

**UNITED STATES PATENT AND TRADEMARK OFFICE**

---

**BEFORE THE PATENT TRIAL AND APPEAL BOARD**

---

ORACLE CORPORATION,  
Petitioner

v.

VIRTAMOVE, CORP.,  
Patent Owner

---

Case No.: IPR2025-01001  
U.S. Patent No. 7,519,814  
Issue Date: April 14, 2009

Title: SYSTEM FOR CONTAINERIZATION OF APPLICATION SETS

---

**PETITION FOR *INTER PARTES* REVIEW  
OF CLAIMS 1, 2, 4, 6, 8-10, AND 13-14 OF U.S. PATENT NO. 7,519,814**

# TABLE OF CONTENTS

	<b>Page</b>
INTRODUCTION .....	1
I. BACKGROUND .....	2
A. Containers Versus Virtual Machines .....	2
B. Containers Versus Shared Application Environment .....	3
C. Containers in the Prior Art .....	3
1. Linux VServer (Gélinas).....	4
2. Solaris Zones (Tucker).....	5
3. Zap Pods (Osman).....	6
II. THE '814 PATENT .....	7
A. Overview .....	7
B. Prosecution History .....	8
III. STATEMENT OF RELIEF REQUESTED .....	9
A. Grounds .....	9
B. The References Are Prior Art.....	10
1. The Patent's Filing Date .....	10
2. Osman .....	12
3. Tucker .....	13
4. Bandhole .....	15
5. Gélinas.....	16
IV. LEVEL OF ORDINARY SKILL.....	17
V. CLAIM CONSTRUCTION .....	18
A. "Container" and "System Files".....	18

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
B. “Disparate Computing Environments” .....	18
VI. GROUNDS OF UNPATENTABILITY.....	20
A. Ground 1: Claims 1, 2, 4, 6, 8-10, and 13-14 are Unpatentable as Obvious in View of Osman.....	20
1. Claim 1 .....	20
a. Limitation 1[pre][i]: “In a system having a plurality of servers” .....	20
b. Limitation 1[pre][ii]: “with operating systems that differ” .....	20
c. Limitation 1[pre][iii]: “operating in disparate computing environments”.....	21
d. Limitation 1[pre][iv]: “wherein each server includes a processor and an operating system including a kernel”.....	22
e. Limitation 1[pre][v]: “a set of associated local system files compatible with the processor” .....	22
f. Limitation 1[pre][vi]: “a method of providing at least some of the servers in the system with secure, executable applications related to a service”.....	23
g. Limitation 1[pre][vii]: “wherein the applications are executed in a secure environment” .....	23
h. Limitation 1[pre][viii]: “wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service” .....	24

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
i.    Limitation 1[a][i] “the method comprising: storing in memory accessible to at least some of the servers a plurality of secure containers of application software” .....	24
j.    Limitation 1[a][ii]: “each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications” .....	26
k.    Limitation 1[a][iii]: “for use with a local kernel residing permanently on one of the servers” .....	26
l.    Limitation 1[a][iv]: “wherein the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems” .....	27
m.    Limitation 1[a][v]: “the containers of application software excluding a kernel” .....	27
n.    Limitation 1[a][vi]: “wherein some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server” .....	28
o.    Limitation 1[a][vii]: “wherein said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server” .....	29
p.    Limitation 1[a][viii]: “wherein the application software cannot be shared between the plurality of secure containers of application software” .....	29

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
q.    Limitation 1[a][ix]: “wherein each of the containers has a unique root file system that is different from an operating system’s root file system” .....	30
2.    Claim 2: “wherein each container has an execution file associated therewith for starting the one or more applications”.....	30
3.    Claim 4: “pre-identifying applications and system files required for association with the one or more containers prior to said storing step” .....	30
4.    Claim 6: “assigning a unique associated identity to each of a plurality of the containers, wherein the identity includes at least one of IP address, host name, and MAC address” .....	31
5.    Claim 8: “wherein the one or more applications and associated system files are retrieved from a computer system having a plurality of secure containers” .....	31
6.    Claim 9: “wherein server information related to hardware resource usage including at least one of CPU memory, network bandwidth, and disk allocation is associated with at least some of the containers prior to the applications within the containers being executed” .....	32
7.    Claim 10: “wherein in operation when an application residing within a container is executed, said application has no access to system files or applications in other containers or to system files within the operating system during execution thereof” .....	33
8.    Claim 13: “associating with a plurality of containers a stored history of when processes related to applications within the container are executed for at least one of, tracking statistics, resource allocation, and for monitoring the status of the applicaion” .....	34

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
9. Claim 14.....	34
a. Limitation 14[a][i]: “creating containers prior to said step of storing containers in memory, wherein containers are created by:” .....	35
b. Limitation 14[a][ii]: “a) running an instance of a service on a server” .....	35
c. Limitation 14[a][iii]: “b) determining which files are being used” .....	35
d. Limitation 14[a][iv]: “c) copying applications and associated system files to memory without overwriting the associated system files so as to provide a second instance of the applications and associated system files” .....	35
B. Ground 2: Claims 1, 2, 4, 6, 8-10, and 13 are Unpatentable as Obvious in View of Tucker and Bandhole.....	36
1. Claim 1 .....	38
a. Limitation 1[pre][i]: “In a system having a plurality of servers” .....	38
b. Limitation 1[pre][ii]: “with operating systems that differ” .....	39
c. Limitation 1[pre][iii]: “operating in disparate computing environments” .....	39
d. Limitation 1[pre][iv]: “wherein each server includes a processor and an operating system including a kernel” .....	40
e. Limitation 1[pre][v]: “a set of associated local system files compatible with the processor” .....	40

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
f. Limitation 1[pre][vi]: “a method of providing at least some of the servers in the system with secure, executable applications related to a service” .....	41
g. Limitation 1[pre][vii]: “wherein the applications are executed in a secure environment” .....	42
h. Limitation 1[pre][viii]: “wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service” .....	42
i. Limitation 1[a][i] “the method comprising: storing in memory accessible to at least some of the servers a plurality of secure containers of application software” .....	42
j. Limitation 1[a][ii]: “each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications” .....	43
k. Limitation 1[a][iii]: “for use with a local kernel residing permanently on one of the servers” .....	44
l. Limitation 1[a][iv]: “wherein the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems” .....	44
m. Limitation 1[a][v]: “the containers of application software excluding a kernel” .....	44
n. Limitation 1[a][vi]: “wherein some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server” .....	44

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
o.    Limitation 1[a][vii]: “wherein said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server” .....	46
p.    Limitation 1[a][viii]: “wherein the application software cannot be shared between the plurality of secure containers of application software” .....	46
q.    Limitation 1[a][ix]: “wherein each of the containers has a unique root file system that is different from an operating system’s root file system” .....	47
2.    Claim 2: “wherein each container has an execution file associated therewith for starting the one or more applications” .....	47
3.    Claim 4: “pre-identifying applications and system files required for association with the one or more containers prior to said storing step” .....	48
4.    Claim 6: “assigning a unique associated identity to each of a plurality of the containers, wherein the identity includes at least one of IP address, host name, and MAC address” .....	48
5.    Claim 8: “wherein the one or more applications and associated system files are retrieved from a computer system having a plurality of secure containers” .....	48
6.    Claim 9: “wherein server information related to hardware resource usage including at least one of CPU memory, network bandwidth, and disk allocation is associated with at least some of the containers prior to the applications within the containers being executed” .....	49

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
7.	Claim 10: “wherein in operation when an application residing within a container is executed, said application has no access to system files or applications in other containers or to system files within the operating system during execution thereof” .....49
8.	Claim 13: “associating with a plurality of containers a stored history of when processes related to applications within the container are executed for at least one of, tracking statistics, resource allocation, and for monitoring the status of the applicaion” .....50
C.	Ground 3: Claims 1, 2, 4, 6, 8-10, and 13-14 are Unpatentable as Obvious in View of Gélinas.....51
1.	Claim 1 .....51
a.	Limitation 1[pre][i]: “In a system having a plurality of servers” .....51
b.	Limitation 1[pre][ii]: “with operating systems that differ” .....51
c.	Limitation 1[pre][iii]: “operating in disparate computing environments” .....52
d.	Limitation 1[pre][iv]: “wherein each server includes a processor and an operating system including a kernel” .....53
e.	Limitation 1[pre][v]: “a set of associated local system files compatible with the processor” .....53
f.	Limitation 1[pre][vii]: “a method of providing at least some of the servers in the system with secure, executable applications related to a service” .....54
g.	Limitation 1[pre][vii]: “wherein the applications are executed in a secure environment” .....54

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
h. Limitation 1[pre][viii]: “wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service” .....	54
i. Limitation 1[a][i] “the method comprising: storing in memory accessible to at least some of the servers a plurality of secure containers of application software” .....	54
j. Limitation 1[a][ii]: “each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications” .....	55
k. Limitation 1[a][iii]: “for use with a local kernel residing permanently on one of the servers” .....	55
l. Limitation 1[a][iv]: “wherein the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems” .....	56
m. Limitation 1[a][v]: “the containers of application software excluding a kernel” .....	56
n. Limitation 1[a][vi]: “wherein some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server” .....	57
o. Limitation 1[a][vii]: “wherein said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server” .....	57

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
p.    Limitation 1[a][viii]: “wherein the application software cannot be shared between the plurality of secure containers of application software” .....	57
q.    Limitation 1[a][ix]: “wherein each of the containers has a unique root file system that is different from an operating system’s root file system” .....	58
2.    Claim 2: “wherein each container has an execution file associated therewith for starting the one or more applications” .....	58
3.    Claim 4: “pre-identifying applications and system files required for association with the one or more containers prior to said storing step” .....	59
4.    Claim 6: “assigning a unique associated identity to each of a plurality of the containers, wherein the identity includes at least one of IP address, host name, and MAC address” .....	59
5.    Claim 8: “wherein the one or more applications and associated system files are retrieved from a computer system having a plurality of secure containers” .....	60
6.    Claim 9: “wherein server information related to hardware resource usage including at least one of CPU memory, network bandwidth, and disk allocation is associated with at least some of the containers prior to the applications within the containers being executed” .....	60
7.    Claim 10: “wherein in operation when an application residing within a container is executed, said application has no access to system files or applications in other containers or to system files within the operating system during execution thereof” .....	61

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
8. Claim 13: “associating with a plurality of containers a stored history of when processes related to applications within the container are executed for at least one of, tracking statistics, resource allocation, and for monitoring the status of the applicaion” .....	61
9. Claim 14 .....	62
a. Limitation 14[a][i]: “creating containers prior to said step of storing containers in memory, wherein containers are created by:” .....	62
b. Limitation 14[a][ii]: “a) running an instance of a service on a server” .....	62
c. Limitation 14[a][iii]: “b) determining which files are being used” .....	62
d. Limitation 14[a][iv]: “c) copying applications and associated system files to memory without overwriting the associated system files so as to provide a second instance of the applications and associated system files” .....	63
VII. SECONDARY CONSIDERATIONS OF NONOBVIOUSNESS .....	63
VIII. DISCRETIONARY DENIAL IS UNWARRANTED .....	64
IX. MANDATORY NOTICES .....	64
A. Real Parties-In-Interest (37 C.F.R. §42.8(b)(1)) .....	64
B. Related Matters .....	65
1. <b>United States Patent &amp; Trademark Office</b> .....	65
2. <b>USPTO Patent Trial and Appeal Board</b> .....	65
3. <b>U.S. District Court for the Eastern District of Texas</b> .....	68
4. <b>U.S. District Court for the Western District of Texas</b> .....	68

**TABLE OF CONTENTS**  
(continued)

	<b>Page</b>
<b>5. U.S. District Court for the Northern District of California .....</b>	<b>68</b>
C. Counsel, Service, and Fee Information .....	69
D. Payment of Fees (37 C.F.R. §42.103).....	70
E. Grounds for Standing (37 C.F.R. §42.104(a)) .....	70
X. CONCLUSION.....	70

