

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)
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 IPR2024-00220
Patent No. 11,431,735

**PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 11,431,735**

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	MANDATORY NOTICES UNDER 37 C.F.R. § 42.8.....	1
III.	CERTIFICATIONS.....	2
IV.	IDENTIFICATION OF CHALLENGE; STATEMENT OF PRECISE RELIEF REQUESTED	2
V.	THE '735 PATENT.....	3
	A. Challenged Claims	4
	B. Prosecution History	5
VI.	NO BASIS EXISTS FOR DENIAL UNDER 35 U.S.C. §325(D)	6
VII.	LEVEL OF ORDINARY SKILL	7
VIII.	CLAIM CONSTRUCTION	7
	A. Determining/Determine a “Location” of a Snapshot	8
	B. “Analyzing the Snapshot”	9
IX.	BACKGROUND	10
	A. Cloud Computing, Virtualization, and Snapshots.....	10
	B. Cyber Security	12
X.	PRIOR ART.....	14
	A. Veselov (U.S. Patent. No. 11,216,563, EX1007).....	14
	B. Basavapatna (U.S. Patent No. 8,595,845, EX1008)	19
XI.	GROUND 1: CLAIMS 1-7, 9-17, AND 19 WERE OBVIOUS OVER VESELOV AND BASAVAPATNA	20
	A. Independent Claims 1, 10, and 11	21
	1. Independent Claim 1	21
	2. Independent Claim 10	41
	3. Independent Claim 11	43
	4. Reasons Why a Person of Ordinary Skill in the Art Would Have Combined Veselov and Basavapatna	46

B.	Dependent Claims	50
1.	Dependent Claims 2 and 12	50
2.	Dependent Claims 3 and 13	53
3.	Dependent Claims 4 and 14	55
4.	Dependent Claims 5 and 15	60
5.	Dependent Claims 6 and 16	65
6.	Dependent Claims 7 and 17	69
7.	Dependent Claims 9 and 19	72
XII.	CONCLUSION.....	73

I. INTRODUCTION

Petitioner Wiz, Inc. (“Wiz”) respectfully requests review of U.S. Patent No. 11,431,735 (“the ’735 patent”), currently assigned to Orca Security Ltd. (“Orca”). This petition demonstrates that claims 1-7, 9-17, and 19 are unpatentable.

The claims of the ’735 patent describe simple, well-known techniques for protecting virtual cloud assets, such as virtual machines (“VMs”), against cyber threats. The security system locates, accesses, and analyzes a “snapshot” of the asset’s virtual disk, then sends prioritized security alerts to a user.

This type of snapshot-based security analysis technique was already well known in the art, as demonstrated by the combination of Veselov and Basavapatna. Veselov teaches every aspect of the independent claims, except that it does not expressly state that its alerts are prioritized. However, alert prioritization was commonplace at the time, as shown for example by Basavapatna, which describes analogous security analysis in which the alert process is based on a determined priority. The remaining dependent claims merely describe other well-known aspects of security analysis and cloud systems.

Accordingly, Wiz respectfully requests institution of *inter partes* review.

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 '735 patent in *Orca Security Ltd. v. Wiz, Inc.*, No. 1-23-cv-00758 (DDE), filed and served on July 12, 2023.

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)

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

Post: WILSON SONSINI GOODRICH & ROSATI, 650 Page Mill Road,

Palo Alto, CA 94304

Tel.: 650-354-4154

Fax: 650-493-6811

III. CERTIFICATIONS

Wiz certifies that the '735 patent is available for IPR and that 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). *See also* EX1002, ¶¶2-5 (qualifications); EX1003 (CV). The claims are unpatentable under 35 U.S.C. §311 and AIA §6 based on at least the following ground:

Ground	Claims	Basis
1	1-7, 9-17, 19	§103(a): obviousness over Veselov (EX1007) ¹ and Basavapatna (EX1008) ²

V. THE '735 PATENT

The '735 patent issued from U.S. Application No. 16/585,967 (“the '967 application”), filed September 27, 2019. EX1001, Face. It claims priority only to Provisional Application No. 62/797,718, filed January 28, 2019. The '735 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 '735 patent describes techniques for securing virtual assets against cyber threats. EX1001, Abstract. The specification describes snapshot-based security analysis of virtual cloud assets, which was a well-known security approach at the time. This straightforward process includes determining the location of a snapshot of a virtual disk of an instantiated virtual cloud asset,

¹ U.S. Patent No. 11,216,563.

² U.S. Patent No. 8,595,845.

accessing and analyzing the snapshot to identify threats, and issuing prioritized alerts for the identified threats. *Id.*, 1:37-49, 5:4-24, 6:38-58; EX1002, ¶¶61-62.

A. Challenged Claims

The '735 patent includes 19 claims. Claims 1, 10, and 11 are independent. EX1002, ¶63. Claim 1 is representative for purposes of this petition and recites:

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

determining a location of 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;

accessing the snapshot of the virtual disk based on the determined location; analyzing the snapshot of the protected virtual cloud asset to detect potential cyber threats risking the protected virtual cloud asset; and

alerting detected potential cyber threats based on a determined priority.

Independent claim 10 is directed to a computer-readable medium, and independent claim 11 is directed to a system, but these claims essentially mirror claim 1. The dependent claims add other conventional aspects of cybersecurity and cloud computing, including risk-based threat prioritization and mitigation (claims 2 and 12), determining an allocated virtual disk (claims 3 and 13), determining a snapshot's location by querying a cloud management console (claims 4 and 14), snapshot parsing and scanning (claims 5 and 15), checking

configuration files, access times, system logs, or machine memory (claims 6 and 16), monitoring the activity an instantiated virtual cloud asset (claims 7 and 17), and basic types of cloud assets such as VMs and software containers (claims 9 and 19).

B. Prosecution History

The '967 application (which issued as the '735 patent) underwent a brief examination and never received a rejection under §102 or §103. The first office action provisionally rejected the claims based on statutory double patenting but indicated that the claims were otherwise allowable. EX1004, 79-80. The double-patenting rejection was based on claims in two related applications, one of which—U.S. Application No. 16/750,556 (“the '556 application”)—was later abandoned after all claims received art-based rejections from a different Examiner.³ *Id.*, 79-80; EX1005, 66-80 (rejections), 107-08 (abandonment); EX1002, ¶64. The '735 patent issued after the applicant amended two claims “for

³ While one of the references in the combination applied against the '556 claims—Derbeko (U.S Patent No. 10,536,471)—was cited by the Examiner of the '967 application, the applicant did not disclose the rejections or the other applied references to the Examiner.

clarity,” canceled two other claims, and filed a terminal disclaimer. EX1004, 115, 127. 139.

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

The same or substantially the same prior art or arguments described herein were not previously presented to the Office. 35 U.S.C. §325(d). The references applied here—Veselov and Basavapatna—were not disclosed to the Office or discussed by the Examiner, nor are they cumulative of references considered during prosecution. As noted above, the ’967 application never received an art-based rejection, and the Examiner did not identify any particular limitations as disclosed or missing in the handful of references identified as the “closest” art. *Supra*, §V.B. The Notice of Allowance simply copies the abstracts from those references and states that these references do not teach the claims as a whole. *See* EX1004, 139-41.

By contrast, the combination of Veselov and Basavapatna does teach the challenged claims as a whole. Veselov teaches every aspect of the independent claims except that it does not expressly describe the commonplace technique of alerting based on a determined “priority,” which is disclosed in Basavapatna. *Infra*, §XI. The claims therefore would not have issued if the Examiner had considered Veselov and Basavapatna. The Office also lacked additional evidence discussed herein, including the declaration provided by Wiz’s expert, Dr. Stavrou

(EX1002), demonstrating that the claims merely recite features already known in the art.

Accordingly, there is no basis to deny institution under §325(d).

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 virtual machines. 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. *See* EX1002, ¶¶2-5, 22. Section IX below summarizes the state of the art, including background knowledge that would have informed a POSA's understanding of the references' teachings applied herein.

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). Unless otherwise stated,

this petition applies the ordinary and customary meaning of the claim terms. *See also* EX1002, ¶65. The following limitations warrant discussion.

A. Determining/Determine a “Location” of a Snapshot

Each independent claim recites determining (or a system configured to determine) “a location of a snapshot” of a virtual disk of a protected virtual cloud asset. A POSA reading the claims in light of the specification would have understood that the recited “location” encompasses at least a virtual location and a non-virtual location.

A POSA would have understood that the ordinary and customary meaning of a “location” in this context broadly encompassed a virtual location and a non-virtual location. EX1002, ¶¶66-68; *see also id.*, ¶¶30 (data locations), 34 (snapshot locations).

The specification confirms this understanding. It states that the “management console 150 may be queried, by the security system 140, about as the **location (e.g., virtual address)** of the virtual disk 118-1 in the storage 117.” EX1001, 4:27-30 (emphasis added). This parenthetical makes it clear that the recited location at least encompasses a virtual address, and the “e.g.” indicates that the location is not limited to a virtual address. EX1002, ¶67. Indeed, snapshots of virtual assets were routinely stored in non-virtual storage and accessed by referencing non-virtual locations. *Id.* A POSA therefore would have interpreted the term “location” to encompass both virtual

and non-virtual locations. *Id.*, ¶¶67-68 (citing EX1009, 242, 246-57; EX1010, 3-4; EX1015, 56; EX1021, 8); *see also* EX1002, ¶¶30, 34 (discussing common locations associated with computer data generally and with snapshots specifically).

B. “Analyzing the Snapshot”

Each independent claim recites “analyzing the snapshot” (or a system configured to “analyze 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, ¶¶69-70. This understanding is confirmed by the specification. *See, e.g.*, EX1001, 5:36-39 (“[T]he security system 140 may be configured to match the application files, either directly (using binary comparison) or by computing a cryptographic hash against database of files in vulnerable applications.”), 6:4-11 (“analyze the machine memory, as reflected in the page file”), 6:34-37 (security system “computing a cryptographic hash of the sensitive areas in the virtual disk and checking for differences over time”), 6:54-58 (system can analyze logs “derived from the snapshot”); *see also* EX1002, ¶70.

