 to Bo

CERTIFICATE OF SERVICE

Pursuant to 37 C.F.R. §§42.6(e) and 42.105, I certify that I caused to be served true and correct copies of the foregoing: **PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 11,868,798 and Exhibits 1001-1004, 1006-1049, 1051-1059, 1063-1081, 1084-1085, 1090-1091, 1095, 1100-1103, 1105, 1111-1114 and 1116-1121** by Federal Express Next Business Day Delivery on January 17, 2025, on the Patent Owner's correspondence address of record for the subject patent as follows:

Finnegan, Henderson, Farabow,
Garrett & Dunner LLP
Attn: Christopher C. Johns
901 New York Avenue, N.W.
Washington, DC 20001-4413

Dated: January 17, 2025

/Matthew A. Argenti/
Matthew A. Argenti, Lead Counsel
Reg. No. 61,836