## TABLE OF AUTHORITIES

	<b>Page(s)</b>
<b>Cases</b>	
<i>Amazon.com, Inc. v. CustomPlay, LLC</i> , IPR2018-01496, Paper 34 (P.T.A.B. Mar. 4, 2020).....	17, 13
<i>Amazon.com, Inc. v. VirtaMove, Corp.</i> , Case No. IPR2025-00561 (PTAB, filed January 31, 2025) .....	66
<i>Amazon.com, Inc. v. VirtaMove, Corp.</i> , Case No. IPR2025-00563 (PTAB, filed January 31, 2025) .....	66
<i>Amazon.com, Inc. v. VirtaMove, Corp.</i> , Case No. IPR2025-00566 (PTAB, filed January 31, 2025) .....	66
<i>Apple v. Fintiv</i> , IPR2020-00019, Paper 11 (P.T.A.B. Mar. 20, 2020).....	17
<i>Dynamic Drinkware, LLC v. Nat’l Graphics, Inc.</i> , 800 F.3d 1375 (Fed. Cir. 2015) .....	17, 10, 13
<i>Google LLC v. Multimodal Media LLC</i> , IPR2024-00063, Paper 10 (P.T.A.B. May 29, 2024) .....	17
<i>Google LLC v. VirtaMove, Corp.</i> , Case No. IPR2025-00487 (PTAB, filed January 31, 2025) .....	66
<i>Google LLC v. VirtaMove, Corp.</i> , Case No. IPR2025-00488 (PTAB, filed January 31, 2025) .....	66
<i>Google LLC v. VirtaMove, Corp.</i> , Case No. IPR2025-00489 (PTAB, filed January 31, 2025) .....	66
<i>Google LLC v. VirtaMove, Corp.</i> , Case No. IPR2025-00490 (PTAB, filed January 30, 2025) .....	66
<i>In re GPAC Inc.</i> , 57 F.3d 1573 (Fed. Cir. 1995) .....	17, 17

<i>Halliburton Energy Servs., Inc. v. Dynamic 3D Geosolutions LLC</i> , IPR 2014-01186, 2015 WL 5565065 (PTAB Dec. 15, 2015) .....	17, 17
<i>Hulu, LLC v. Sound View Innovations, LLC</i> , IPR2018-01039, Paper 29 (P.T.A.B. Dec. 20, 2019) .....	17, 16
<i>Intel Corp. v. Qualcomm Inc.</i> , 21 F.4th 801 (Fed. Cir. 2021) .....	17, 19
<i>International Business Machines Corp. v. VirtaMove, Corp.</i> , Case No. IPR2025-00591 (PTAB, filed February 6, 2025) .....	66
<i>International Business Machines Corp. v. VirtaMove, Corp.</i> , Case No. IPR2025-00599 (PTAB, filed February 7, 2025) .....	67
<i>Keysight Techs., Inc. v. Centripetal Networks, LLC</i> , IPR2022-01525, Paper 27 (P.T.A.B. Apr. 15, 2024) .....	18, 16
<i>Kolcraft Enters., Inc. v. Graco Children’s Prods., Inc.</i> , 927 F.3d 1320 (Fed. Cir. 2019) .....	18, 12
<i>Leapfrog Enters. v. Fisher-Price, Inc.</i> , 485 F.3d 1157 (Fed. Cir. 2007) .....	18, 63
<i>Medtronic, Inc. v. Teleflex Innovations S.A.R.L.</i> , 68 F.4th 1298 (Fed. Cir. 2023) .....	18, 12
<i>Microsoft Corp. v. VirtaMove, Corp.</i> , Case No. IPR2025-00849 (PTAB, filed April 18, 2025) .....	67
<i>Microsoft Corp. v. VirtaMove, Corp.</i> , Case No. IPR2025-00850 (PTAB, filed April 18, 2025) .....	67
<i>Microsoft Corp. v. VirtaMove, Corp.</i> , Case No. IPR2025-00851 (PTAB, filed April 18, 2025) .....	67
<i>Microsoft Corp. v. VirtaMove, Corp.</i> , Case No. IPR2025-00852 (PTAB, filed April 18, 2025) .....	67
<i>Microsoft Corp. v. VirtaMove, Corp.</i> , Case No. IPR2025-00853 (PTAB, filed April 18, 2025) .....	67

<i>Microsoft Corp. v. VirtaMove, Corp.</i> , Case No. IPR2025-00854 (PTAB, filed April 18, 2025) .....	67
<i>Microsoft Corp. v. VirtaMove, Corp.</i> , Case No. IPR2025-00855 (PTAB, filed April 18, 2025) .....	67
<i>Microsoft Corp. v. VirtaMove, Corp.</i> , IPR2025-00851, Paper 3 (PTAB Apr. 18, 2025) .....	18, 64
<i>New Railhead Mfg., LLC v. Vermeer Mfg. Co.</i> , 298 F.3d 1290 (Fed. Cir. 2002) .....	18, 10
<i>Newell Cos. v. Kenney Mfg. Co.</i> , 864 F.2d 757 (Fed. Cir. 1988) .....	18, 63
<i>Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co. Ltd.</i> , 868 F.3d 1013 (Fed. Cir. 2017) .....	18, 18
<i>Red Hat, Inc. v. VirtaMove, Corp.</i> , Case No. 5:24-cv-04740 .....	18, 68
<i>Sotera Wireless, Inc. v. Masimo Corp.</i> , IPR2020-01019, Paper 12 (P.T.A.B. Dec. 1, 2020) .....	18
<i>VirtaMove, Corp. v. Amazon.com, Inc. et al</i> , Case No. 7:24-cv-00030 (pending transfer to Northern District of California per Order dated February 19, 2025, <i>see</i> Docket Entry No. 94) .....	68
<i>VirtaMove, Corp. v. Amazon.com, Inc.</i> , No. 7:24-cv-00030 .....	18, 18
<i>VirtaMove, Corp. v. Google LLC</i> , 7:24-cv-00033-DC-DTG (W.D. Tex.) (May 7, 2025).....	18, 2
<i>VirtaMove, Corp. v. Google LLC</i> , Case No. 5:25-cv-00860 .....	68
<i>VirtaMove, Corp. v. Google LLC</i> , Case No. 7:24-cv-00033 (transferred to Northern District of California per Order dated May 7, 2025, <i>see</i> Ex. 1028).....	68

<i>VirtaMove, Corp. v. Hewlett Packard Enterprise Company,</i> Case No. 2:24-cv-00093 .....	68
<i>VirtaMove, Corp. v. International Business Machines Corp.,</i> Case No. 2:24-cv-00064 .....	68
<i>VirtaMove, Corp. v. Microsoft Corp.,</i> Case No. 7:24-cv-00338 .....	68
<i>VirtaMove, Corp. v. Oracle Corp.,</i> 7:24-cv-00339 (W.D. Tex.) (Mar. 28, 2025).....	18, 2, 3
<i>VirtaMove, Corp. v. Oracle Corp.,</i> Case No. 7:24-cv-00339 .....	68
<i>VirtaMove, Corp. v. Oracle Corporation,</i> Case No. 7:24-cv-00339-ADA .....	66
<i>Vivid Techs., Inc. v. Am. Sci. &amp; Eng’g, Inc.,</i> 200 F.3d 795 (Fed. Cir. 1999) .....	18, 18
<i>Voter Verified, Inc. v. Premier Election Sols., Inc.,</i> 698 F.3d 1374 (Fed. Cir. 2012) .....	19, 16

**Statutes**

35 U.S.C. §102.....	19, 15
35 U.S.C. §103.....	19, 9
35 U.S.C. §314.....	19
35 U.S.C. §325.....	19

**Other Authorities**

37 C.F.R. §42.8 .....	19
37 C.F.R. §42.8(b)(1).....	19, 64
37 C.F.R. §42.10 .....	19
37 C.F.R. §42.10(b) .....	19, 70

37 C.F.R. §42.15 .....	19
37 C.F.R. §42.15(a).....	19, 70
37 C.F.R. §42.103 .....	19, 70
37 C.F.R. §42.104 .....	19
37 C.F.R. §42.104(a).....	19, 70
77 Fed. Reg. 48759 .....	19, 65
Osman et al., <i>The Design and Implementation of</i> .....	2

## TABLE OF EXHIBITS

Exhibit No.	Description
1001	U.S. Patent No. 7,519,814 (“the ’814 patent”)
1002	Declaration of Dr. Darrell Long, Ph.D.
1003	Osman et al., <i>The Design and Implementation of Zap: A System for Migrating Computing Environments</i> , 5 Proc. of the Symposium on Operating Systems Design and Implementation (2002) (“Osman”)
1004	U.S. Patent No. 7,437,556 (“Tucker”)
1005	U.S. Provisional Patent Application No. 60/469,558 (“Tucker Provisional”)
1006	U.S. Patent Publication No. 2002/0171678A1 (“Bandhole”)
1007	<i>Virtual Private Servers and Security Contexts</i> (“Gélinas”)
1008	File history of the ’814 patent
1009	Solaris 9 press release from Sun Microsystems
1010	B. Walters, “VmWare Virtual Platform.” <i>Linux Journal</i> , 1999.
1011	Soltesz et al., <i>Container-based operating system virtualization: a scalable, high-performance alternative to hypervisors</i> (2007)
1012	D. Price and A. Tucker. <i>Solaris zones: Operating system support for consolidating commercial workloads</i> . In Proceedings of the 18th Usenix LISA Conference, 2004.
1013	U.S. Provisional Patent Application No. 60/502,619
1014	U.S. Provisional Patent Application No. 60/512,103
1015	Declaration of Rachel Watters Regarding Osman
1016	Declaration of Jacques Gélinas Regarding Linux VServer

Exhibit No.	Description
1017	Message to Linux Kernel Mailing List Regarding Linux VServer
1018	Slashdot post Regarding Linux VServer
1019	Amazon’s Opening Claim Construction Brief in <i>VirtaMove, Corp. v. Amazon.com, Inc. et al.</i> , No. 7:24-cv-30-ADA-DTG (W.D. Tex.) (the “Amazon Litigation”)
1020	Patent Owner’s Sur-Reply Claim Construction Brief from the Amazon Litigation
1021	Excerpts from deposition of named inventor Donn Rochette from the Amazon Litigation
1022	J. Ball, “Managing Initscripts with Red Hat’s chkconfig.” <i>Linux Journal</i> , 2001.
1023	Kravetz, et al. “Enhancing Linux scheduler scalability.” <i>Proceedings of the Ottawa Linux Symposium, Ottawa, CA</i> . 2001.
1024	Scheduling order from the Amazon Litigation
1025	Order cancelling <i>Markman</i> hearing in the Amazon Litigation
1026	Order granting transfer in the Amazon Litigation
1027	<i>Curriculum vitae</i> of Dr. Darrell Long
1028	Order Overruling Objections and Transferring Case to Northern District of California, in <i>VirtaMove, Corp. v. Google LLC</i> , 7:24-cv-00033-DC-DTG (W.D. Tex.) (May 7, 2025)
1029	Plaintiff VirtaMove Corp.’s Preliminary Disclosure of Asserted Claims and Infringement Contentions, in <i>VirtaMove, Corp. v. Oracle Corp.</i> , 7:24-cv-00339 (W.D. Tex.) (Mar. 28, 2025)

Petitioner Oracle Corporation (“Petitioner” or “Oracle”) requests *inter partes* review of claims 1, 2, 4, 6, 8-10, and 13-14 (the “challenged claims”) of U.S. Patent No. 7,519,814 (“the ’814 patent”), which VirtaMove, Corp. (“Patent Owner” or “PO”) purportedly owns.

## **INTRODUCTION**

The challenged claims recite methods for running software applications in “containers.” A container is a set of files needed to execute an application on a computer. The files in a container are grouped together and isolated from other files and applications on the same computer. Containers prevent different applications on the same computer from interfering with each other.

A host of prior-art container technologies—Solaris zones, Zap pods, Linux VServers, and more—provide the same functionality described in the challenged claims. All of these container technologies (and several others) were available and well known before the ’814 patent’s earliest claimed priority date in September 2003. Yet the patent fails to acknowledge these earlier container technologies.

The Examiner was aware of at least one of these technologies—VServer—and expressly recognized its relevance to the patent claims. But the only reference the Examiner cited concerning VServer published in 2007 and thus was not prior art. This 2007 publication omits aspects of VServer that are material to the patent claims.

The Examiner never considered the 2002 VServer reference raised in this Petition, which is prior art. Nor did the Examiner review the prior art references that this Petition relies on concerning Zap pods and Solaris zones. Each of these references discloses the elements that were missing from the Examiner's prior art and each of them renders the challenged claims unpatentable.

The '814 patent claims exclusive rights to container technology that belongs in the public domain. Because the patent contributed nothing to the art, PO is not entitled to exclude the public from practicing the challenged claims. Thus, the Board should cancel the claims.

## **I. BACKGROUND**

### **A. Containers Versus Virtual Machines**

Instead of addressing the many container systems in the prior art, the '814 patent frames its contribution as an advance over “Virtual Machine technology, pioneered by VmWare” and released commercially in 1999. (Ex. 1001, 1:51-56; Ex. 1010.). Like containers, multiple virtual machines can be hosted on a single physical computer and each one can be customized to meet the unique needs of the applications it contains. (Ex. 1001, 1:27-56.). But the patent identifies a “key difference” between the Virtual Machine (“VM”) approach and the patent's container-based approach. (Ex. 1001, 1:56-61.) While VMs require a copy of the operating system “for each application,” the container approach required only one

copy of the operating system (OS) “regardless of the number of application containers deployed.” (*Id.*) Because containers do not require multiple copies of the OS, they avoid the “performance overhead” associated with VMs. (*Id.*, 1:62-63.)