In the related litigation (*supra*, §II), Orca appears to treat similar limitations in related patents as also encompassing analysis of a VM instantiated from a snapshot. For example, Orca alleges that the accused product satisfies “analyz[ing] the at least one snapshot,” as recited in claim 1 of U.S. Patent No. 11,693,685, because it

“analyzes [the] operating system, application layer, and data layer’ of virtual machines to provide full visibility into vulnerabilities across the cloud computing environment.” EX1006, 23, 57-58. For purposes of this IPR, Wiz also applies Orca’s interpretation. *See also* EX1002, ¶71.

Accordingly, the discussion below applies a construction of “analyzing the snapshot” encompassing both direct analysis of the snapshot data and analysis of a VM instantiated from the snapshot. EX1002, ¶72. Veselov describes both of these approaches. *Infra*, §§XI.A.1.d, XI.A.2.d, XI.A.3.f.


IX. BACKGROUND

A. Cloud Computing, Virtualization, and Snapshots

Cloud computing was well known long before 2019. It was commonly used to provide services over a network to end users, who were typically unaware of the underlying computer infrastructure running those services. EX1002, ¶¶23, 36-38; EX1021, 18-19. The physical infrastructure was often provided by data centers that included large collections of physical resources. EX1002, ¶40; EX1013, 229; EX1021, 18-19.

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, ¶39; EX1010, 2; EX1011, 35; EX1021, 19. In this paradigm, each physical server could emulate one or more virtualized

computer systems (“virtual machines” or “VMs”), each of which ran its own operating system and applications like a non-virtual system, as shown below:



EX1009, 505 (Fig. A-5); *see also* EX1002, ¶¶24-25; EX1010, 2; EX1012, 556; EX1013, 229. Virtual machines were commonly managed by a virtual machine manager (“VMM”)—sometimes called a hypervisor or virtual machine monitor—which, among other functions, allocated storage to guest VMs (e.g., a virtual disk or portion thereof) and provided tools for interacting with guest VMs. EX1002, ¶¶26-29; EX1009, 505; EX1015, 55-58, 62-66, 163; EX1016, ii; EX1017, 1:16-35; EX1021, 94-95; EX1022, 29. Virtualized resources—including virtual machines themselves, the virtual disks on which the virtual machines reside, and more granular virtual 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, ¶¶29-31, 34; EX1009, 2, 242, 246-57;

EX1010, 3-4; EX1012, 553; EX1014, 22; EX1015, 56, 124; EX1021, 8; EX1031, 1.

As early as 2005, virtualized systems employed backup techniques involving “snapshots,” which saved data from the VM’s memory, as well as changes to the VM’s virtual disks, allowing the VM to revert to a previous state. EX1002, ¶¶32-33; EX1009, 257; EX1015, 164-66; EX1018, 2-6. Snapshots were widely used before 2019 and were commonly managed with utilities such as VMware’s vSphere Snapshot Manager. EX1002, ¶¶34-35; EX1009, 221, 257-60; EX1015, 164-66; EX1019, Abstract; EX1020, Abstract, 21:42-22:58.

B. Cyber Security

The security of virtual and non-virtual computer systems had long been a concern of computer engineers and customers. Traditional security systems sought to make computer systems more secure by identifying security risks including threats (e.g., malicious code or unauthorized entrants) and vulnerabilities (e.g., exploitable hardware, software, or system configurations). EX1002, ¶¶41-42, 50-51; EX1023, 1-2; EX1024, 1. Security scans commonly searched for known risks—e.g., those in published lists such as the Common Vulnerabilities and Exposures (“CVE”)—and unknown risks that might be indicated, for example, by behavioral abnormalities. EX1002, ¶43; EX1023, 1-3; EX1025, 12, 21; EX1026, 9, 24-25; EX1044, 3:50-54, 5:28-32. Virtual machines were known to be subject to at least

the same security risks as nonvirtualized machines. EX1002, ¶¶44-45; EX1018, 3-1, ES-2, 4-1, 4-3; EX1021, 45-46; EX1027, 24-26; EX1028, 6-7; EX1029, 1.

A POSA would have been familiar with a variety of scanning techniques applied in both virtual and nonvirtual environments. 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 at the time. EX1002, ¶¶46-48, 50-52; EX1030, 116; EX1031, 22; EX1032, ¶¶27, 43; EX1033, 1-2; EX1034, 5643; EX1037, 2; EX1038, 1-2; EX1040, Abstract (passive monitoring of virtual server “without utilizing agents executing within the virtual server”), ¶¶8-12, 23, 43-45, 48, 52. 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, ¶48; EX1033, 1; EX1034, 2; EX1035, 389; EX1037, 5; EX1038, 1-2. For virtualized systems, security scans were often performed on a snapshot of the protected resource rather than directly on the resource itself. EX1002, ¶48; EX1014, 24-25, 55; EX1038, 2-3; EX1039, 9. Another common technique involved “parsing” snapshots or other data (e.g., adding metadata or otherwise reformatting the data) before scanning it. EX1002, ¶49; EX1018, 3-3; EX1035, 389; EX1036, 3:56-65; EX1039, 14; EX1041, 3:1-9; EX1042, 10:45-12:60; EX1043, 12:10.

Security systems typically alerted users of the identified risks. EX1002, ¶54. Given the large number of risks identified in a typical scan and the widely varying importance of such risks, alerts were commonly prioritized based on various factors. EX1002, ¶¶53, 55-58; EX1008, Abstract; EX1025, 32; EX1042, 2:10-27; EX1044, 6:9-18, 11:51-67; EX1045, ¶7; EX1046, 11:52-12:14. Cybersecurity systems also commonly suggested or provided remedial actions to mitigate the identified risks (e.g., installing software patches, establishing firewalls, or blocking data containing signatures of known threats). EX1002, ¶¶59-60; EX1025, 21, 24; EX1032, ¶33.

X. PRIOR ART


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

Veselov,⁴ a patent titled “Security Assessment of Virtual Computing Environment Using Logical Volume Image,” describes methods and systems for performing a security assessment of a target computing resource, such as a virtual machine or an instance of a virtual machine, includ[ing] a scanning service that facilitates duplication of all or a portion of the target computing resource, and then performs the security assessment on the duplicate computing resource.

⁴ Veselov was filed on May 19, 2017, and issued January 4, 2022. It is therefore prior art under 35 U.S.C. §102(a)(2).


EX1007, Abstract. This “duplicate” can be a snapshot of a resource such as a virtual disk allocated to a VM. *Id.*, Abstract, 3:20-40, 11:28-33, Figs. 2, 4, 6; EX1002, ¶73. As Veselov explains, performing the security assessment on a snapshot, rather than the target computing resource itself, “avoid[s] consuming processing time, processing power, and storage space of the target computing resource.” EX1007, Abstract. Veselov uses its security assessment to identify various types of threats and vulnerabilities. *Id.*, 1:41-49 (known security assessment services identify “vulnerabilities, bad configurations, and the like”), 2:35-41 (“The present disclosure contemplates implementation of any suitable security assessment, including 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.”); EX1002, ¶77.

Veselov notes that its techniques are particularly useful in cloud computing environments such as data centers. EX1007, 2:22-3:19; EX1002, ¶74. It also describes applying its techniques to distributed, virtualized systems commonly used for cloud computing:



EX1007, Fig. 1; *see also id.*, 1:5-50 (discussing virtualization in data centers); *supra*, §IX.A (discussing cloud computing and virtualization).

Veselov describes security analysis systems and techniques mirroring those of the '735 patent. For example, the system obtains (or obtains access to) a snapshot—a routine process that typically involved locating the snapshot and either obtaining a copy of the snapshot in a test environment using that location or accessing the snapshot at its location in that environment (*see, e.g., supra*, §§IX.A-B)—prepares the snapshot for scanning, performs the scan, and provides assessment results to a user:



EX1007, Fig. 2; *see also id.*, 8:52-10:36 (describing Fig. 2), Figs. 4, 6; EX1002, ¶¶75, 77.

Veselov describes a variety of analysis techniques. For example, it describes analysis using 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, ¶76. Veselov also


describes analysis techniques that do not rely on any agent, including agentless analysis of the assessment VM and analysis of the snapshot itself as a data file (without instantiating an assessment VM). *E.g.*, EX1007, 3:61-67 (agentless analysis of snapshot data), 16:67-17:3 (“In the implementations of FIGS. 5A-B, rather than recreating and instantiating a duplicate VM, the scanning service 510 is configured to analyze the volume image 508 as a data file.”), Figs. 3A (depicting assessment VM without agent), 5A-B (agentless analysis of snapshot as data file); *see also id.*, 14:13-22 (appropriately authorized scanning service “may perform any suitable security assessment” on assessment VM and its underlying resources); EX1002, ¶76; *supra*, §IX.B (agent-based and agentless security assessments both well-known).

Veselov describes performing the security assessment in the target resource’s computing environment and in a test environment. *E.g.*, EX1007, Figs. 3A, 5A (scanning service analyzes data in test environment), 3B, 5B (scanning service analyzes data in target resource’s environment). For each of these approaches, the scanning service can analyze the snapshot data as a data file and/or analyze an assessment VM instantiated from the snapshot. *E.g.*, *id.*, Figs. 3A-B (scanning service analyzes assessment VM), 5A-B (scanning service analyzes snapshot as data file).

B. Basavapatna (U.S. Patent No. 8,595,845, EX1008)

Basavapatna⁵ describes examples of the commonplace technique of prioritizing security alerts. EX1008, Abstract; *supra*, §IX.B (discussing alert prioritization); EX1002, ¶78. For example, Basavapatna discusses prioritizing alerts with risk metrics calculated from vulnerability scores. EX1008, Abstract, ¶¶138-40. These scores are determined by assessing, for example, the severity of a vulnerability, the probability that the vulnerability affects a particular asset, and the probability that a countermeasure would mitigate the risk to the asset posed by the vulnerability. *Id.*, claim 1. “Ranking assets, vulnerabilities, and/or threats according to aggregate risk scores allows a user to quickly identify which assets are most at risk, or which vulnerabilities and threats are most dangerous for a system.” *Id.*, ¶138. Basavapatna also describes alerting users regarding the identified, prioritized risks in a variety of ways, including reporting the ten greatest risks, only reporting items above a minimum risk threshold, and only reporting risks for a particular asset. *Id.*, ¶¶138-39, Fig. 6. The following figure depicts exemplary system components and various factors that can be evaluated as part of the security analysis and alerting process:

⁵ Basavapatna was filed on January 19, 2012, published July 25, 2013, and issued November 26, 2013. It is therefore prior art under §§102(a)(1)-(2).



EX1008, Fig. 2.

XI. GROUND 1: CLAIMS 1-7, 9-17, AND 19 WERE OBVIOUS OVER VESELOV AND BASAVAPATNA

The combination of Veselov and Basavapatna teaches or suggests all limitations of claims 1-7, 9-17, and 19 as a whole. EX1002, ¶¶79-81. Veselov discloses systems and methods that locate a snapshot of a virtual disk of a virtual machine, access or obtain a copy of the snapshot, analyze the snapshot to detect threats and vulnerabilities, and then report the results of the security scan as alerts to a user. *Id.*, ¶¶73-77, 80; *supra*, §X.A. While Veselov does not expressly describe alerting the detected threats and vulnerabilities “based on a determined

priority,” prioritizing security alerts in systems such as Veselov’s was commonplace. *Supra*, §IX.B. For example, Basavapatna teaches various ways of prioritizing alerts during security scans analogous to those in Veselov. EX1002, ¶¶78, 81; *supra*, §X.B. To the extent Veselov does not independently suggest prioritizing the alerts its techniques generate, a POSA reading Veselov and Basavapatna had reason to do so with a reasonable expectation of success. *Infra*, §XI.A.4.

A. Independent Claims 1, 10, and 11

Claims 1, 10, and 11 recite a method, a computer-readable medium, and a system with basic computer components, respectively, but otherwise recite analogous limitations.

1. Independent Claim 1

a. Claim 1 Preamble

The preamble of claim 1 recites “[a] method for securing virtual cloud assets in a cloud computing environment against cyber threats, comprising:” To the extent the preamble is limiting, Veselov teaches it. EX1002, ¶¶82-86.

First, Veselov teaches “securing” assets “against cyber threats” as recited. *Id.*, ¶¶83-85. Veselov secures such assets by performing a “security assessment” that identifies cyber threats, including those identified in “rules packages, such as Common Vulnerabilities and Exposures (CVEs), Center for Internet Security (CIS)

benchmarks, ‘best practice’ packages, static or runtime behavior analysis, host configuration assessments, and the like.” EX1007, 2:36-41 (emphasis added); *see also id.*, 4:1-8 (describing results of security assessment of a virtual machine), 7:1-5 (scanning service identifies events that “exposed or repaired a security vulnerability”); EX1002, ¶83. The assets are secured by, for example, generating assessment results and providing the results and corresponding remedial options to a user. EX1007, 10:17-36 (“[I]f vulnerabilities are identified in the assessment results, compar[e] the assessment results to a remediation framework to identify one or more actions the user can take to address the vulnerabilities, and provid[e] the identified actions to the user....”); EX1002, ¶83.

Second, Veselov teaches, or at a minimum suggests, that its methods secure “virtual cloud assets in a cloud computing environment,” as recited. EX1002, ¶¶83-85. Veselov describes securing virtual assets such as virtual machines. EX1007, Abstract, 2:48-51 (protected assets include “a virtual machine, a group of interconnected virtual machines, a container of a virtual machine, a software application, a service, a server, a hard disk drive or logical volume, etc.”); EX1002, ¶¶83-85. Veselov teaches that its methods are particularly useful in a “distributed computing environment (e.g., a **‘cloud’ computing environment**),” such as a data center, and describes embodiments securing virtual resources in such environments. EX1007, 2:42-67 (emphasis added); *see also id.*, 2:22-35 (methods “particularly

suitable for security assessments performed in a data center at any level of abstraction,” including “virtual machines and virtual environments implementing virtual machines”), 2:67-3:19 (discussing security issues in cloud environments operated by service provider running “distributed application”), 12:21-36 (describing performance impact of scanning a live virtual machine in a production environment), 19:30-37 (data stores may include devices “in any standard, distributed, virtual or clustered environment”), 19:60-64 (operations may be “performed collectively by multiple devices, which may form a distributed and/or virtual system”), 20:46-56 (describing embodiment in which environment is a “distributed and/or virtual computing environment”); EX1002, ¶¶84-85; *supra*, §IX.A (data centers were typical components of cloud computing environments). Thus, Veselov discloses methods for securing *virtual cloud assets in a cloud computing environment* against cyber threats.

Accordingly, Veselov teaches all aspects of the preamble of claim 1.

b. Element 1.1

Element 1.1 recites “determining a location of 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 this element. EX1002, ¶¶87-102.

Veselov teaches performing its security assessment on “a snapshot of at least one virtual disk of a protected virtual cloud asset,” as recited. As Veselov explains, “the scanning system obtains, or obtains access to, a state of the resource at a point in time (e.g., a ‘snapshot’) prior to, or in conjunction with, initiating the security assessment.” EX1007, 3:20-27 (emphasis added); *see also id.*, 2:42-63 (describing benefits of snapshot-based analysis), 4:8-14 (service may delete or store snapshot after analysis), 4:53-54 (service may “obtain and analyze snapshot data 146”), 6:25-30 (service allows “analysis of various snapshot data 146”); EX1002, ¶¶88, 94.

Veselov teaches analyzing a wide variety of snapshots, including a snapshot of a logical volume of a virtual machine (i.e., a “virtual disk of a protected virtual cloud asset”). *E.g.*, EX1007, Abstract (analyzed snapshot can be of VM’s logical volume), 3:29-40 (snapshot may be “a copy of the state of memory, 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 resource “includes without limitation one or more physical and/or virtual computing resources or segment(s) thereof, such as a virtual machine, virtual machine instance, ... logical volume, physical or virtual storage device, etc.”), 11:27-33 (logical volume “also known as ... ‘virtual disk drive’”); *see also id.*, 2:27-35 (assessment suited for VMs/logical volumes), 3:27-40 (target resource can be VM, and snapshot may be,

inter alia, “the state of any devices (virtual or physical) allocated to the resource” or image of “logical volume”), 6:7-13 (protected systems contain “virtual disk storage”), 11:5-67 (discussing types of snapshots), 16:15-26 (same), 18:16-26 (same), Figs. 3A-B; EX1002, ¶¶88-90. Veselov further teaches that these VMs are protected virtual *cloud* assets since it teaches securing such resources in cloud computing environments. *Supra*, §§ IX.A (cloud computing), XI.A.1.a (preamble); EX1007 2:22-63 (techniques “particularly suited for security assessments performed in a data center at any level of abstraction”), 2:64-3:19 (service provider operates infrastructure running “distributed application”), 19:30-64 (data store and other devices may be part of distributed system); EX1002, ¶¶91-92.

Veselov also teaches “determining a location of” such a snapshot in numerous ways. This limitation encompasses determining a virtual or non-virtual location of the snapshot (*supra*, §VIII.A), and Veselov teaches both. EX1002, ¶¶93-94. As explained below, Veselov broadly teaches determining several types of virtual and non-virtual locations of the snapshot, including the snapshot’s general location, more specific locations of the snapshot data within its environment, and locations in a test environment. *See also supra*, §VIII.A (discussing different types of well understood “locations”); EX1002, ¶¶30, 34 (same).

Veselov explains that the scanning service can obtain the snapshot (e.g., bringing the snapshot from the VM's host environment or storage service into a testing environment) or obtain *access* to the snapshot (e.g., remotely interacting with the snapshot in the VM's own computing environment). EX1007, 3:20-27 (scanning system "obtains, or obtains access to" snapshot), 4:49-54 (scanning service "may obtain and analyze snapshot data"), 5:18-22 (service provider may "provide the snapshot data 146 to the scanning service 110, or may store the snapshot data 146 or otherwise make the snapshot data 146 accessible by the scanning service 110"), 6:25-49 (scanning service may "transmit a command to virtualization layer 144 to provide snapshot data 146 (or to provide access to snapshot data 146) corresponding to the virtual machine 112"), 9:14-25, 15:29-34, 16:21-26, 18:16-21; *see also id.*, 3:41-51 (describing duplicate VM "within" or "outside of the original virtual computing environment"), Figs. 2, 4, 6 (depicting obtaining steps). Obtaining the snapshot is depicted, for example, in Figs. 3A and 5A, while obtaining *access* to the snapshot is depicted, for example, in Figs. 3B and 5B. *Supra*, §X.A.

With respect to obtaining *access* to the snapshot, Veselov teaches or suggests determining both general and specific locations of the snapshot. For example, discussing Fig. 1 (teachings that apply to the later embodiments, including those depicted in Figs. 2-7) Veselov states that the scanning service may

“determine the corresponding virtualization layer 144” for the VM before transmitting a command to provide snapshot data or to provide access to the snapshot data. *Id.*, 6:25-49; *see also* EX1002, ¶94 (citing 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) (teachings apply to later embodiments). A POSA would have viewed determining the target resource’s virtualization layer as teaching, or at least suggesting, determining a “location” of the snapshot as recited, since the snapshot is located there. EX1002, ¶97.