The claims of the ’814 patent capture this distinction over VMs by specifying that containers do not contain their own “kernel” (which is the core of an OS). (*E.g.*, *id.*, 17:41-50 (claim 1); *see also id.*, 2:39-42 (containers share a kernel from the underlying OS). However, the patent’s distinction does not apply to the prior-art systems presented in this Petition—all of which used containers rather than VMs.

#### **B. Containers Versus Shared Application Environment**

The ’814 patent acknowledged that, outside of VMs, multiple applications could share an operating system in a prior-art environment called “SoftGrid.” (*Id.*, 2:4-12.) The patent distinguished SoftGrid in that it did “not isolate applications into distinct environments.” (*Id.*) In contrast, containers isolate applications to prevent them from interfering with each other. (*Id.*, 4:39-42.)

#### **C. Containers in the Prior Art**

The ’814 patent fails to distinguish the many prior-art container systems that provided the same capabilities the patent describes—including the isolation and kernel sharing that the patent relies on to distinguish other prior art.

## 1. Linux VServer (Gélinas)

In October 2001, Jacques Gélinas disclosed a system for Linux computers that would allow “several independant [sic] virtual servers running on the same box (*sharing the same kernel as well*).” (Ex. 1017 (emphasis added).) This system became known as Linux VServer. (Ex. 1002 ¶38.) By 2002, Gélinas’s website described that VServer “split a Linux server into virtual ones with as much *isolation* as possible between each one, looking like real servers, yet sharing some common tasks.” (Ex. 1007, 1 (emphasis added).)

To provide isolation, VServer built on a well-known computer command called “chroot,” which had been part of conventional operating systems for decades. (Ex. 1002 ¶39.) The chroot command confined a particular application process to a limited view of the computer’s file system, locking it out of other areas of the computer. (*Id.*) VServer built on chroot to also prevent processes from communicating with each other and isolate them from network resources. (Ex. 1007, 2-3; Ex. 1002 ¶39.) Thus, VServer created secure containers that isolate and confine processes while allowing them to share the kernel of the underlying server. (Ex. 1002 ¶39.)

Although the Examiner cited a 2007 article that mentioned VServer (Ex. 1008, 15), the Examiner never considered the 2002 publication on which this Petition relies. (Ex. 1001, 1-2.) The 2007 article omits key features of VServer that

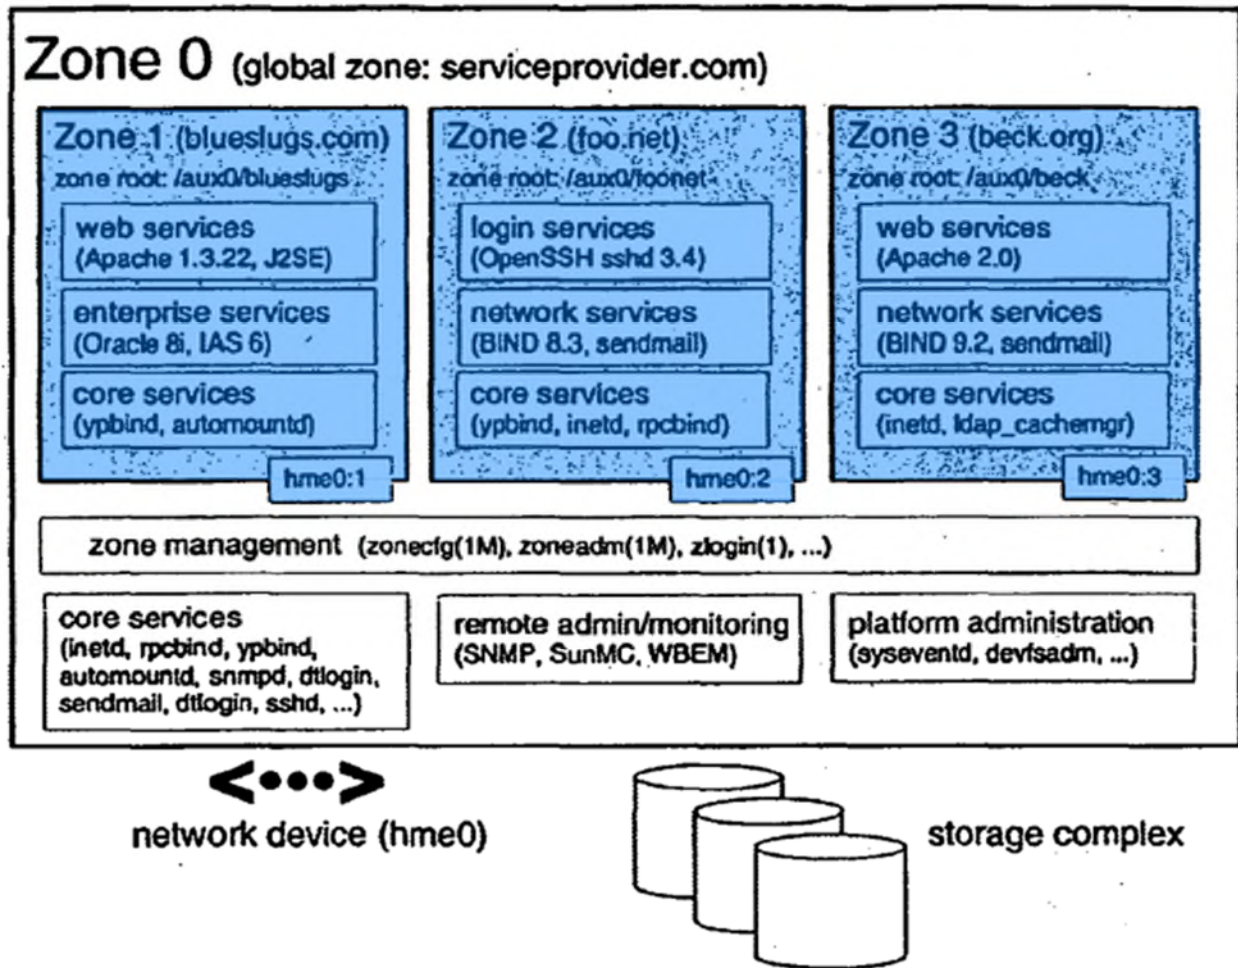
satisfy elements of the asserted patent claims and that are disclosed in the 2002 publication presented here. (*Infra* §VI.C.)

## 2. Solaris Zones (Tucker)

While Gélinas was developing VServer, Sun Microsystems was working on containers for its popular Solaris operating system. Sun published a press release in May 2002 that announced “containers” as part of Solaris 9. (Ex. 1009, 2-3.) Sun’s containers allowed “customers to run multiple applications on a single server, with fault, security and resource containment built-in.” (*Id.*, 3.)

As Sun continued to develop its container technology, it filed a provisional patent application describing a system called Solaris Zones in May 2003. (Ex. 1005.) Sun’s provisional explained that a “zone is an application container[.]” (*Id.*, 92.) That provisional eventually matured into an issued patent. (Ex. 1004.)

Sun’s technology allowed an operating system to be partitioned into a “global zone” and one or more “non-global zones”—containers that would isolate groups of processes from each other and from the underlying operating system. (Ex. 1005, 1-5.) Sun’s provisional also described non-global zones as isolated “‘sandbox[es]’ within which one or more applications can run without affecting or interacting with the rest of the system.” (*Id.*, 1.) Figure 1.1 from Sun’s provisional shows non-global zones 1, 2, and 3 ([containers](#)) running services to host three different web sites:



(Ex. 1005, 3 (annotated).) Thus, Solaris Zones are also secure containers that isolate and confine processes that share an underlying operating system. (Ex. 1002 ¶42.)

### 3. Zap Pods (Osman)

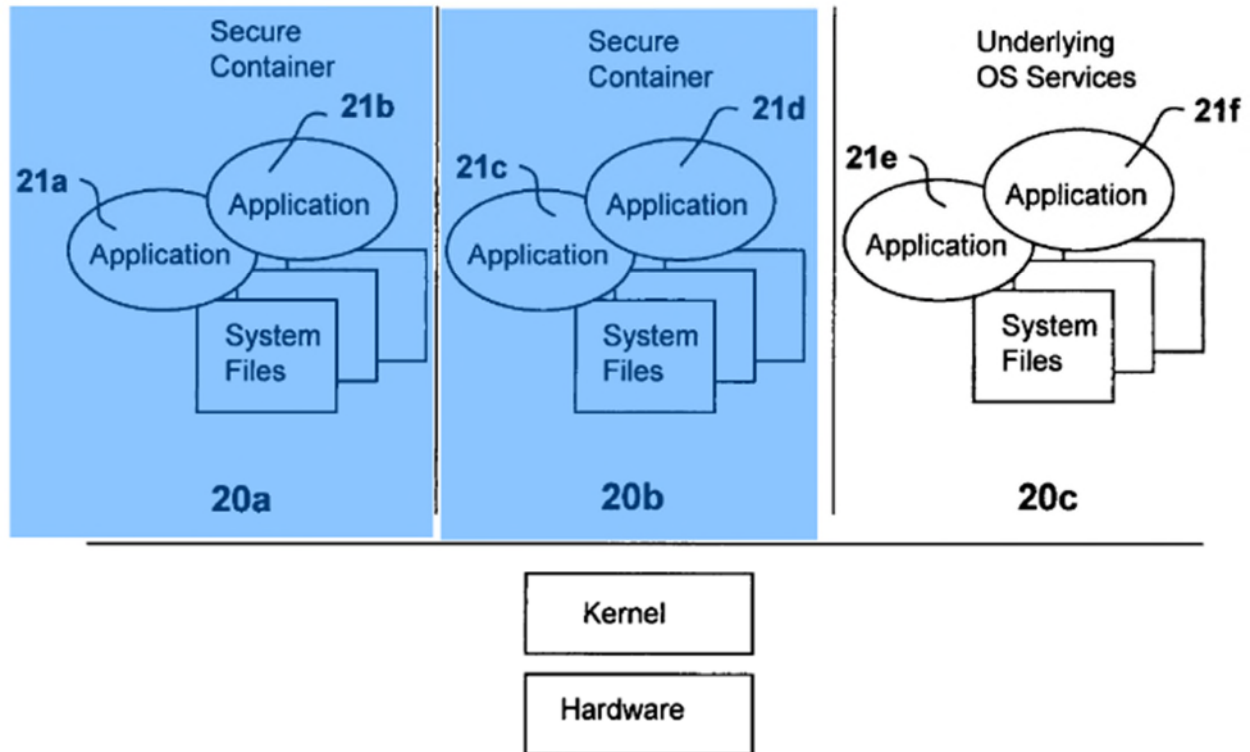
Containers were also the subject of academic research. In December 2002, researchers from Columbia University presented a paper on “Zap,” a system designed to allow groups of application processes to be easily moved from server to server. (E.g., Ex. 1003, 361, 367.) Zap isolated groups of processes into “pods,” which are “a group of processes with a private namespace that presents the process

group with the same virtualized view of the system.... This decouples processes in a pod from dependencies on the host operating system and from other processes in the system.” (*Id.*) Zap’s pods were isolated from the underlying server and from other pods so that they could be paused, saved, and resumed on another server without interruption. (*See id.*; Ex. 1002 ¶43.)

## **II. THE ’814 PATENT**

### **A. Overview**

The ’814 patent describes a system of “associating applications with secure containers.” (Ex. 1001, 7:4-15.) As Figure 2 shows, each **container** (20a, 20b) is comprised of applications and system files, and the container is thus “an aggregate of all files required to successfully execute a set of software applications on a computing platform.” (*Id.*, 7:22-29.)



**Figure 2**

(*Id.*, Fig. 2.)

As in the prior art, “each application executing associated with a container ... is able to access [sic] files that are dedicated to the container” but is “not able to access the files contained in another container” or “the file provided with the underlying operating system” that are outside of any container. (*Id.*, 8:66-9:7.)

### **B. Prosecution History**

The Examiner cited, among other references, a 2007 article by Soltesz et al. that was not prior art. (Ex. 1008, 15.) Soltesz discussed various aspects of the prior-art VServer system, but omitted aspects material to the claims at issue here. For example, the Examiner incorrectly determined that the prior art failed to disclose (1)

“applications software not being sharable between the plurality of secure (and isolated) containers”; (2) “unique root file systems different from an operating system’s root file system”; and (3) “different versions of the same operating system running on the same system/server environment.” (Ex. 1008, 15-16.) These features were all disclosed in Gélinas (Ex. 1007), which describes VServer in far more detail than Soltesz and was not before the Examiner (Ex. 1001, 1-2).

Finally, the Soltesz article also cites (among a list of 25 references) the Osman paper relied on in this Petition and a 2004 paper about Solaris zones. (Ex. 1011, 286-87 [Soltesz].) That Solaris paper is not prior art because it published in November 2004. (Ex. 1012.) As for the Osman paper, the Examiner never cited it and the patent applicant never submitted it. (Ex. 1001, 1-2.) Accordingly, the Patent Office never compared the claims of the ’814 patent to any of the prior art presented here.