Veselov also expressly teaches accessing the snapshot “at a storage **location** of the snapshot data.” EX1007, 9:14-25 (emphasis added). This statement is presented in the discussion of Fig. 2, and these teachings apply to the other embodiments. *E.g., id.*, 8:60-62 (“The process 200 may be performed by any suitable system such as a scanning service or component thereof as described above in connection with FIG. 1.”); EX1002, ¶94. At a minimum, a POSA would have understood this storage “location” as encompassing the general location where the snapshot is stored. EX1002, ¶95. This “location” also encompasses a virtual location such as a virtual address, since Veselov describes snapshots located in virtual environments, and resources in virtual environments were commonly accessed via virtual locations such as virtual addresses. EX1002, ¶¶95-96; EX1007, Figs. 3A (snapshot, here volume image 308, initially located in virtual

computing environment 304A), 3B (snapshot located in virtual computing environment 304B), 5A (snapshot, here volume image 508, initially located in virtual computing environment 504A), 5B (snapshot located in virtual computing environment 504B); *see also* EX1007, 10:63-11:4 (VM “may be associated with a virtual network address”). A POSA would have understood that obtaining access to the snapshot via its location (e.g., its virtual address) involves determining that location in order to allow the snapshot to be accessed there. EX1002, ¶¶94-96; *infra*, §XI.A.1.c. At a minimum, determining the location (e.g., the virtual address) of the snapshot would have been obvious. EX1002, ¶96. For example, as Dr. Stavrou explains, it was common for remote systems to make queries to determine a resource’s location, such as a virtual address, in order to access the resource. *Id.*, ¶96; EX1009, 257; EX1015, 56, 164-66; EX1031, 1.

With respect to the other approach described in Veselov—obtaining the snapshot (e.g., bringing the snapshot into a testing environment)—Veselov teaches determining the location of the snapshot for similar reasons. The previously discussed disclosure of determining the virtualization layer applies to this approach as well. EX1007, 6:25-49; EX1002, ¶94; *see also* EX1007, Figs. 3A, 5A (depicting snapshot brought into test environment). Veselov also describes obtaining the snapshot “by accessing a data store (e.g., the snapshot storage service) containing the snapshot.” EX1007, 9:10-41; *see also id.*, 7:23-8:21 (discussing snapshot

storage service), Fig. 1 (depicting snapshot storage service). The general location of the data store itself and the reference to a particular snapshot within that data store are both “locations” of a snapshot, and a POSA would have understood that obtaining the snapshot from a snapshot storage service involves determining these locations (virtual or non-virtual) so that the snapshot can be obtained. EX1002, ¶¶96-98; *see also* EX1007, 7:43-8:21 (describing querying snapshot storage service), 19:21-37 (data store may operate “in any standard, distributed, virtual or clustered environment”). The disclosed process also teaches, or at least suggests, determining a specific location (e.g., virtual address) of the snapshot itself for the same reasons discussed above with respect to the process of obtaining access to the snapshot. *Id.*, ¶97.

Additionally, a POSA would have understood that bringing a snapshot into the test environment and interacting with the snapshot would both involve determining its location in the test environment. EX1002, ¶¶96, 99. For example, Fig. 3A shows that the snapshot is brought into testing environment 305, which includes its own virtualized environment managed by hypervisor 317. EX1007, 12:21-63, Fig. 3A; EX1002, ¶99. Figure 5A similarly shows the snapshot brought into testing environment 505, which a POSA would have understood to be either a virtualized environment or non-virtualized environment. EX1007, 17:9-29, Fig. 5A; EX1002, ¶99. A POSA would have understood that determining the location

of the snapshot in the test environment would occur when the snapshot is initially placed into the test environment (since storing the data at a location typically involved determining that location so that it can be stored there) and when the snapshot is analyzed during the assessment (e.g., reading the new location from memory to facilitate subsequent interaction with the snapshot). EX1002, ¶99. An example of both such determinations—involving both virtual and non-virtual locations—is discussed in the context of Fig. 3A, and a POSA would have understood that this process also applies to Fig. 5A. *Id.*; EX1007, 16:56-67 (Fig. 3A-B disclosures apply to Figs. 5A-B). Here, Veselov describes maintaining an index “associating data locations of the second logical volume 316 with the actual locations of the data within the physical storage device.” EX1007, 12:44-63. Then, during the scan, “the hypervisor 317 uses the index to identify which data is being read by the instance, and **where on the physical storage device the data is stored**,” then “may retrieve the data from the physical storage device and deliver the data to the requesting virtual machine instance.” *Id.* (emphasis added). At a minimum, determining the snapshot’s location in this context would have been obvious since determining data’s location (virtual or non-virtual, depending on the environment) was a typical precursor for interacting with the data. EX1002, ¶99.

Veselov thus teaches “determining a location of a snapshot of at least one virtual disk of a protected virtual cloud asset” in multiple ways. *Id.*, ¶100.

Lastly, Veselov teaches making this determination “wherein the virtual cloud asset is instantiated in the cloud computing environment.” *Id.*, ¶101.

Veselov notes that its systems perform “a security assessment of a target computing resource, such as a virtual machine or an **instance of a virtual machine**.” EX1007, Abstract (emphasis added). Veselov further explains that the servers may host VMs “instantiated” from volume images. *Id.*, 6:9-18; *see also id.*, 6:32-34 (user may request “a particular security assessment of one or more instances of a virtual machine”), 8:52-60 (target resource may include “a virtual machine, virtual machine instance, virtual computing environment implementing a plurality of virtual machine instances, logical volume, physical or virtual storage device, etc.”), 10:37-48 (discussing protected VM instances). Veselov notes that its techniques are beneficial because they allow VMs to remain running while their snapshots are being analyzed, which at a minimum suggests that the assessed VM remains instantiated in the cloud computing environment in some cases. EX1007, 2:42-63 (methods avoid “unavailability” of target resources); EX1002, ¶101.

Accordingly, Veselov teaches or suggests that the virtual cloud asset is “instantiated in the cloud computing environment,” as recited. EX1002, ¶101-02.

Accordingly, Veselov teaches all aspects of element 1.1.

c. Element 1.2

Veselov teaches element 1.2, which recites “accessing the snapshot of the virtual disk based on the determined location.” EX1002, ¶¶103-09.

As discussed above, Veselov teaches several ways of performing security assessments on snapshots, including obtaining access to a snapshot in its original computing environment and bringing the snapshot into a test environment (from the original computing environment or from a storage service). *Supra*, §XI.A.1.b; EX1002, ¶¶94, 104. These approaches teach or suggest determining a location of the snapshot in numerous ways, including determining the virtual or non-virtual location of the snapshot itself (within the original computing environment or storage service), determining the location of the virtualization layer of the original computing environment or the location of the data store (e.g., a snapshot storage service), and determining the location of the snapshot within the testing environment. *Supra*, § XI.A.1.b. For each of these “locations,” Veselov further teaches “accessing” the snapshot “based on the determined location,” as recited.

With respect to determining the snapshot’s location in the original computing environment or a snapshot storage service, Veselov teaches that the scanning system obtains “access to the snapshot data, such as at a storage location of the snapshot data.” EX1007, 9:14-25 (emphasis added); *see also id.*, 3:20-27 (scanning service obtains “access” to snapshot), 4:53-55 (service obtains

and analyzes snapshot data), 5:18-22 (VM's computing resource provider makes snapshot data "accessible" to scanning service), 6:38-46 (scanning service may transmit command to host virtualization layer to "provide access to snapshot data" (emphasis added)), 7:34-61 (scanning service accessing snapshot data by querying snapshot storage service), 8:16-21 (scanning service "may obtain the snapshot data from the snapshot storage service 146"), 11:34-36 (scanning service obtains snapshot data "as described above"), 22:60-23:3 (computing resource provider "make[s] the first data accessible" to scanning system), Figs. 2, 4, 6 (depicting obtaining step), claim 1 (processor configured to "obtain access to a first volume image"); EX1002, ¶¶105-07. At a minimum, in such embodiments it would have been obvious for the scanning service to access the snapshot based on its determined location in the virtualized computer environment, since it was well known that resources in such environments were typically accessed by referencing their virtual locations. EX1002, ¶106.

With respect to determining the location of the remote virtualization layer or storage service where the snapshot is stored, Veselov also teaches accessing the snapshot based on this determined location. *Id.*, ¶¶105-07. For example, the snapshot may be stored in a "data store or data storage service (e.g., snapshot storage service 108) accessible by the scanning service 110." EX1007, 5:9-22 (emphasis added)); *see also id.*, 6:25-38 (user API call may include "information

corresponding to ... the on-demand storage service 108, and/or the particular virtual machine 112”), 7:34-37 (“The snapshot storage service 108 may ... make the snapshot data 146 **accessible** to one or more other computer systems, such as the scanning service 110.” (emphasis added)), 8:16-34 (“data store 130 accessible by the scanning service 110”). Likewise, Veselov explains that the scanning service may obtain the snapshot data by “**accessing a data store** (e.g., the snapshot storage service) **containing the snapshot**.” EX1007, 9:14-25 (emphasis added). At a minimum, it would have been obvious for the scanning service to access the snapshot based on the determined location of the virtualization layer or storage service, since the disclosed process involves communicating with such systems. EX1002, ¶106; *see, e.g.*, EX1007, 6:38-46 (scanning service may “transmit a command to the virtualization layer 144”).

With respect to determining the location of the snapshot within the test environment, Veselov teaches accessing the snapshot based on the determined location for similar reasons. EX1002, ¶108. For example, Veselov explains that the hypervisor in the test environment may “read the volume image 308” (e.g., to instantiate a duplicate VM in the test environment). EX1007, 12:37-43; *see also id.*, 12:51-63 (hypervisor 317 in virtualized test environment “may create and/or maintain an index associating data locations of the second logical volume 316 with the actual locations of the data within the physical storage device”), Figs. 3A, 5A

(depicting snapshot in test environments). Likewise, Veselov describes the scanning service “mount[ing] the volume image 508 in the new logical volume 516 so that the scanning service 510 can access the volume image 508....” *Id.*, 17:9-29. At a minimum, accessing the snapshot based on the determined location in this context would have been obvious since it was well known that analyzing data in a test environment typically involved accessing that resource at its location (virtual or non-virtual, depending on the environment). EX1002, ¶108.

Accordingly, Veselov teaches all aspects of element 1.2.

d. Element 1.3

Element 1.3 recites “analyzing the snapshot of the protected virtual cloud asset to detect potential cyber threats risking the protected virtual cloud asset.” Veselov teaches this element. EX1002, ¶¶110-16.

Regardless of the technique used to locate and access the snapshot (*supra*, §§XI.A.1.b-c), Veselov teaches “analyzing the snapshot of the protected virtual cloud asset” using its scanning service per the construction discussed in §VIII.B. EX1002, ¶¶111-14. An exemplary scanning service 110 is shown in Figure 1:




FIG. 1

EX1007, Fig. 1 (scanning service 110 analyzing snapshot data 146 and producing assessment results 132); *see also* Figs. 3A-B (depicting scanning service 310), 5A-B (depicting scanning service 510). As previously discussed, these teachings apply to later embodiments. EX1002, ¶94. 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); *see also id.*, Figs. 2 (depicting “Perform Assessment” step), 4 (same), 6 (same).

Veselov discloses several ways of analyzing the snapshot. For example, Veselov explains that the scanning service may “determine whether to reconstitute an active duplicate virtual machine instance or to analyze the snapshot as a data file.” EX1007, 8:40-42; *see also id.*, 8:30-45 (service determining which type of analysis to perform on snapshot), 10:1-16 (describing types of analyses). In other words, Veselov teaches (1) instantiating a duplicate of the protected virtual machine from the snapshot and analyzing the duplicate and (2) mounting the snapshot as a data file and assessing the snapshot at the file system level. EX1002, ¶¶112-14. Each of these approaches satisfies the previously discussed construction of element 1.3. *See supra*, §VIII.B.

In the first approach, the snapshot is analyzed by using it to launch an active VM instance. *See, e.g.*, EX1007, 3:54-61 (service launching an active duplicate instance from snapshot data), 12:1-20 (same); EX1002, ¶113. “[T]he system may then perform the security assessment on the active virtual machine instance in a static or runtime environment.” EX1007, 3:54-61. Veselov describes instantiating a duplicate VM from the snapshot in the original environment and in a test environment. *Id.*, 3:44-54; *see also id.*, Figs. 3A (assessment VM 312A in original computing environment), 3B (assessment VM 312B in testing environment). This analysis may be performed with or without an agent installed in the assessment VM. *See, e.g., id.*, Fig. 3A (no agent on assessment VM), 3B (security agent 330

installed on assessment VM); *see also id.*, 14:13-22 (scanning service “may perform any suitable security assessment” on assessment VM and underlying resources); EX1002, ¶113; *supra*, §X.A (Veselov summary).

In the second approach, the scanning service analyzes the snapshot data itself rather than instantiating a VM from the snapshot. EX1002, ¶114. For example, Veselov explains that the scanning system can “mount the logical volume (or transform the snapshot into a data file by identifying the file system of the original logical volume and then mount the snapshot in the duplicate logical volume).” EX1007, 3:61-67 (service mounting and scanning snapshot as a hard drive), 17:10-29 (service mounts volume image and reads files). The scanning system “may then perform the security assessment on the data at the file system level.” *Id.*, 3:61-67; *see also id.*, 8:30-45 (scanning service may “analyze the snapshot as a data file”), 17:10-29 (scanning service can mount volume image to access it “as if it were a disk drive”). This type of security analysis is not performed by an agent installed on assessment VM (i.e., it is agentless), since the security service analyzes the snapshot data directly without instantiating an assessment VM. EX1002, ¶114.

Finally, Veselov analyzes snapshots “to detect potential cyber threats risking the protected virtual cloud asset,” as recited. EX1002, ¶115; *see also supra*, §XI.A.1.a (references teach preamble, including securing virtual cloud assets

against “cyber threats”). Veselov embodiments may use “any suitable security assessment,” including assessments defined by rules packages defining particular threats (e.g., CVEs), in order to “identify vulnerabilities, bad configurations, and the like.” EX1007, 2:22-41; *see also id.*, 1:46-49 (service analyzes behavior of protected resources to identify vulnerabilities and bad configurations), 4:57-59 (assessment identifies “security risk information”), 5:9-11 (snapshot analysis produces assessment results). The identified threats are those risking the target VM (i.e., the protected virtual cloud asset). *E.g., id.*, 4:1-4 (results from snapshot-based analysis comparable to those from analysis of original VM instance); EX1002, ¶115.

Accordingly, Veselov teaches all aspects of element 1.3.

e. Element 1.4

Element 1.4 recites “alerting detected potential cyber threats based on a determined priority.” The combination of Veselov and Basavapatna teaches this element. EX1002, ¶¶117-22; *see also infra*, §XI.A.4 (reasons to combine).

Veselov teaches “alerting detected potential cyber threats,” as recited. Veselov explains that the assessment results describe the identified risks. *E.g.*, EX1007, 10:17-36 (service produces “assessment results that describe ‘findings’ relative to the assessed properties” and identifies remedial actions), 20:57-21:13 (system produces “assessment result describing a security vulnerability”), 23:19-47

(same); *see also id.*, 5:9-13 (scanning service produces assessment results), 8:19-21 (same), claims 1, 16-17 (generating assessment results); EX1002, ¶¶118-20. Veselov further teaches providing these assessment results to the user (i.e., “alerting” the detected potential cyber threats). EX1007, 10:30-36 (system “provid[es] the assessment results in a user-readable format to a user interface to be displayed on the user device” and identifies remedial actions), 18:53-56 (service performs actions on assessment results, “such as storing or delivering the results”); *see also id.*, Figs. 2, 4, 6 (depicting step of providing assessment results).

Although Veselov does not expressly state that its alerts are “based on a determined priority,” prioritization of cybersecurity alerts was commonplace well before 2019, as shown, for example, in Basavapatna. EX1002, ¶121; *see also supra*, §IX.B (alert prioritization well known). Basavapatna discusses the usefulness of prioritization, noting that “[r]anking assets, vulnerabilities, and/or threats according to aggregate risk scores allows a user to quickly identify which assets are most at risk, or which vulnerabilities and threats are most dangerous for a system.” EX1008, ¶138. Basavapatna also discloses multiple algorithms for prioritizing alerts, based on, for example, the “risk associated with the vulnerability relative to other vulnerabilities” and the “criticality of the particular asset.” *Id.*, ¶¶14-15; *see also infra*, §XI.B.1.a (element 2.1). Basavapatna also describes displaying the prioritized alerts to the user. *E.g.*, EX1008, ¶¶138 (system “can list

all assets, vulnerabilities, or threats, sorted by aggregate risk metric, or can list a top number, e.g., top ten, of the assets, vulnerabilities, or threats”), 139 (user interface can list threatened assets sorted by risk level), 140 (system may limit alerts to those rising above specified risk threshold). Basavapatna thus teaches “alerting detected potential cyber threats based on a determined priority.”

EX1002, ¶121. A POSA had reason to combine these teachings with those of Veselov as discussed below. *Infra*, §XI.A.4.

Accordingly, the combination of Veselov and Basavapatna teaches all aspects of element 1.4.

2. Independent Claim 10

Independent claim 10 uses language identical to claim 1 with the exception that its preamble recites a computer-readable medium instead of a method (as in claim 1). The combination of Veselov and Basavapatna teaches claim 10 as a whole for substantially the same reasons as claim 1. EX1002, ¶123.

a. Claim 10 Preamble

The preamble of claim 10 recites “[a] non-transitory computer readable medium having stored thereon instructions for causing a processing circuitry to execute a process for securing virtual cloud assets in a cloud computing environment against cyber threats, the process comprising:” To the extent the preamble is limiting, Veselov teaches it. EX1002, ¶¶124-27.

Veselov teaches that its system includes a “non-transitory computer readable medium having stored thereon instructions for causing a processing circuitry to execute” its process. EX1007, 20:57-60 (system includes “memory and a processor, the memory including computer-executable instructions”), 20:60-23:18 (scanning system’s actions caused by executing instructions), claim 1; EX1002, ¶126. That process is “for securing virtual cloud assets in a cloud computing environment against cyber threats,” as explained above. *Supra*, §XI.A.1.a.

Accordingly, Veselov teaches the preamble of claim 10.

b. Element 10.1

Like element 1.1, element 10.1 recites “determining a location of 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 this limitation as explained above. *Supra*, § XI.A.1.b; EX1002, ¶128.

c. Element 10.2

Like element 1.2, element 10.2 recites “accessing the snapshot of the virtual disk based on the determined location.” Veselov teaches this limitation as explained above. *Supra*, § XI.A.1.c; EX1002, ¶129.

d. Element 10.3

Like element 1.3, element 10.3 recites “analyzing the snapshot of the protected virtual cloud asset to detect potential cyber threats risking the protected

virtual cloud asset.” Veselov teaches this limitation as explained above. *Supra*, § XI.A.1.d; EX1002, ¶130.

e. Element 10.4

Like element 1.4, element 10.4 recites “alerting detected potential cyber threats based on a determined priority.” The combination of Veselov and Basavapatna teaches this limitation as explained above. *Infra*, § XI.A.4 (reasons to combine); *Supra*, § XI.A.1.e (element 1.4); EX1002, ¶131.

3. Independent Claim 11

Independent claim 11 uses language identical to claim 1 with the exception that its preamble and first two elements recite a system of conventional computer programming components instead of a method (as in claim 1). The combination of Veselov and Basavapatna teaches claim 11 as a whole for substantially the same reasons as claim 1. EX1002, ¶132.

a. Claim 11 Preamble

The preamble of claim 11 recites “[a] system for securing virtual cloud assets in a cloud computing environment against cyber threats, comprising:” To the extent the preamble is limiting, Veselov teaches it. EX1002, ¶¶133-35. The preamble of claim 11 is identical to the preamble of claim 1 except that it recites a “system” instead of a “method.” Veselov discloses both systems and methods, and a POSA would have recognized that Veselov teaches a system performing its

techniques. EX1002, ¶134. Veselov teaches all aspects of the preamble of claim 11, as explained above. *Supra*, § XI.A.1.a; EX1002, ¶¶133-35.

b. Element 11.i

Veselov teaches element 11.i, which recites “a processing circuitry.” EX1002, ¶¶136-38; *see also supra*, § XI.A.2.a (claim 10 preamble recites “a processing circuitry”). Veselov explains that the “physical hardware implementing” the invention “may include one or more server computers,” then elaborates that “[a] server computer may be equipped with any needed processing capability including one or more processors” EX1007, 5:23-37. Veselov also describes processing circuitry in its claim 1, which recites a system including “a processor.”