### **III. STATEMENT OF RELIEF REQUESTED**

#### **A. Grounds**

The challenged claims should be canceled under 35 U.S.C. §103 as follows:

<b>Ground</b>	<b>Claims Challenged</b>	<b>References</b>
1	1, 2, 4, 6, 8-10, 13-14	Osman
2	1, 2, 4, 6, 8-10, 13	Tucker & Bandhole
3	1, 2, 4, 5, 8-10, 13-14	Gélinas

## **B. The References Are Prior Art**

### **1. The Patent's Filing Date**

The '814 patent claims priority to two provisional applications. (Ex. 1001.) But the challenged claims are not entitled to the filing date of either provisional application. For a patent to effectively claim priority to a provisional application, the specification of the provisional must contain a written description of the invention that enables a POSITA to practice the invention claimed in the patent. *Dynamic Drinkware, LLC v. Nat'l Graphics, Inc.*, 800 F.3d 1375, 1378 (Fed. Cir. 2015); *New Railhead Mfg., LLC v. Vermeer Mfg. Co.*, 298 F.3d 1290, 1294 (Fed. Cir. 2002). Neither provisional to which the '814 patent claims priority satisfies this standard.

First, the '619 provisional deals with subject matter entirely different from that addressed by challenged claim 1. (*Compare* Ex. 1013 *with* Ex. 1001, claim 1.) The '619 provisional is entitled "Drag & Drop Application Management" and deals with "software that manages the operation of a data center" through drag-and-drop installation. (Ex. 1013, 1-7.) The word "container" appears nowhere in this provisional. (*Id.*)

Second, while the '103 provisional deals with a container system, it fails to disclose all of claim 1. For example, it does not disclose claim 1's limitation that "said associated system files utilized in place of the associated local system files are

copies or modified copies of the associated local system files *that remain resident on the server.*” (Ex. 1001, 17:53-57.) The antecedent basis for “the server” is “one of the servers” in the claimed system on which a “local kernel” resides. (*Id.*, 17:43-46.) This “local kernel” is used to execute the applications and system files in each container. (*Id.*) Thus, claim 1 requires system files in a container to be copies of system files that “remain resident” on the server where the containers run.

In contrast, the ’103 provisional discloses that “data and configuration information are copied from the computing platform from which they originate” and packaged into a container before being installed on a “remote computing platform.” (Ex. 1014, 18-19). The provisional *does not* disclose that the “platform from which they originate” and the “remote” platform are the same server. (*Id.*) For at least this reason, the ’103 provisional fails to support claim 1 and its dependents.

The prosecution history of the ’814 patent confirms that the ’103 provisional fails to support claim 1. The applicant added the “copies or modified copies” limitation by amendment. (Ex. 1008, 28-29.) In doing so, the applicant identified alleged support for the limitation. (*Id.*, 37-41) Specifically, the applicant identified Figure 9 and paragraph 91 of the non-provisional application. (*Id.*, 38.) But this figure and paragraph do not appear in the ’103 provisional. (Ex. 1014.) Thus, the prosecution history confirms that the patent is not entitled to the filing date of the ’103 provisional.

In litigation, VirtaMove has asserted that the challenged claims are entitled to effective filing dates in September 2002, earlier than either of the provisional applications. To establish the earlier date, VirtaMove would have to show that the '814 patent's inventors conceived of the claimed invention by that date and either reduced it to practice by that date or exercised reasonable diligence thereafter. *Medtronic, Inc. v. Teleflex Innovations S.A.R.L.*, 68 F.4th 1298, 1302-03 (Fed. Cir. 2023). Such an assertion must be shown by independent evidence corroborating an inventor's testimony. *Kolcraft Enters., Inc. v. Graco Children's Prods., Inc.*, 927 F.3d 1320, 1324 (Fed. Cir. 2019). VirtaMove has identified no such evidence.

## **2. Osman**

Osman is prior art under at least 35 U.S.C. §102(b) because it is a printed publication that published no later than December 11, 2002—when it was presented at a major industry conference. (Ex. 1003, cover; Ex. 1015, 1 (citation listing conference date).) The publication was also received by the Kurt F. Wendt Library at the University of Wisconsin-Madison on or before April 14, 2003 and available to library patrons by May 2003—more than a year before the '814 patent's filing date. (Ex. 1015, 1-2.) Even if the '814 patent were entitled to its earliest provisional priority date, Osman would still be prior art under at least 35 U.S.C. §102(a).

### 3. Tucker

Tucker is prior art under at least 35 U.S.C. §102(e) because it is an issued patent filed on January 21, 2004. (Ex. 1004.) Further, Tucker is entitled to the priority date of its own provisional application (the “Tucker Provisional”). The Tucker Provisional was filed May 9, 2003 (Ex. 1005). Tucker was filed within one year of the Tucker Provisional’s filing, names at least one inventor in common, and specifically references the Tucker Provisional. (Ex. 1004, 1, 1:6-10; Ex. 1005.)

As shown below, the Tucker Provisional fully supports claim 1 of Tucker. Thus, Tucker is entitled to the Tucker Provisional’s May 9, 2003 priority date. *Dynamic Drinkware*, 800 F.3d at 1381-82; *Amazon.com, Inc. v. CustomPlay, LLC*, IPR2018-01496, Paper 34, at 44-48 (P.T.A.B. Mar. 4, 2020).

Tucker, Claim 1	Exemplary Support in Tucker Provisional (Ex. 1005)
A method comprising:	1 (“Zones provide a means of virtualizing operating system services[.]”)
establishing a global zone, wherein the global zone is a global operating system environment that can support execution of one or more processes;	2 (“Zone 0, enclosing zones 1-3, is the <i>global</i> zone. Processes running in this zone have (by and large) the same set of privileges available on a Solaris system today[.]”), 3 (“The global zone always exists, and acts as the ‘default’ zone in which all processes run if no zones have been explicitly created by the administrator.”), 4-5, 9, 34-35, Fig. 1.1.
establishing a non-global zone within the global zone, wherein the non-global zone is a partition of the global operating system environment, wherein	1 (“A zone is a ‘sandbox’ within which one or more applications can run without affecting or interacting with the rest of the system.”), 2 (“Zones 1, 2 and

<b>Tucker, Claim 1</b>	<b>Exemplary Support in Tucker Provisional (Ex. 1005)</b>
<p>the non-global zone operates as a separate and distinct operating system environment, and wherein the non-global zone can support execution of one or more processes;</p>	<p>3 are each running a disjoint workload in a sample consolidated environment. ... Each zone can provide an almost arbitrarily rich and customized set of services.”), 2 (“each zone is provided a portion of the file system hierarchy”), 3-5, 9 (“The first step toward bringing up a zone is to create a configuration which specifies various parameters for that zone. Once a configuration is complete, it can be installed onto the system, then it can be booted, logged into, and administered.”), 9-20 (discussing creation and configuration of non-global zones), 25, 33-37, 71-74, Fig. 1.1 (showing non-global zones within the global zone), Fig. 3.1.</p>
<p>isolating a first process executing within the non-global zone to the non-global zone so that the first process does not have visibility or access to processes and objects that are not associated with the non-global zone;</p>	<p>1, 2 (“Basic process isolation is also demonstrated. Each zone is given access to a logical network interface; applications running in distinct zones cannot observe the network traffic of the other[.] ... Finally, each zone is provided a portion of the file system hierarchy. Because each zone is confined to its subtree of the file system hierarchy, the workloads cannot access each other’s on-disk data.”), 3, 4 (“No zone (other than the global zone) should be able to access objects belonging to another zone (including the global zone), either to control or modify those objects in some way, or to simply monitor or read them.”), 4 (“processes in a non-global zone should not be able to interfere with the execution of</p>

Tucker, Claim 1	Exemplary Support in Tucker Provisional (Ex. 1005)
	processes in other zones in the system”), 5, 25-31, 33-42, 68, 71-73, Fig. 1.1.
permitting a second process executing within the global zone to have visibility and access to processes and objects associated with the global zone; and	2 (processes in the global zone “have (by and large) the same set of privileges available on a Solaris system today”), 3-5, 25 (“Processes running in the global zone will still have the full set of privileges, allowing them to affect activity in any zone of the system.”), 33, 34 (“By default, access to system objects from the global zone will be restricted to those objects associated with the global zone.”), 35-36, Fig. 1.1.
permitting the second process executing within the global zone to have access to processes and objects associated with the non-global zone, if the second process has a privilege to cross zone boundaries.	3, 4 (“Appropriately privileged processes in the global zone can access objects associated with other zones.”), 5, 25 (“Processes running in the global zone will still have the full set of privileges, allowing them to affect activity in any zone of the system.”), 33-36, Fig. 1.1.

This Petition cites to the Tucker Provisional in the obviousness ground below (*infra*, §VI.B) because the entirety of the Tucker Provisional was incorporated by reference into Tucker. (Ex. 1004, 1:6-10). This incorporated disclosure qualifies as prior art for the reasons set forth above.

**4. Bandhole**

Bandhole is prior art under at least 35 U.S.C. §102(b) because it published on November 21, 2002. (Ex. 1006.) Even if the '814 patent were entitled to its earliest provisional priority date, Bandhole would be prior art under 35 U.S.C. §102(a). It

is also prior art under §102(e) because it is a published patent application filed January 30, 2002. (*Id.*)

## 5. Gélinas

Gélinas (Ex. 1007) is prior art under at least 35 U.S.C. §102(b) because it is a printed publication that published no later than Aug. 14, 2002. (Ex. 1016 ¶¶4-9.) Gélinas is a set of pages from a public website. (*Id.* ¶9.) The website's address was circulated as part of the VServer project's announcement on the Linux Kernel Mailing List. (Ex. 1016 ¶5; Ex. 1017.)

Subscription to the Linux Kernel Mailing List was common for skilled artisans at the time, so the Gélinas website was disseminated to the interested public via the Mailing List and a POSITA therefore would have been independently aware of and able to locate Gélinas. (Ex. 1016 ¶¶5-7.)

The Mailing List post was also publicized on Slashdot, a popular link aggregation service with numerous daily readers, many of whom were members of the interested public. (Ex. 1016 ¶¶6-7; Ex. 1018.) Thus, Gélinas is a prior-art printed publication. *Hulu, LLC v. Sound View Innovations, LLC*, IPR2018-01039, Paper 29 at 8-12 (P.T.A.B. Dec. 20, 2019) (precedential); *Voter Verified, Inc. v. Premier Election Sols., Inc.*, 698 F.3d 1374 (Fed. Cir. 2012); *Keysight Techs., Inc. v. Centripetal Networks, LLC*, IPR2022-01525, Paper 27 (P.T.A.B. Apr. 15, 2024).

Gélinas consists of four pages from the VServer website—a main page titled “Virtual private servers and security contexts” and three pages in a “Howto/FAQ” section. (Ex. 1016 ¶¶ 8-9.) The main page includes a hyperlink to the Howto/FAQ section. (Ex. 1007 at 28 (“A FAQ can be found at [link]”).) Because all four pages were on the same website (hosted at [www.solucorp.qc.ca](http://www.solucorp.qc.ca)) discussing the same software (VServer), they constitute a single publication. *Halliburton Energy Servs., Inc. v. Dynamic 3D Geosolutions LLC*, IPR 2014-01186, 2015 WL 5565065, \*10 (PTAB Dec. 15, 2015). Regardless, even if the Gélinas pages were separate references, it would have been obvious to combine their teachings because a POSITA “would have been interested in reading the [Gélinas] pages, all collected at a single location, to obtain an understanding of a single coherent system, the [VServer] system.” *Id.*

#### **IV. LEVEL OF ORDINARY SKILL**

Based on the relevant factors, *In re GPAC Inc.*, 57 F.3d 1573, 1579 (Fed. Cir. 1995), a person of ordinary skill in the art (“POSITA”) would have had a minimum of a bachelor’s degree in computer science, computer engineering, software engineering, or a similar field, and approximately two years of industry or academic experience in a field related to operating system design. (Ex. 1002 ¶31.) Work experience could substitute for formal education and additional formal education could substitute for work experience. (*Id.*)

## V. CLAIM CONSTRUCTION

Petitioner submits that no claim terms require construction to resolve the obviousness challenges here; however, this is not a waiver of potential claim construction arguments. *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co. Ltd.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017); *Vivid Techs., Inc. v. Am. Sci. & Eng'g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999).

### A. “Container” and “System Files”

Petitioner notes that the terms “container” and “system files,” which are used in the claims, are defined in the patent’s specification. (Ex. 1001, 2:16-3:19.) No explicit construction for these terms is necessary because the prior art relied on herein discloses the claims under the definitions in the specification and the plain meaning of the claims. (Ex. 1002 ¶48.)