Thus, Veselov discloses all aspects of element 11.i.

c. Element 11.ii

Element 11.ii recites “a memory, the memory containing instructions that, when executed by the processing circuitry, configure the system to” Veselov explains that its system includes “a memory and a processor, the memory including computer-executable instructions that upon execution cause the system to” perform the previously discussed actions. EX1007, 20:57-23:18; *see also id.*, claim 1 (same); EX1002, ¶¶139-41.

Thus, Veselov discloses all aspects of element 11.ii.

d. Element 11.1

Like element 1.1, element 11.1 recites “determine a location of 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 this limitation as explained above. *Supra*, § XI.A.1.b; EX1002, ¶142.

e. Element 11.2

Like element 1.2, element 11.2 recites “access the snapshot of the virtual disk based on the determined location.” Veselov teaches this limitation as explained above. *Supra*, § XI.A.1.c; EX1002, ¶143.

f. Element 11.3

Like element 1.3, element 11.3 recites “analyze the snapshot of the protected virtual cloud asset to detect potential cyber threats risking the protected virtual cloud asset.” Veselov teaches this limitation as explained above. *Supra*, § XI.A.1.d; EX1002, ¶144.

g. Element 11.4

Like element 1.4, element 11.4 recites “alert detected potential cyber threats based on a determined priority.” The combination of Veselov and Basavapatna teaches this limitation as explained above. *Infra*, § XI.A.4 (reasons to combine); *Supra*, § XI.A.1.e (element 1.4); EX1002, ¶145.

4. Reasons Why a Person of Ordinary Skill in the Art Would Have Combined Veselov and Basavapatna

A POSA had reason to combine the previously discussed teachings of Veselov and Basavapatna to arrive at the claimed inventions. EX1002, ¶¶146-53.

While not dispositive of the obviousness inquiry, Veselov and Basavapatna describe systems and methods that are analogous to each other and to the '735 patent. EX1002, ¶147. Veselov describes detecting and alerting cyber threats involving virtual machines, including those in cloud environments. *See, e.g.*, EX1007, 4:5-8 (“The scanning system may store the assessment results, or provide them to an owner or administrator of the virtual machine in order to evaluate the vulnerabilities of the target resources.”); *see also id.*, Abstract (describing security assessment of virtual machines), 8:46-50 (describing generating of assessment results and storage or delivery of results to user); *supra*, §X.A (Veselov summary). Similarly, Basavapatna “relates in general to the field of computer security assessment and, more particularly, to calculating risk metrics for assets in a system of computing assets.” EX1008, ¶1; *see also id.*, ¶12 (describing standardized vulnerability score and relative risk levels); *supra*, §X.B (Basavapatna summary). As Dr. Stavrou explains, both references are in the same field as the '735 patent and are pertinent to the problem addressed by the '735 patent’s disclosure. EX1002, ¶147; *see also* EX1001, 1:10-12 (“This disclosure relates generally to

cyber-security systems and, more specifically, to techniques for securing virtual machines.”).

Moreover, arriving at the claimed invention involved nothing more than the use of known elements according to their known functions to achieve predictable results. EX1002, ¶148; *See KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 416-17 (2007). Performing a security assessment on a snapshot rather than the target computing resource itself, as taught by Veselov, was both well-known and predictable. *Supra*, §IX.B; EX1002, ¶148 (citing EX1014, 24-25). Prioritizing alerts generated by cybersecurity assessments, as taught in Basavapatna, was similarly common and predictable. *Supra*, §IX.B; EX1002, ¶148 (citing EX1044, 6:9-18, 11:51-67). These teachings were easily implemented and equally useful in Veselov’s similar context, particularly since virtual and non-virtual machines have analogous components, functionality, and security concerns. EX1002, ¶148; *supra*, §IX.B.

A POSA also would have been motivated to combine Veselov and Basavapatna for multiple reasons. EX1002, ¶¶149-53; EX1046, 11:52-12:14.

First, Basavapatna’s prioritization techniques were a natural fit and directly applicable to Veselov’s teachings regarding snapshot-based security assessments. Veselov provides a high-level description of its alerting process, while Basavapatna presents a more detailed discussion of the common technique of

optimizing alerts generated by a cybersecurity analysis. EX1002, ¶150; *Supra*, §§X.A-B. A POSA would have recognized that, while Veselov teaches snapshot-based analysis of VMs, the assessment results it produces would have been comparable to those produced by other types of security assessments (on virtual and non-virtual systems alike), meaning the application of familiar, generalized alert-prioritization schemes, such as those taught by Basavapatna, would have been straightforward and equally beneficial in Veselov’s context. EX1002, ¶150; *Supra*, §IX.B.

Moreover, a POSA would have been motivated to combine the teachings of Basavapatna with those of Veselov to obtain the known benefits of prioritized alerts. EX1002, ¶¶151-52. It was well known, and expressly disclosed in Basavapatna, that prioritizing alerts was beneficial for security assessments in several ways, including rapid identification of which assets are most at risk and which threats and vulnerabilities are the most dangerous. EX1008, ¶138 (“Ranking assets, vulnerabilities, and/or threats according to aggregate risk scores allows a user to quickly identify which assets are most at risk, or which vulnerabilities and threats are most dangerous for a system.”); EX1002, ¶152 (citing EX1045, ¶7); *supra*, §IX.B (known benefits of prioritized alerts). Given the comparable security concerns in virtual and non-virtual computer systems, a POSA would have expected to similarly achieve those benefits by alerting based on

determined priority when identifying cyber threats during snapshot-based analysis as taught by Veselov. EX1002, ¶¶151-52; *supra*, §IX.B (similar security concerns).

Finally, a POSA would have had a reasonable expectation of success in combining the above teachings of Veselov and Basavapatna to arrive at the claimed invention. EX1002, ¶153; *supra*, §IX.B. The snapshot-based techniques disclosed in Veselov were well understood and routinely practiced. EX1002, ¶153; *supra*, §IX.B. Alert prioritization was similarly well-known and commonplace. EX1002, ¶153 (citing EX1042, 2:10-27); *supra*, §IX.B. As Dr. Stavrou explains, a POSA would have been familiar with the types of techniques discussed in Veselov and Basavapatna—including determining snapshot locations, accessing snapshots based such locations, analyzing snapshots to identify cyber threats, and alerting such threats based on determined priorities—and combining the references’ teachings as discussed above would not have presented meaningful technical challenges or produced unexpected results. EX1002, ¶153; *supra*, §§IX.A-B.

Accordingly, independent claims 1 and 10-11 as a whole were obvious.

B. Dependent Claims

1. Dependent Claims 2 and 12

Claims 2 and 12 depend from independent claims 1 and 11, respectively, and recite further limitations taught by the combination of Veselov and Basavapatna. EX1002, ¶¶154-65.

a. Elements 2.1 and 12.1

Elements 2.1 and 12.1 recite “prioritiz[ing] each of the detected potential cyber threats based on their respective risk to the protected virtual cloud asset.” The combination of Veselov and Basavapatna teaches these limitations. EX1002, ¶¶155-60.

As discussed above for elements 1.4 and 11.4, the combination teaches alerting based on a determined priority. *Supra*, §§XI.A.1.e, XI.A.3.g; EX1002, ¶156.

Both references also teach providing alerts for “each of the detected potential cyber threats,” as recited. EX1002, ¶157. For example, Veselov describes assessment results “describing one or more security vulnerabilities of the first computing resource; and, associating the first assessment result with the first computing resource.” EX1007, 23:19-47. Basavapatna similarly teaches risk metrics “for each of the one or more assets for each of the one or more vulnerabilities.” EX1008, ¶14.

Moreover, Basavapatna teaches prioritization of alerts “based on their respective risk to the protected virtual cloud asset,” as recited. Basavapatna describes risk metrics composed of several scores. For example, Basavapatna calculates a “standardized vulnerability score indicating a **relative level of risk** associated with the particular vulnerability.” EX1008, ¶12 (emphasis added). Basavapatna then calculates a “vulnerability detection score indicating an estimated probability that a particular asset possess[es] the particular vulnerability.” *Id.* Basavapatna combines this with a countermeasure component score indicating a probability that a countermeasure will mitigate the risk to the affected asset. *Id.* These scores are used to rank, or prioritize, the alerts sent to a user. *See, e.g., id.*, ¶¶138 (discussing ranking threats), 140 (risk metric rising above a specified threshold can trigger an alert); EX1002, ¶158.

Finally, a POSA had reason to combine these teachings of Veselov and Basavapatna for the reasons previously discussed for the independent claims, since the teachings of Basavapatna discussed above are the same types of well-known alert-prioritization techniques. *Supra*, §XI.A.4 (reasons to apply Basavapatna’s teachings regarding prioritization of alerts); EX1002, ¶159.

Accordingly, the combination of Veselov and Basavapatna teaches all aspects of elements 2.1 and 12.1.

b. Elements 2.2 and 12.2

Elements 2.2 and 12.2 recite “mitigat[ing] a potential cyber threat posing a risk to the protected virtual cloud asset.” The combination of Veselov and Basavapatna teaches the additional limitations in elements 2.2 and 12.2. EX1002, ¶¶161-64.

First, Veselov teaches these limitations. EX1002, ¶162. For example, Veselov explains that its scanning service “may take various actions on the assessment results including ... comparing the assessment results to a **remediation** framework to identify one or more actions the user can take to address the vulnerabilities, **and providing the identified actions** to the user; and the like.” EX1007, 10:17-36 (emphasis added).

To the extent Veselov does not teach these limitations, they were obvious in view of Basavapatna. Similar to Veselov, Basavapatna describes multiple types of “**countermeasures** that are part of the sensor.” EX1008, ¶24 (emphasis added). Such a “countermeasure protects against a threat by detecting and stopping an attack associated with the threat, by detecting and stopping activities associated with the attack, or by **mitigating damage caused by an attack.**” *Id.* (emphasis added); *see also id.*, ¶20 (describing passive and active countermeasures), ¶29 (describing active countermeasures). As Dr. Stavrou explains, a POSA would have understood that the disclosed countermeasures mitigate potential cyber threats

posing a risk to the protected virtual cloud assets. EX1002, ¶¶163-65. A POSA had reason to combine the teachings of Veselov and Basavapatna regarding mitigating actions, since Veselov already describes identifying remedial actions and providing them to the user, and the types of threat-mitigating countermeasures described in Basavapatna are comparable to Veselov's remedial actions. EX1002, ¶¶163-64. Indeed, a POSA would have recognized that the main purpose of identifying cyber threats is to facilitate mitigation. *Id.*; *see also supra*, §XI.A.4 (reasons to combine Veselov and Basavapatna).

Accordingly, claims 2 and 12 were obvious over the combination of Veselov and Basavapatna.

2. Dependent Claims 3 and 13

Claims 3 and 13 depend from claims 1 and 11, respectively, and further recite: “wherein determining the location of the snapshot of at least one virtual disk further comprises: determining a virtual disk allocated to the protected virtual cloud asset.” Veselov teaches the additional limitations in claims 3 and 13. EX1002, ¶¶166-69.

Veselov teaches “determining a virtual disk allocated to the protected virtual cloud asset.” As previously explained, Veselov teaches determining the location of a snapshot of a logical volume (i.e., a virtual disk). *Supra*, §XI.A.1.b. For example, Veselov explains that the “snapshot may be a copy of the state of

memory, the **state of any devices (virtual or physical) allocated to** the resource.” EX1007, 3:29-40 (emphasis added); *see also id.*, 6:7-10 (virtual computing environment may contain virtual machines and virtual disk storage), 11:36-41 (“[T]he snapshot data of the virtual machine 312 may be a volume image 308 of at least that portion (e.g., a partition) of the physical storage device that is associated with the logical volume 306 and/or that contains the virtual machine 312.”). As Dr. Stavrou explains, a POSA would have understood that performing a security assessment by obtaining a snapshot of a virtual disk that was allocated to a virtual machine would involve determining *which* virtual disk(s) were allocated to the protected virtual machine. EX1002, ¶¶166-67; *supra*, §XI.A.1.b (element 1.1).

At a minimum, it would have been obvious to determine a virtual disk allocated to the VM. EX1002, ¶¶167-68. Indeed, the snapshot discussed above is a snapshot *of* that virtual disk. Moreover, as Dr. Stavrou explains, it was common for systems seeking to interact with remote VMs or associated resources to query the remote hypervisors (such as the hypervisors in the computer systems hosting the target VMs in Veselov) in order to determine the virtual disks allocated to such VMs. EX1002, ¶168 (citing EX1015, 164-66). Determining a virtual disk allocated to the VM (i.e., the protected virtual cloud asset) was thus a routine part of locating and accessing a snapshot of that virtual disk. EX1002, ¶168.


Accordingly, claims 3 and 13 were obvious over the combination of Veselov and Basavapatna

3. Dependent Claims 4 and 14

Claims 4 and 14 depend from claims 2 and 12, respectively, and recite that the method further comprises “querying” (or that the system is further configured to “query”) “a cloud management console of the cloud computing environment to determine the location of the snapshot and the location of the virtual disk.”

Veselov teaches the additional limitations in claims 4 and 14. EX1002, ¶¶170-74.

First, Veselov teaches that determining the snapshot’s location can be accomplished by “querying a cloud management console of the cloud computing environment,” as recited. Discussing Fig. 1—teachings that apply to all embodiments (EX1002, ¶94; EX1007, 4:19-29)—Veselov discloses a user interface (“UI”) for accessing snapshots via an application programming interface (“API”) of the protected cloud environment. EX1007, 5:1-8 (“The scanning service 110 may access the computing resource service provider 104, or any associated virtual computing environments and/or computing resources, via the API 120 or another API.”); *see also id.*, 4:19-48 (describing UI for similar API access). For example, Figure 1 depicts the scanning service interacting with the computing resource service provider hosting the target VMs via an API:



EX1007, Fig. 1; *see also id.*, Figs. 3A-B (depicting API 338). While other figures do not depict an API between the scanning service and the hypervisor (*see id.*, Figs. 5A-B), Veselov states that “the user device 502, virtual computing environment 504A, ... as well as their implementation and interactions, are the same as described above for the associated components ... of FIGS. 3A-B.” *Id.*, 16:56-17:8. Further, a POSA would have recognized that API-mediated interactions between the scanning service and the virtualization layer of the target VMs was equally applicable to those embodiments. EX1002, ¶¶170-71.

As Dr. Stavrou explains, a POSA would have understood that this UI (alone or in conjunction with its associated API) operates as a “cloud management console,” as recited, or at least suggests such a console. EX1002, ¶172. As discussed above in §§XI.A.1.a-b, Veselov teaches, or least suggests, that the target resources are part of a cloud system. *See also supra*, §IX.A (cloud computing

background). Veselov further teaches that the API is used to manage various aspects of the computer environment hosting the target VMs. *See* EX1007, 4:19-48 (API may be “web-based interface implemented on a web server of the computing resource service provider,” in which “one or more user interfaces may be transmitted to the user device 102 and displayed thereon,” allowing environment configuration and other functions), 4:49-5:8 (“In some embodiments, the scanning service 110 may use the API 120 to provide user interfaces to the user device 102, enabling the user to configure settings of the scanning service 110”), 5:55-60 (“The virtualization layer 144 may also receive and process API calls from external devices or services.”), 14:37-65 (describing “using the user device 302 to access the scanning service 310 via one or more **APIs 338, web interfaces, and the like**” (emphasis added)). Veselov thus teaches interacting with this API via a UI for managing a cloud system (i.e., querying a “cloud management console”). At a minimum, this would have been obvious in view of Veselov’s teachings, since the disclosed API-mediated interactions are at least functionally comparable to the querying of a cloud management console, and such consoles were routinely used to facilitate remote management of cloud systems in analogous contexts. EX1002, ¶172 (citing EX1009, 221, 257-60; EX1015, 164-66).

Veselov also teaches, or at least suggests, querying this cloud management console “to determine the location of the snapshot and the location of the virtual

disk,” as recited. EX1002, ¶173. The user can initiate the snapshot-analysis process discussed with respect to claims 1 and 11, by using the UI to make an API call identifying the target VM, which a POSA would have understood to involve, or at least suggest, querying a cloud management console. *Id.*, 6:25-49; EX1002, ¶173; *see also id.*, 5:1-5 (describing use of environment API to access environment and protected resources), 5:59-60 (host system receiving and processing API calls). As discussed above with respect to element 1.1, this process involves determining the snapshot’s location. *Supra*, §XI.A.1.b. (*supra*, §XI.A.1.b). Querying consoles to determine resource locations was also a common step when managing and securing computer resources. *Id.*; EX1009, 221, 257-60; EX1015, 56, 124, 164-66, EX1031, 1 (Executive Summary), ii, 1-4. Accordingly, Veselov teaches, or at least suggests, that the previously discussed determination of the snapshot’s location (*supra*, §XI.A.1.b) is achieved by querying a cloud management console. EX1002, ¶¶172-73.

While the claims do not previously recite a “location of the virtual disk,” querying a cloud management console to determine such a location was nevertheless obvious. For example, Veselov shows the snapshot and its associated virtual disk initially located in the same computing environment, which at least suggests that querying the console to determine the snapshot’s general location (e.g., virtualization layer) would also determine the virtual disk’s general location.

EX1007, Figs. 3A-B, 5A-B (image 308/508 and logical volume 306/506 in same computing environment); EX1002, ¶174; *supra*, §XI.A.1.b (discussing general locations). Veselov also teaches using the UI/API to display and configure aspects of the assessment, including selection of the assessment target (e.g., the virtual disk). EX1007, 4:19-48, 8:52-62 (target resource may be logical volume), 11:28-33 (logical volume known as virtual disk drive), 14:37-63 (configuration/display of assessment information); *see also id.*, 16:31-37 (scanning service may provide logical volume's parameters to hypervisor). This teaches or at a minimum suggests querying the console to determine the virtual disk's location, because that location would facilitate the subsequent assessment and because determining the location of such virtual resources was a common feature of these consoles. EX1002, ¶174 (citing EX1009, 221, 257-60; EX1015, 16-23, 55-58, 115-20, 124-26, 150-52, 155-56; EX1031, 1 (Executive Summary), ii, 1-4). This also would have been obvious as part of the remediation process discussed in §XI.B.1.b (elements 2.2 and 12.2), since identifying/providing remedial actions, as taught by Veselov, at least suggests determining the location of the virtual disk whose vulnerability will be addressed. EX1007, 9:9-41, 10:17-36; EX1002, ¶174.

Accordingly, claims 4 and 14 were obvious over the combination of Veselov and Basavapatna.

4. Dependent Claims 5 and 15

Dependent claims 5 and 15 depend from claims 1 and 11, respectively, but recite substantially similar limitations. The preambles of claims 5 and 15 recite the “method of claim 1” and the “system of claim 11,” respectively, “wherein analyzing the snapshot of the protected virtual cloud asset further comprises....” Both claims then recite identical limitations, which are taught by the combination of Veselov and Basavapatna. EX1002, ¶175.

a. Elements 5.1 and 15.1

Elements 5.1 and 15.1 recite “parsing a copy of the snapshot.” Veselov teaches the additional limitations in elements 5.1 and 15.1. EX1002, ¶¶176-80.

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.A.1.b-c; *see, e.g.*, EX1007, 9:16-25 (emphasis added), Figs. 3A (snapshot copy brought to test environment 305), 5A (snapshot copy brought to test environment 505); EX1002, ¶177. Thus, Veselov teaches methods and systems for intaking and using a copy of a snapshot as a part of analyzing the snapshot.

Veselov also teaches “parsing” a copy of a snapshot as part of its security analysis. Veselov states that its 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 “[t]he scanning service 510 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). A POSA would have understood that these transformations involve “parsing” the copy of the snapshot, as recited. EX1002, ¶178. At a minimum, such parsing would have been obvious, since parsing was common, well-known technique for performing the types of data transformations described in Veselov. *Id.*, ¶¶178-79; EX1033, 1; EX1034, 2; EX1035, 389; EX1037, 11-15; EX1038, 2; EX1039, 14; EX1041, 3:1-9; EX1042, 10:45-55, 12:54-60; EX1043, 12:10; *supra*, §IX.B.