### B. “Disparate Computing Environments”

Petitioner notes that in a related district court case, *VirtaMove, Corp. v. Amazon.com, Inc.*, No. 7:24-cv-00030 (“Amazon Litigation”), Defendant (“Amazon”) asserted that the phrase “disparate computing environments,” recited in claim 1, is indefinite. The patent defines “disparate computing environments” as “Environments where computers are *stand-alone* or where there are plural computers and where they are *unrelated*.” (Ex. 1001, 2:17-19 (emphases added).) Defendant argued in the Amazon Litigation that “unrelated” is subjective. (Ex. 1019, 3-4.) However, PO argued that the district court need not resolve the

meaning of “unrelated” because, in the context of claim 1, the computers cannot be “unrelated.” (Ex. 1020, at 3.) In PO’s view, “disparate computing environments” would thus be limited to “environments run by ... standalone computers.” (*Id.* at 4.)

PO further asserted that a computer can be “standalone” even if connected to other computers in a larger system, so long as the computer is “independently operable.” (*Id.*) The prior art here discloses “disparate computing environments” under PO’s interpretation. (Ex. 1002 ¶50; *infra* §VI.A.1.c, §VI.B.1.c, §VI.C.1.c.)

Another potential interpretation of “disparate computing environments” was suggested by the patent’s named inventor, Donn Rochette. In his deposition in the Amazon litigation, Rochette testified that the patent’s definition of “disparate computing environments” would cover computers that each “have a different network address,” which “would make them standalone and separate.” (Ex. 1021, 71:17-19, 73:18-74:18.) The prior art discloses “disparate computing environments” under Rochette’s interpretation also. (Ex. 1002 ¶51; *infra* §VI.A.1.c, §VI.B.1.c, §VI.C.1.c.)

Because the prior art discloses the claim limitations under any available interpretation, no limitations require express construction here. *Intel Corp. v. Qualcomm Inc.*, 21 F.4th 801, 812-13 (Fed. Cir. 2021).

## VI. GROUNDS OF UNPATENTABILITY

### A. Ground 1: Claims 1, 2, 4, 6, 8-10, and 13-14 are Unpatentable as Obvious in View of Osman.

#### 1. Claim 1

##### a. Limitation 1[pre][i]: “In a system having a plurality of servers”

Osman discloses that its system is designed for an environment where “compute machines [plural] typically run completely independently of one another” and applications are migrated from one machine to another. (Ex. 1003, 363.) Osman’s “compute machines” include “servers, where application processing takes place.” (*Id.*) Osman also discloses multiple “network file servers” for storing “applications and user data.” (*Id.*) Thus, Osman discloses this limitation. (Ex. 1002 ¶54.)

##### b. Limitation 1[pre][ii]: “with operating systems that differ”

Osman discloses that its system was “successfully used ... to migrate pods across Linux kernels with minor version differences.” (Ex. 1003, 372.) The Linux kernel is the core of the Linux operating system, so a POSITA would understand from this disclosure that Osman’s system was used with Linux operating systems having minor kernel-version differences. (Ex. 1002 ¶55.) Thus, Osman discloses this limitation. (Ex. 1002 ¶55.)

Additionally, Osman discloses that its system lacks “dependencies on the host operating system.” (Ex. 1003, 361 (system “should not necessitate use of new operating systems or substantial modifications to existing ones”), 362 (discussing system’s “compatib[ility] with existing operating systems”—plural), 363, 373.) Based on this express teaching, a POSITA would have found it obvious to use Osman’s system with other operating systems in addition to Linux. (Ex. 1002 ¶56.)

**c. Limitation 1[pre][iii]: “operating in disparate computing environments”**

As discussed above, Osman discloses that its system is used in environments where “compute machines typically run completely independently of one another, each running its own independent operating system.” (Ex. 1003, 363.) This disclosure satisfies the claim limitation under PO’s interpretation, which requires “independently operable” computers. (*Supra*, §V.B.)

Osman also discloses that pods running on different host machines have different network addresses. (Ex. 1003, 369 (“external IP address ... changes as the given pod moves from host to host”).) This disclosure satisfies the claim limitation under the inventor’s interpretation, which is satisfied by separate computers that each “have a different network address.” (*Supra*, §V.B.)

Thus, Osman discloses this limitation. (Ex. 1002 ¶59.)

**d. Limitation 1[pre][iv]: “wherein each server includes a processor and an operating system including a kernel”**

As discussed above, Osman discloses servers that each run their own independent operating system. (Ex. 1003, 363.) A POSITA would have understood that each such operating system included a kernel—which is the core of an operating system. (Ex. 1002 ¶60; *see also* Ex. 1003, 362 (Osman’s system is “implemented . . . as a loadable kernel module”).) Furthermore, Osman discloses implementing its system on servers with “933 MHz Intel Pentium-III” processors. (Ex. 1003, 372.) Thus, Osman discloses this limitation. (Ex. 1002 ¶60.)

**e. Limitation 1[pre][v]: “a set of associated local system files compatible with the processor”**

The ’814 patent defines “System files” as “files provided within an operating system and which are available to applications as shared libraries and configuration files.” (Ex. 1001, 2:52-54.) As an example, the ’814 patent lists “configuration files” from the “Apache” web server. (*Id.*, 3:6-17.) Osman discloses that pods may contain such configuration files. (Ex. 1003, 367 (describing example where “each pod maintains its own configuration” used to run “apache”).) These configuration files are in “private pod directories” that can be stored “locally on the host machine.” (*Id.*) Thus, Osman discloses this limitation. (Ex. 1002 ¶61.)

**f. Limitation 1[pre][vi]: “a method of providing at least some of the servers in the system with secure, executable applications related to a service”**

Osman discloses executing eleven different applications related to a “VNC” service. (Ex. 1003, 374.) VNC is a service used to connect to a remote computer and run graphical applications such as a “web browser” or “PDF viewer” in a pod. (*Id.*) As another example, Osman discloses executing “Apache” in a pod. (*Id.*) The ’814 patent admits that Apache is an application related to a service. (Ex. 1001, 4:26-31.)

Further, applications in Osman’s pods are secure because only authorized users “are allowed to manipulate the pod[.]” (Ex. 1003, 371.) Thus, Osman discloses this limitation. (Ex. 1002 ¶63.)

**g. Limitation 1[pre][vii]: “wherein the applications are executed in a secure environment”**

Osman discloses that its “processes are created inside of a pod and spend their entire lifetimes in the context of that pod; they are not allowed to leave” the pod. (Ex. 1003, 364.) “Other processes outside a pod ... are therefore not able to interact with processes inside a pod” as they otherwise would. (*Id.*) Because application processes are isolated in Osman’s pods, the applications are executed in a secure environment. (Ex. 1002 ¶64; *see also infra*, §VI.A.1.p, §VI.A.1.q.) Thus, Osman discloses this limitation. (*Id.* ¶64.)

**h. Limitation 1[pre][viii]: “wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service”**

Osman discloses that its applications perform tasks related to a service. (Ex. 1003, 374 (listing tasks performed by Apache and VNC applications).) For example, Osman discloses that its system may run a web server, which a POSITA would have understood necessarily includes at least one object executable by the server’s operating system for performing a task related to web hosting. (*Id.*, 367, 373 (web server delivered “54 web pages”); Ex. 1002 ¶65.) Thus, Osman discloses this limitation. (Ex. 1002 ¶65.)

**i. Limitation 1[a][i] “the method comprising: storing in memory accessible to at least some of the servers a plurality of secure containers of application software”**

The ’814 patent defines a “container” as an “aggregate of files required to successfully execute a set of software applications on a computing platform[.]” (Ex. 1001, 2:23-25.) Osman’s pods are containers because they comprise a set of files needed to execute the applications running in the pod. (Ex. 1003, 367-68.)

Osman discloses, for example, that when a pod is to be migrated from one place to another, the system “saves and restores the information it needs to reconstruct the pod virtual file system, including a list of all files opened by the processes within a pod and the access rights with which the files were opened.” (*Id.*, 368.) For any files that are not already saved in network file storage, the system

“saves the contents of the file and recreates the file when the pod is restarted.” (*Id.*, 365, 368; Ex. 1002 ¶67.)

Osman also discloses storing multiple pods together in memory accessible to one or more servers. (Ex. 1003, 367-68 (on “network file server” and local host’s file system), 372, 374.) Thus, Osman discloses storing a “plurality” of containers. (Ex. 1002 ¶68.)

In addition to the basic definition of “container” addressed above, the ’814 patent further describes the attributes of containers as follows:

Each container for use on a server is mutually exclusive of the other containers, such that read/write files within a container cannot be shared with other containers. . . . A container comprises one or more application programs including one or more processes, and associated system files for use in executing the one or more processes; but containers do not comprise a kernel; each container has its own execution file associated therewith for starting one or more applications. In operation, each container utilizes a kernel resident on the server that is part of the operating system (OS) the container is running under to execute its applications.

(Ex. 1001, 2:29-42.) To the extent “container” is construed to require any of these attributes, Osman discloses them. For example, Osman discloses that files within a container cannot be shared (unless the container is specifically configured to allow sharing). (Ex. 1003, 367; *infra*, §VI.A.1.p, §VI.A.7.) Osman also discloses the other attributes of containers, as explained below in connection with other claim

limitations. (*Infra*, §VI.A.1.j (applications and system files in containers); §VI.A.1.m (no kernel in containers); §VI.A.2 (execution file in containers); §VI.A.1.k (containers use server’s kernel).) Thus, Osman discloses this limitation under any construction of “container.” (Ex. 1002 ¶¶66-68.)

**j. Limitation 1[a][ii]: “each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications”**

Osman discloses that each pod (container) may comprise one or more applications. For example, Osman describes testing that involved two pods—one for VNC and one for Apache. (Ex. 1003, 374.) The VNC pod contained eleven applications and the Apache pod contained one application. (*Id.*) Apache relies on system files as explained above. (*Supra*, §VI.A.1.e.) Additionally, a POSITA would understand that applications in the VNC pod relied on shared libraries and configuration files, both of which are “system files” under the ’814 patent’s definition (Ex. 1001, 2:52-54), because shared libraries and configuration files were routinely used in complex applications like the ones installed in the VNC pod. (Ex. 1002 ¶¶69.) Accordingly, Osman discloses this limitation. (*Id.*)

**k. Limitation 1[a][iii]: “for use with a local kernel residing permanently on one of the servers”**

Osman states that its system “can be dynamically loaded on a running Linux system to provide ... functionality without modifying, recompiling, or reinstalling

the Linux kernel.” (Ex. 1003, 372.) Osman accomplishes this by “leveraging the loadable kernel module interface available in many commodity operating systems” and implementing its system “as a Linux kernel module.” (*Id.*) Thus, Osman discloses this limitation. (Ex. 1002 ¶70.)

**l. Limitation 1[a][iv]: “wherein the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems”**

Osman discloses that its system has been used “to migrate pods across Linux kernels with minor version differences.” (Ex. 1003, 372.) A POSITA would have understood this to mean that associated system files within the pods were compatible with local kernels of at least several different Linux operating systems. (Ex. 1002 ¶71.) More specifically, the system files in Osman’s pods were compatible with the Linux kernels in the operating systems running on both the source computer and the destination computer. (*Id.*) Otherwise, the migration of the pods would not have been successful because the applications in the pod could not have been running both before and after the migration. (*Id.*) Thus, Osman discloses, or at least suggests, this limitation. (*Id.*)

**m. Limitation 1[a][v]: “the containers of application software excluding a kernel”**

As discussed above, Osman discloses that its system makes use of a local kernel. (*Supra* §VI.A.1.k.) Further, Osman discloses that its system may “migrate

Pods across Linux kernels.” (Ex. 1003, 372.) Such migration is accomplished by saving information that the pod was using to run on a host with a first kernel and loading that information when the pod is restarted on another host with a different kernel. (*Id.*) Thus, a POSITA would understand that the pod (container) does not include the kernel. (Ex. 1002 ¶72.) Accordingly, Osman discloses this limitation. (*Id.*)

**n. Limitation 1[a][vi]: “wherein some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server”**

Osman discloses an example in which one copy of the Apache web server runs inside a pod and another copy runs outside the pod on the same computer. (Ex. 1003, 373.) As explained above, Apache uses system files to control its configuration. (*Supra*, §VI.A.1.e.) When one copy of Apache is inside a pod, that copy uses the system files inside the pod instead of the system files outside the pod—because system files outside of the pod are inaccessible from inside the pod. (Ex. 1002 ¶73; Ex. 1003, 367.) Thus, Osman discloses this limitation. (Ex. 1002 ¶73.)

- o. Limitation 1[a][vii]: “wherein said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server”**

When copies of Apache run both inside and outside of a pod (*supra*, §VI.A.1.n), the Apache configuration files inside the pod are copies or modified copies of the Apache configuration files outside the pod. (Ex. 1002 ¶74; Ex. 1003, 367, 374.) Thus, Osman discloses this limitation. (Ex. 1002 ¶74.)