Accordingly, Veselov teaches all aspects of elements 5.1 and 15.1.

b. Elements 5.2 and 15.2

Veselov teaches elements 5.2 and 15.2, which recite “scanning the parsed copy to detect the potential cyber threats.” EX1002, ¶¶181-83. Veselov states that its scanning service may perform the security assessment by “reading the files of the mounted volume image” and performing “string or other pattern comparisons” to identify files, software, and other data. EX1007, 18:26-44; *see also id.*, 3:62-67

(following parsing, “the system may then perform the security assessment on the data at the file system level”). Accordingly, Veselov teaches “scanning a parsed copy.”

Veselov further teaches that this scanning is employed “to detect the potential cyber threats,” as recited. EX1002, ¶¶182-83. For example, Veselov describes generating assessment results by performing security checks depending on what software applications the scan reveals to be installed on the parsed copy of the snapshot. EX1007, 18:44-53. As previously discussed, these assessment results identify the detected cyber threats. *Supra*, §XI.A.1.e.

Accordingly, Veselov teaches all aspects of elements 5.2 and 15.2.

c. Elements 5.3 and 15.3

Elements 5.3 and 15.3 recite “wherein the potential cyber threats include known and unknown vulnerabilities.” Both Veselov and Basavapatna independently teach the additional limitations in elements 5.3 and 15.3. EX1002, ¶¶184-90.

First, Veselov teaches that the cyber threats it addresses include “known vulnerabilities” as recited. *Id.*, ¶¶185-86. 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). The ’735 patent itself 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).

Second, Veselov teaches, or at a minimum suggests, that the cyber threats it addresses include “unknown vulnerabilities,” as recited. It describes various analysis techniques that go beyond looking for specific, known vulnerabilities. For example, it explains that its techniques build on prior art 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, ¶¶185, 87 (citing EX1025, 12); *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.”

To the extent Veselov itself does not teach or suggest this limitation, it would have been obvious in view of Basavapatna. EX1002, ¶¶188-89.

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, ¶¶185, 88-90; *supra*, §IX.B (background).

Accordingly, the combination of Veselov and Basavapatna teaches all aspects of elements 5.3 and 15.3.

d. Elements 5.4 and 15.4

Elements 5.4 and 15.4 recite “wherein the detection is based on a type of vulnerability.” Veselov teaches the additional limitations in elements 5.4 and 15.4. EX1002, ¶¶191-93.

Veselov discloses detecting vulnerabilities “based on a type of vulnerability” as recited. 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). Moreover, as Dr. Stavrou explains, a POSA would have understood that a security assessment based on rules packages such as CVEs would be based on the types of vulnerabilities. EX1002, ¶192; *supra*, §XI.A.1.d (Veselov teaches CVE-based analysis). Veselov thus teaches elements 5.4 and 15.4.

Accordingly, claims 5 and 15 were obvious over the combination of Veselov and Basavapatna.

5. Dependent Claims 6 and 16

Dependent claims 6 and 16 depend from claims 5 and 15, respectively, but recite substantially similar limitations. The preamble of claim 6 recites the

“method of claim 5, wherein scanning the parsed copy further comprises any one of ...” and the preamble of claim 16 recites the “system of claim 15, wherein scanning the parsed copy further comprises any one of....” The claims then list the same four alternative limitations, any one of which is sufficient to satisfy the claims. Veselov teaches at least the first and fourth listed elements. EX1002, ¶¶194-204.

a. Elements 6.1 and 16.1

Elements 6.1 and 16.1 recite “checking configuration files of applications and operating system installed in the protected virtual cloud asset.” Veselov teaches these limitations. EX1002, ¶¶196-200.

Veselov teaches “checking configuration files of applications and operating system” installed in the VM (i.e., the protected virtual cloud asset). Veselov states that “[t]he ‘state’ of the target resource represented by the snapshot data may include any properties of the target resource . . . that are relevant to the security assessment.” EX1007, 9:55-67. These relevant properties include:

file system and file information; **operating system information**; system and software patches applied; configurations of network protocols and services; properties of connections to other computing resources; permission settings; **software packages installed, and configurations thereof**; system log properties; and the like.

Id. (emphasis added); *see also id.*, Abstract (reproducing a target virtual machine based on its configuration), 1:22-34 (describing configurations of operating system and installed applications), 7:16-22 (snapshot may include “virtual machine package manager database, which lists all software packages installed on the virtual machine and their configurations”); EX1002, ¶¶197-98.

Veslov further teaches checking these configuration files when scanning the parsed copy of the snapshot. As part of the previously discussed scan of the parsed copy, Veslov states:

once the file system of the mounted volume image can be read, the scanning service may compare the names of directories and/or files in the image to character patterns representing known directory and/or filenames **to identify** ‘signature’ directories and **files that identify the software application, version, installation attributes, etc.**”

EX1007, 18:32-44 (emphasis added). Veslov explains that the scanning service can maintain a database of assessment tasks “associated with particular software applications and versions thereof” to “identify which assessment tasks to run based on the installed software packages.” *Id.*, 18:44-50; *see also id.*, 1:43-50 (analysis includes identification of “bad configurations”), 2:35-41 (service may apply rules packages including “host configuration assessments”), 7:16-22 (discussing use of configuration data in snapshot), 21:61-22:6 (security assessment may include

analysis of VM configuration), claim 9 (same); *supra*, §§XI.B.4.a-b; EX1002, ¶199.

Accordingly, Veselov teaches all aspects of elements 6.1 and 16.1.

b. Elements 6.4 and 16.4

Veselov teaches, or at least suggests, elements 6.4 and 16.4, which recite “analyzing machine memory stored in the snapshot to deduce what applications and modules executed in the protected virtual cloud asset.” EX1002, ¶¶201-04. Veselov 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. The previously discussed scanning of the parsed copy thus includes “analyzing machine memory stored in the snapshot,” as recited. EX1002, ¶¶202-03; *supra*, §§XI.B.4.a-b. Veselov further explains that the snapshot (which includes the state of memory) includes properties 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, 18:44-50; *see also id.*, 7:15-22 (assessment includes determination of installed software), 9:55-67 (snapshot may include “software packages installed” on protected resource). Veselov thus teaches, or at least suggests, analyzing machine memory stored in the snapshot “to deduce what applications and modules executed in the protected virtual cloud asset,” as recited.

EX1002, ¶¶202-04. Accordingly, Veselov teaches or suggests all aspects of elements 6.4 and 16.4.

Claims 6 and 16 were therefore obvious over the combination of Veselov and Basavapatna.

6. Dependent Claims 7 and 17

Claims 7 (method) and 17 (system) depend from dependent claims 5 and 15, respectively, and recite further limitations taught by Veselov. EX1002, ¶¶205-12.

a. Elements 7.1 and 17.1


Elements 7.1 and 17.1 recite “instantiat[ing] an instance copy of the protected virtual cloud asset from the snapshot.” Veselov teaches the additional limitations in elements 7.1 and 17.1. EX1002, ¶¶206-09.

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 (Veselov), 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.”). Veselov also teaches or suggests performing multiple scans involving the same snapshot. *Id.*, 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....”). 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 5 and 15—from which claims 7 and 17 respectively depend—and the VM-instantiation techniques discussed below. EX1002, ¶206; *supra*, §XI.B.5 (claims 5 and 15).

Turning to the elements 7.1 and 17.1, 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).

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, ¶¶207-08.

Accordingly, Veselov teaches all aspects of elements 7.1 and 17.1.

b. Elements 7.2 and 17.2

Veselov teaches the additional limitations of elements 7.2 and 17.2, which recite “monitor[ing] all activity performed by the instance copy of the protected virtual cloud asset.” EX1002, ¶¶210-12. 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.*, 15:60-16:1; *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 copy, as recited in elements 7.2 and 17.2.

Accordingly, claims 7 and 17 were obvious over the combination of Veselov and Basavapatna.

7. Dependent Claims 9 and 19

Claims 9 (method) and 19 (system) depend from independent claims 1 and 11, respectively, and further recite: “wherein the protected virtual cloud asset

includes **any one of**: a virtual machine, a software container, a micro-service.”
(emphasis added). Satisfying any one of the listed elements is therefore sufficient to render these claims obvious.

Veselov at least teaches that the protected virtual cloud asset includes a “virtual machine” and a “software container.” EX1002, ¶¶213-15. First, Veselov describes “[s]ystems for performing a security assessment of a target computing resource, such as a **virtual machine** or an instance of a virtual machine” EX1007, Abstract (emphasis added); *see also id.*, 3:30-32 (“target resource may be an instance of a virtual machine”); *supra*, §XI.A.1.b. Veselov also teaches that the target computing resource can be a software container. EX1007, 2:44-51 (target can be “container of a virtual machine”); *see also id.*, 1:34-40 (describing containers).

Accordingly, claims 9 and 19 were obvious over the combination of Veselov and Basavapatna.

XII. CONCLUSION

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

Dated: January 8, 2024

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 13,979 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 8, 2024

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. 11,431,735 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. 16/585,967 ('735 Pat.)
1005	File History of U.S. Patent Application No. 16/750,556
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 et al., "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	Microsoft Computer Dictionary, 5 th ed., (2002)
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 et al.
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.20171.026
1035	Lengyel, “Scalability, Fidelity and Stealth in the DRAKVUF Dynamic Malware Analysis System (2014) https://doi.org/10.1145/2664243.2664252
1036	U.S. Patent No. 8k613,080 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) http://dx.doi.org/10.1145/2731186.2731196
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, Published May 2, 2013, 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, Published December 12, 2019, to Canzanese Jr.
1046	U.S. Patent 10,783,241 B2 to Crabtree

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. 11,431,735 and Exhibits 1001-1046** by Federal Express Next Business Day Delivery on January 8, 2024, 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 8, 2024

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