- p. Limitation 1[a][viii]: “wherein the application software cannot be shared between the plurality of secure containers of application software”**

Osman discloses creating a private directory for each pod and ensuring that “this directory is not accessible by processes on the host machine that are not in the given pod.” (Ex. 1003, 367.) These private directories prevent applications running in different pods from interfering with each other. (*Id.*)

Osman also discloses that pods *may* be configured to share application software in a common location. (*Id.*) But absent such configuration, a POSITA would understand that application software in one pod would not be shared with other pods because Osman’s system would “prevent processes within a pod from breaking out of their virtual file system environment.” (*Id.*; Ex. 1002 ¶76.) Osman also discloses that processes in a pod are prevented from sharing memory with processes outside the pod. (Ex. 1003, 364.) Thus, Osman discloses this limitation. (Ex. 1002 ¶76.)

**q. Limitation 1[a][ix]: “wherein each of the containers has a unique root file system that is different from an operating system’s root file system”**

Osman discloses that when a pod is created, the system “then uses the **chroot** call to set the staging area [for the pod’s virtual file system] as the root directory for the pod,” thereby giving the pod its own root file system separate from the operating system’s root file system. (Ex. 1003, 367; Ex. 1002 ¶77.) Thus, Osman discloses this limitation.

For the reasons above, Osman renders claim 1 obvious. (Ex. 1002 ¶¶54-78.)

**2. Claim 2: “wherein each container has an execution file associated therewith for starting the one or more applications”**

Claim 2 depends from claim 1. Osman discloses claim 2’s additional limitation because its pods include an execution file that sets “what applications if any should be launched once the pod is created[.]” (Ex. 1003, 371.) Thus, Osman renders claim 2 obvious. (Ex. 1002 ¶79.)

**3. Claim 4: “pre-identifying applications and system files required for association with the one or more containers prior to said storing step”**

Claim 4 depends from claim 1. Osman discloses claim 4’s additional limitation because applications and system files associated with a pod are pre-identified when a pod is suspended on a first computer before being migrated and thus stored in memory a second computer. (Ex. 1002 ¶81.) Specifically, when

a pod is to be suspended, Osman’s system saves “the information it needs to reconstruct the pod virtual file system, including a list of all files opened by the processes within a pod.” (Ex. 1003, 368.) The “processes within a pod” include any applications running in the pod and the “files opened by the processes” include the system files that control the configuration of such applications. (Ex. 1002 ¶81.) This information is identified on a first computer before the pod “migrates,” and is therefore stored in memory on a second computer. (Ex. 1003, 368.) Thus, Osman renders claim 4 obvious. (Ex. 1002 ¶81.)

**4. Claim 6: “assigning a unique associated identity to each of a plurality of the containers, wherein the identity includes at least one of IP address, host name, and MAC address”**

Claim 6 depends from claim 2. Osman discloses the additional limitation of claim 6 because each pod is assigned a “[p]er-pod IP address.” (Ex. 1003, 365; *id.*, 369.) Thus, Osman renders claim 6 obvious. (Ex. 1002 ¶83.)

**5. Claim 8: “wherein the one or more applications and associated system files are retrieved from a computer system having a plurality of secure containers”**

Claim 8 depends from claim 1. Osman discloses claim 8’s further limitation because its system “can use network file servers to support many pods running on many machines at the same time.” (Ex. 1003, 367.) The network file server can store applications and other “common executables” used by the pods. (*Id.*) The network file server can also host the “private pod directory” for each pod, which

stores system files (such as Apache configuration files). (*Id.*) The network file server (computer system) can also store all other information associated with a pod. (*Id.* at 365 (pods can be “checkpointed, suspended to secondary storage, and stored for future use”), 367 (“checkpointed data” can be stored on a network file server).) Thus, Osman renders claim 8 obvious. (Ex. 1002 ¶85.)

**6. Claim 9: “wherein server information related to hardware resource usage including at least one of CPU memory, network bandwidth, and disk allocation is associated with at least some of the containers prior to the applications within the containers being executed”**

Claim 9 depends from claim 2. The claim contains a typographical error; it is missing a comma between “CPU” and “memory.” (*Compare* Ex. 1001, claim 9 *with id.*, 8:47-51 (“CPU, memory, network bandwidth and disk usage”), Fig. 5 (listing “CPU” and “memory” separately).)

Osman discloses the additional limitation of claim 9 because, when a pod is migrated, its host saves “all process states, including memory, CPU registers, open file handles, etc.” (Ex. 1003, 365.) Accordingly, information relating to hardware resource usage, including at least CPU and memory usage, is associated with a pod before the pod is restarted. (Ex. 1002 ¶87.) And when the pod restarts, applications in it are executed. (Ex. 1003, 371, 374.)

Additionally, Osman discloses “mounting various NFS mount points from a file server for pods ... on the local machine.” (Ex. 1003, 372.) Based on this

disclosure a POSITA would understand that the system associates disk allocation information with a pod before starting the pod and executing applications. (Ex. 1002 ¶88.)

Finally, Osman discloses experiments in which the disk allocation of particular pods was measured (e.g., “23 MB of image data”) before such pods were started and applications executed. (Ex. 1003, 374-375; Ex. 1002 ¶89.)

Thus, Osman renders claim 9 obvious. (Ex. 1002 ¶90.)

**7. Claim 10: “wherein in operation when an application residing within a container is executed, said application has no access to system files or applications in other containers or to system files within the operating system during execution thereof”**

Claim 10 depends from claim 2. Osman discloses the additional limitation of claim 10 because its system “uses the chroot call to set the [pod-specific] staging area as the root directory for the pod,” which “prevent[s] processes with a pod from breaking out of their virtual file system environment.” (Ex. 1003, 367.) This call prevents processes running in the pod from accessing system files or applications in other pods or outside of the pods (unless the administrator specifically configured the pods to allow such outside access). (Ex. 1002 ¶91; *supra*, §VI.A.1.p.) Thus, Osman renders claim 10 obvious. (*Id.* ¶91.)

**8. Claim 13: “associating with a plurality of containers a stored history of when processes related to applications within the container are executed for at least one of, tracking statistics, resource allocation, and for monitoring the status of the application”**

Claim 13 depends from claim 1. Osman discloses the additional limitation of claim 13 because Osman’s system maintains a “hash table” that contains identifiers for every process running in a pod. (Ex. 1003, 365-66.) The system uses the information in this table to allocate “process resources” such as “shared memory.” (*Id.*)

The hash table is also used to monitor the status of applications—for example, to determine whether an application is within a pod. (*Id.*; Ex. 1002 ¶95.) By monitoring this information, the system prevents requests originating outside a pod from interfering with applications inside a pod. (*Id.* at 366 (“Outside of the pod environment, the system call will reject any requests to physical resources which belong to pods”).) Thus, Osman renders claim 13 obvious. (*Id.* ¶95.)

**9. Claim 14**

Claim 14 depends from claim 1. Osman discloses each further limitation of claim 14. (Ex. 1002 ¶¶96-101.)

- a. **Limitation 14[a][i]: “creating containers prior to said step of storing containers in memory, wherein containers are created by:”**

Osman discloses pods running on a computer before they are suspended and stored, e.g., in network file storage. (Ex. 1003, 365-371; *supra* §VI.A.3, §VI.A.5.)

The pods are created as follows.

- b. **Limitation 14[a][ii]: “a) running an instance of a service on a server”**

Osman discloses executing applications in a pod to provide a service. (*Supra* §VI.A.1.f.) For example, Osman discloses executing an Apache web server. (*Id.*)

- c. **Limitation 14[a][iii]: “b) determining which files are being used”**

Osman discloses that when a pod is suspended, the system saves “a list of all files opened by the processes within a pod.” (Ex. 1003, 368.)

- d. **Limitation 14[a][iv]: “c) copying applications and associated system files to memory without overwriting the associated system files so as to provide a second instance of the applications and associated system files”**

Osman discloses an example in which a first copy of Apache runs outside a pod and a second copy of Apache runs inside a pod. (*Supra*, §VI.A.1.n, §VI.A.1.o.)

In this example, a POSITA would understand that the Apache application and its associated system files were copied into the pod without overwriting the copy of Apache outside the pod. (Ex. 1002 ¶100.) Osman also discloses migrating the

Apache pod. (Ex. 1003, 374-75.) And Osman discloses other applications with associated system files operating both inside and outside of a pod. (*Id.*, 373.)

Thus, Osman renders claim 14 obvious. (Ex. 1002 ¶101.)

**B. Ground 2: Claims 1, 2, 4, 6, 8-10, and 13 are Unpatentable as Obvious in View of Tucker and Bandhole.**

Tucker discloses “zones,” which are application containers that Sun Microsystems implemented in its Solaris operating system. (*Supra*, §I.C.2.) Tucker focuses on how zones work on a single server, with limited discussion of how that single server might be used with other servers. (Ex. 1005.)<sup>1</sup> But a POSITA would have known at the time that Solaris servers were routinely networked with other servers to provide a wide array of services. (Ex. 1002 ¶132.)

Bandhole describes a system in which servers and other computing resources are combined to provide various services. (Ex. 1006.) Several examples in Bandhole involve a Solaris server combined with another server running a different operating system. (*Id.* ¶¶[0005],[0034].)

In one of Bandhole’s examples, the Solaris server runs both “custom application server software” and “Oracle database software” to “provide a search service on the Web.” (*Id.* ¶[0005].) A separate server using the Linux operating

---

<sup>1</sup> Citations to the Tucker Provisional (Ex. 1005) are also citations to Tucker, which incorporates the entire Tucker Provisional by reference (Ex. 1004, 1:6-10).

system runs “Apache web server software” as part of the same search service. (*Id.*) This example appears in Bandhole’s background section, indicating that architectures like this were well known before Bandhole’s 2002 filing date. (*Id.*)

A POSITA would have been motivated to use Tucker’s zones to implement the types of services described in Bandhole’s background section. (Ex. 1002 ¶135.) For example, a POSITA would have found it obvious to run Bandhole’s application server and database using two separate zones in accordance with Tucker. (*Id.* ¶135.) This approach would have been obvious for several reasons.

First, Tucker discloses that zones “limit[] the damage possible in the event of a security violation.” (Ex. 1005, 1.) Zones “allow the deployment of multiple applications on the same machine,” even when those applications have differing requirements. (*Id.*) Zones also “can be useful to support rapid deployment (and redeployment) of applications.” (*Id.*, 2.) A POSITA would have been motivated to use zones to run Bandhole’s services to gain the security and deployment benefits that Tucker discloses. (Ex. 1002 ¶136.)

Second, the combination is a simple addition of one known element (Tucker’s zones) to another known element (Bandhole’s software) to obtain predictable results (software running in zones). (Ex. 1002 ¶137.) Moreover, the combination uses a known technique (zones on Solaris) to improve a similar device and method (Bandhole’s Solaris server and search service) in the same way. (*Id.*) The

combination also applies a known technique (zones on Solaris) to a known device and method (Bandhole's Solaris server and search service) that is ready for improvement and yields predictable results. (*Id.*)

Finally, because Bandhole's software was compatible with Solaris (Ex. 1006 ¶[0005]) and Tucker's zones ran on Solaris (*supra*, §I.C.2.), a POSITA would have had a reasonable expectation of success in making the combination. (Ex. 1002 ¶138.)

## 1. Claim 1

### a. Limitation 1[pre][i]: "In a system having a plurality of servers"

The '814 patent states that the terms "computing platform, server and computer are used interchangeably throughout this specification." (Ex. 1001, 12:5-7; *see also, e.g., id.* at 10:56-58 ("an existing server, for example a computer").)

Tucker discloses that "the same application environment can be maintained on different physical machines" (plural), to support rapid deployment of applications across such machines. (Ex. 1005, 2.) Thus, Tucker discloses this limitation. (Ex. 1002 ¶140.)

Bandhole also discloses a plurality of servers: "a Linux server running Apache web server software" and "a Solaris server running a custom application server software and Oracle database software." (Ex. 1006 ¶[0005]; Ex. 1002 ¶141.)

**b. Limitation 1[pre][ii]: “with operating systems that differ”**

Bandhole’s system includes a Linux server and a Solaris server. (Ex. 1006 ¶¶0005.) Linux and Solaris are different operating systems. (Ex. 1002 ¶142.)

**c. Limitation 1[pre][iii]: “operating in disparate computing environments”**

Tucker and Bandhole disclose this limitation under both VirtaMove’s interpretation (requiring “independently operable” computers) and the inventor’s interpretation (requiring computers with different network addresses). (*Supra*, §V.B.)

First, Tucker discloses implementing zones on computers running Solaris. (Ex. 1005, 1-3.) Computers running Solaris were independently operable because one computer running Solaris could operate without being controlled by any other computer. (Ex. 1002 ¶144.) Tucker shows an example of this in Figure 1.1, where zones run on a single Solaris computer operating independently. (Ex. 1005, 1-3.)

Second, Tucker discloses implementing zones on computers with different network addresses. (Ex. 1002 ¶145.) For example, Tucker discloses that a system running zones would have its own “primary IP address.” (Ex. 1005, 1-2.) Further, different zones within the system could have “distinct IP addresses.” (*Id.*)

Additionally, Bandhole discloses independently operable computers with different network addresses. For example, Bandhole discloses one server running

Linux and another running Solaris, communicating with each other over an “Ethernet LAN.” (Ex. 1006 ¶[0005].) The Linux server and Solaris server are independently operable because they use distinct operating systems that do not control each other. (Ex. 1002 ¶146.) Additionally, a POSITA would understand that the servers use different network addresses because they communicate with each other over the disclosed Ethernet LAN. (*Id.*)

**d. Limitation 1[pre][iv]: “wherein each server includes a processor and an operating system including a kernel”**

Tucker discloses that a server running zones on the Solaris operating system uses one processor or multiple processors. (Ex. 1005, 2.) Further, Tucker discloses that Solaris includes a kernel. (*E.g., id.*, 4, 10, 21.)

Bandhole’s servers run either Solaris or Linux, both of which are operating systems that include a kernel. (Ex. 1006 ¶[0005]; Ex. 1002 ¶148.) Thus, Bandhole and Tucker each disclose this limitation. (Ex. 1002 ¶148.)

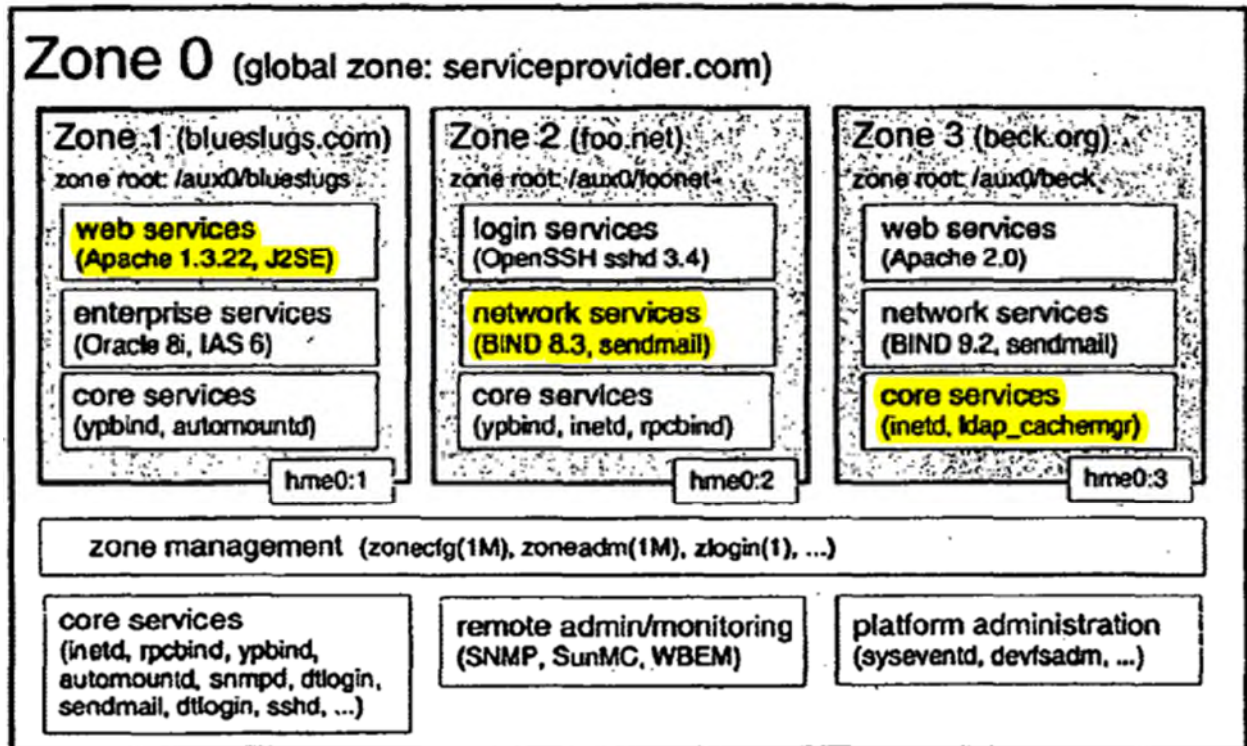
**e. Limitation 1[pre][v]: “a set of associated local system files compatible with the processor”**

The ’814 patent defines system files as including “shared libraries” and “configuration files.” (Ex. 1001, 2:52-54.) Tucker discloses both. (Ex. 1005, 43-44, 71-72.) Additionally, the ’814 patent admits that Apache running on Linux used both shared libraries and configuration files. (Ex. 1001, 2:55-3:19.) Bandhole discloses the use of Apache on Linux. (Ex. 1006 ¶¶[0005],[0034].) A POSITA

would have understood that these system files are compatible with the processor used to execute them. (Ex. 1002 ¶149.)

f. **Limitation 1[pre][vi]: “a method of providing at least some of the servers in the system with secure, executable applications related to a service”**

Tucker discloses that its zones feature provides a way in which “one or more applications can run without affecting or interacting with the rest of the system,” thereby providing “Security” and “Isolation” benefits. (Ex. 1005, 1.) Tucker further discloses that such applications may relate to services, including web services, enterprise services, login services, network services, and the like:



(Id., 3; Ex. 1002 ¶150.)

**g. Limitation 1[pre][vii]: “wherein the applications are executed in a secure environment”**

Tucker also discloses that “applications running in distinct zones cannot observe the network traffic of the other” and “cannot access each other’s on-disk data.” (Ex. 1005, 1-2.) Thus, zones are a secure environment. (Ex. 1002 ¶151.)

**h. Limitation 1[pre][viii]: “wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service”**

Tucker discloses many applications that run in zones, including an Apache web server, Oracle database, OpenSSH login, and email server. (Ex. 1005, 3, Fig. 1.1.) These applications each included at least one object executable by the operating system for performing a task related to the types of services listed in Tucker’s Figure 1.1. (Ex. 1002 ¶152.)

**i. Limitation 1[a][i] “the method comprising: storing in memory accessible to at least some of the servers a plurality of secure containers of application software”**

Tucker discloses a four-step process for installing and configuring zones. (Ex. 1005, 9-10.) First, the zone is “configured,” meaning “a zone’s configuration has been completely specified and committed to stable storage.” (*Id.*, 9.) Second, the zone is “installed,” meaning “a zone’s configuration has been laid out on the system: packages [e.g., application software] have been installed under the zone’s root path.” (*Id.*; Ex. 1002 ¶153.) Third, the zone is “ready,” meaning “the kernel has created the zone, ... file systems have been mounted, ... but no processes have

been started.” (Ex. 1005, 10.) Finally, the zone is “running,” meaning “processes have been started.” (*Id.*) Thus, Tucker discloses (in at least the zone-installing step) storing a container as this limitation requires. (Ex. 1002 ¶153; *supra*, §I.C.2 (non-global zones are containers).) Tucker also discloses that a plurality of containers are stored together. (*E.g., id.*, 2-3 (three non-global zones).)

**j. Limitation 1[a][ii]: “each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications”**

Tucker discloses that its zones run various applications and use various system files. (*Supra* §VI.B.1.f, §VI.B.1.e.) Tucker also discloses an approach in which “all packages required for a zone that make up a particular metacluster will be installed as well as any other packages selected by the global administrator.” (Ex. 1005, 72.) This approach “will provide the maximum configurability” because “all of the required and any selected optional Solaris packages” are installed within the zone. (*Id.*) A POSITA would understand that, in this arrangement, the applications and system files used by a zone are installed into the zone using the Solaris packages. (Ex. 1002 ¶154 (explaining that Solaris packages were the standard way of installing applications and their system files in Solaris).)

**k. Limitation 1[a][iii]: “for use with a local kernel residing permanently on one of the servers”**

Tucker discloses that zones are created and managed by the kernel. (Ex. 1005, 10-11, 4; Ex. 1002 ¶155.) The kernel resides permanently on the servers because it remains in place as zones are added or removed. (Ex. 1005, 9-15; Ex. 1002 ¶155.)

**l. Limitation 1[a][iv]: “wherein the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems”**

Because applications installed in zones are executed by the local kernel (*supra* §VI.B.1.k; Ex. 1005, 4, 74.), a POSITA would have understood that the associated system files (e.g., libraries and configuration files required to execute the applications) are likewise compatible with the local kernel. (Ex. 1002 ¶156.)

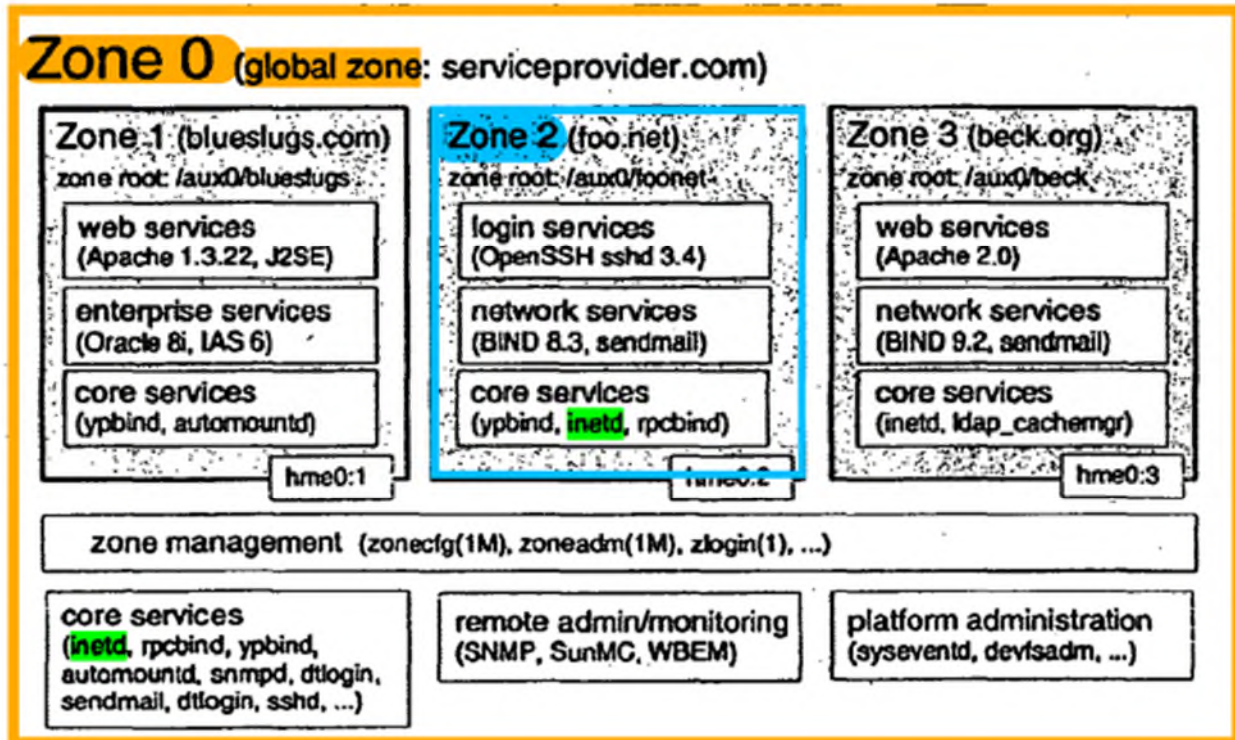
**m. Limitation 1[a][v]: “the containers of application software excluding a kernel”**

Tucker discloses that the operating system’s single kernel keeps track of multiple non-global zones. (Ex. 1005, 11, 21.) The kernel also creates each of these zones. (*Id.*, 10.) Based on this disclosure, a POSITA would have understood that the kernel is outside of the non-global zones. (Ex. 1002 ¶157.)

**n. Limitation 1[a][vi]: “wherein some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server”**

Tucker discloses that various applications in non-global zones (containers) may also be executed in the global zone (underlying server). For example, Tucker

discloses that the service “**inetd**” may be executed in both a **non-global zone** (container) and a **global zone** (underlying server):



(Ex. 1005, 3.) Tucker further discloses that the “**inetd**” service uses a configuration file (system file), and different configuration files may be used for **inetd** in each zone. (*Id.* at 43-44 (“each zone can, for example, run its own `inetd(1M)` with a full configuration file”).) Thus, **inetd** in the non-global zone uses the system file from its own zone in place of the **inetd** system file in the global zone. (Ex. 1002 ¶158.) The same is true for other applications like “**ypbind**” and “**sendmail**,” which use system files and are installed in both the global and non-global zones. (*Id.*)

- o. Limitation 1[a][vii]: “wherein said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server”**

As discussed above, Tucker discloses that the underlying server and one or more non-global zones may run the same application, utilizing the non-global zone’s associated system files in place of the server’s associated local system files. (*Supra* §VI.B.1.n.) Because the same application is running in both locations, at least some of the associated system files (configuration files or shared libraries) are copies or modified copies of the associated local system files on the underlying server. (Ex. 1002 ¶159.) Indeed, Tucker confirms that one zone will sometimes duplicate system files from other zones leading to a “heavier disk footprint.” (Ex. 1005, 72.)

Tucker also discloses that “different versions of the same application may be run without negative consequence in different zones.” (Ex. 1005, 2.) Based on this disclosure, a POSITA would find it obvious that at least some system files in the non-global zone (e.g., files associated with a newer application version) would be modified copies of system files in the global zone (e.g., files associated with an older application version). (Ex. 1002 ¶160.)

- p. Limitation 1[a][viii]: “wherein the application software cannot be shared between the plurality of secure containers of application software”**

Tucker discloses that “[v]irtualization of storage in a zone is achieved via a restricted root.... Processes running within a zone will be limited to files and file

systems that can be accessed from the restricted root.” (Ex. 1005, 37; *id.*, 71 (“zones require a separate root file systems [sic] for isolation”).) Thus, application software located in one zone cannot be shared with other zones, which will be unable to access that application’s files. (*Id.*; Ex. 1002 ¶161.)<sup>2</sup>

**q. Limitation 1[a][ix]: “wherein each of the containers has a unique root file system that is different from an operating system’s root file system”**

Tucker discloses that “zones require a separate root file systems [sic] for isolation.” (Ex. 1005, 71; *see also id.*, 37; *supra*, §VI.B.1.p; Ex. 1002 ¶162.)

For the reasons above, claim 1 would have been obvious to a POSITA in view of Tucker and Bandhole. (Ex. 1002 ¶163.)

**2. Claim 2: “wherein each container has an execution file associated therewith for starting the one or more applications”**

Tucker discloses the further limitation of claim 2 because “[a]pplications that are installed into a zone and deliver start/stop scripts into /etc/init.d and one or more of the /etc/rc?.d directories will be started when the zone is booted.” (Ex. 1005, 77; *see also id.*, 14.) The “start/stop scripts” in each zone are execution files because they cause applications to execute. (Ex. 1002 ¶164.)

---

<sup>2</sup> Zones may be able to share files if they are specifically configured to do so, but absent such configuration no sharing can occur. (Ex. 1005, 72.)

**3. Claim 4: “pre-identifying applications and system files required for association with the one or more containers prior to said storing step”**

Tucker discloses the further limitation of claim 4 because when a zone is installed using a particular command, the system “will prompt for certain options including ... packages to install.” (Ex. 1005, 74.) Accordingly, the server pre-identifies (e.g., by prompting the user) which applications and associated system files (e.g., in packages) are to be installed prior to the zone’s storing. (*Id.*; Ex. 1002 ¶165.)

**4. Claim 6: “assigning a unique associated identity to each of a plurality of the containers, wherein the identity includes at least one of IP address, host name, and MAC address”**

Tucker discloses the further limitation of claim 6 because “[e]ach zone which requires network connectivity will have one or more dedicated IP addresses.” (Ex. 1005, 44; Ex. 1002 ¶166.)

**5. Claim 8: “wherein the one or more applications and associated system files are retrieved from a computer system having a plurality of secure containers”**

Tucker discloses the further limitation of claim 8 by describing a model in which “applications can be freely moved between systems” by encapsulating the applications in zones. (Ex. 1005, 77; *id.*, 2 (zones “support rapid deployment (and redeployment) of applications, since the same application environment can be maintained on different physical machines”).) Moving or redeploying a zone requires retrieving it from the computer system where it was previously deployed.

(Ex. 1002 ¶167.) And it would have been obvious for the computer system where the zone was previously deployed to have “a plurality of secure containers” because Tucker discloses multiple non-global zones on the same system. (*Supra*, §I.C.2.)

6. **Claim 9: “wherein server information related to hardware resource usage including at least one of CPU memory, network bandwidth, and disk allocation is associated with at least some of the containers prior to the applications within the containers being executed”**

Tucker discloses the further limitation of claim 9 by way of “zone-wide limits [on] applicable resources,” which limit “the total resource usage of all process entities within a zone.” (Ex. 1005, 64.) For example, Tucker discloses that “Cpu shares” can be allocated to zones. (*Id.*, 64-65.) Further, Tucker discloses that “[t]he global administrator will be able to specify these limits in the zonecfg configuration file.” (Ex. 1005, 64.) Because configuration occurs before a zone is running, the limits in this configuration file are associated with a zone prior to applications being executed within the zone. (Ex. 1002 ¶168.)

7. **Claim 10: “wherein in operation when an application residing within a container is executed, said application has no access to system files or applications in other containers or to system files within the operating system during execution thereof”**

Tucker discloses the further limitation of claim 10 because “[p]rocesses running within a zone will be limited to files and file systems that can be accessed from the restricted root.” (Ex. 1005, 37; *see also id.*, 1 (“A zone is a ‘sandbox’

within which one or more applications can run without affecting or interacting with the rest of the system” and “isolation prevents processes running within a given zone from monitoring or affecting processes running in other zones”), 71 (“zones require a separate root file systems [sic] for isolation”).) Because Tucker discloses that each zone is limited to its own files (unless file sharing is specifically configured), applications in the zone have no access to system files or applications in other zones or in the operating system. (*Id.*; Ex. 1002 ¶169.)

**8. Claim 13: “associating with a plurality of containers a stored history of when processes related to applications within the container are executed for at least one of, tracking statistics, resource allocation, and for monitoring the status of the application”**

Claim 13 depends from claim 1. Tucker discloses the further limitation of claim 13 via an “extended accounting subsystem” that “is virtualized to permit different accounting settings and files on a per-zone basis for process- and task-based accounting.” (Ex. 1005, 64.) Records of processes “can be tagged with a zone id ... allowing the global administrator to determine resource consumption on a per-zone basis.” (*Id.*) Accordingly, Tucker discloses associating with zones a stored history of executed processes belonging to applications within the zone for the purpose of at least tracking statistics and resource allocation. (*Id.*; Ex. 1002 ¶170.)

**C. Ground 3: Claims 1, 2, 4, 6, 8-10, and 13-14 are Unpatentable as Obvious in View of Gélinas.**

Gélinas introduces Linux VServer, which provides secure containers that isolate and confine processes within the containers while allowing them to share the kernel of the underlying server. (*Supra*, §I.C.1.)

**1. Claim 1**

**a. Limitation 1[pre][i]: “In a system having a plurality of servers”**

Gélinas discloses that its containers (called “vservers”) operate in a system with multiple physical servers. (Ex. 1002 ¶173.) For example, it is “possible to move a vserver from one physical server to another.” (Ex. 1007, 30.)

**b. Limitation 1[pre][ii]: “with operating systems that differ”**

Gélinas discloses that its containers can run on differing versions of the Linux operating system, such as Debian and Redhat. (Ex. 1002 ¶174; Ex. 1007, 29.) Each of these different versions of Linux is called a “distribution” (or “distro”). (*See* Ex. 1007, 29; Ex. 1002 ¶174.) Gélinas also recognizes that “you may want to create several vservers to tests various distributions.” (Ex. 1007, 31; *id.* 27 (vserver should “work on any recent distribution”).) These various Linux distributions are operating systems that differ. (Ex. 1002 ¶174.)

**c. Limitation 1[pre][iii]: “operating in disparate computing environments”**

Gélinas discloses this limitation under both VirtaMove’s interpretation (requiring “independently operable” computers) and the inventor’s interpretation (requiring computers with different network addresses). (*Supra*, §V.B.)

First, Gélinas discloses implementing containers on computers running Linux. (*Supra*, §I.C.1.) Linux computers were independently operable because one Linux computer could operate without being controlled by any other computer. (Ex. 1002 ¶176.)

Second, Gélinas discloses implementing containers on computers with different network addresses. For example, Gélinas discloses that each container “is assigned a host name and an IP number.” (Ex. 1007, 2.) Each container “is tied to one IP [so] several servers may run the same services without conflict.” (*Id.*, 16-18.)

To the extent “disparate computing environments” requires some other type of difference, Gélinas would have rendered such an environment obvious. Gélinas discloses that “[m]ost hardware issues,” such as disk configuration, network configuration, processor type, and number of processors, “are irrelevant for the virtual server installation”—meaning that Gélinas’s system can operate in environments with a variety of different computer configurations. (Ex. 1007, 12; Ex. 1002 ¶178.) For example, Gélinas’ containers can operate in environments

comprising “an IDE + uniprocessor server” or “SCSI + multiprocessor server.” (Ex. 1007, 30.)

**d. Limitation 1[pre][iv]: “wherein each server includes a processor and an operating system including a kernel”**

Gélinas uses servers with processors. (Ex. 1007, 30 (“uniprocessor” or “multiprocessor”).) Further, the servers run Linux, which includes a kernel. (*Id.*, 12-13; Ex. 1002 ¶179.)

**e. Limitation 1[pre][v]: “a set of associated local system files compatible with the processor”**

“System files” include “shared libraries” and “configuration files.” (Ex. 1001, 2:52-54.) Gélinas discloses both. (Ex. 1007, 10 (shared libraries and configuration files); *id.*, 16, 24 (configuration files).) A POSITA would understand that the disclosed system files were compatible with the processor on which Linux was running; otherwise the applications that use the files would not work. (Ex. 1002 ¶180.)

Additionally, the ’814 patent admits that Apache used shared libraries and configuration files. (Ex. 1001, 2:55-3:20.) Gélinas discloses Apache on its servers. (Ex. 1007, 17, 21.)

Finally, Gélinas discloses running Linux distributions such as “debian” or “redhat.” (*Id.*, 29.) These distributions contained numerous shared libraries and configuration files compatible with processors on which they ran. (Ex. 1002 ¶182.)

- f. **Limitation 1[pre][vii]: “a method of providing at least some of the servers in the system with secure, executable applications related to a service”**

Gélinas’ containers “run pretty much any services,” including “telnet, mail servers, web servers, [and] SQL servers.” (Ex. 1007, 19, 11.) Each container acts as a “security box” that confines applications running in it. (*Id.*, 1; Ex. 1002 ¶183.)

- g. **Limitation 1[pre][vii]: “wherein the applications are executed in a secure environment”**

Gélinas’ containers “provide a higher level of security.” (Ex. 1007, 27; Ex. 1002 ¶184.) They can “[r]un un-trusted applications with complete control over their interaction with the rest of the computer and the network.” (Ex. 1007, 1.)

- h. **Limitation 1[pre][viii]: “wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service”**

As discussed above, Gélinas’ containers ran mail servers, web servers, and other services. (*Supra* §VI.C.1.f.) The applications used to run such services included executable objects that perform tasks related to the services. (Ex. 1002 ¶185.)

- i. **Limitation 1[a][i] “the method comprising: storing in memory accessible to at least some of the servers a plurality of secure containers of application software”**

Gélinas discloses VServer virtual servers, which are containers. (*Supra*, §I.C.1.) Installing such a container “copies a linux installation inside a sub-directory.” (Ex. 1007, 9.) This sub-directory is stored on a disk, which is a type of

memory. (*Id.*; Ex. 1002 ¶186.) Gélinas further discloses that one may “run several [containers] on the same box,” which “you will certainly do.” (Ex. 1007, 9.)

**j. Limitation 1[a][ii]: “each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications”**

As discussed above, Gélinas’ containers run “mail servers, web servers, SQL servers,” and the like. (Ex. 1007, 11.) A POSITA would have understood that such services comprise one or more executable applications and a set of associated system files because libraries and configuration files were needed to execute the applications. (Ex. 1002 ¶187.) For example, the ’814 patent admits that “Linux Apache [a common web server application] uses the following shared libraries ... which are ‘system’ files,” and then lists 21 libraries and configuration files. (Ex. 1001, 2:55-3:19.) Gélinas discloses Apache in its containers. (Ex. 1007, 17, 21; Ex. 1002 ¶187.)

**k. Limitation 1[a][iii]: “for use with a local kernel residing permanently on one of the servers”**

Gélinas discloses that its containers are “sharing the same kernel” and “do not need kernel packages,” which saves disk space and promotes efficiency. (Ex. 1007, 26, 29.) Based on this disclosure, a POSITA would understand that the containers make use of a local kernel on the underlying server. (Ex. 1002 ¶188.)

Further, Gélinas discloses how to install a kernel (outside of the containers) but says nothing about moving or uninstalling it. (Ex. 1007, 13-15.) Thus, it would have been obvious to a POSITA that the kernel resides permanently on the server. (Ex. 1002 ¶189.)

**l. Limitation 1[a][iv]: “wherein the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems”**

Gélinas discloses that its containers operate “like a normal Linux server” and run “normal services such as telnet, mail servers, web servers, SQL servers.” (Ex. 1007, 11.) A POSITA would understand that these services would be unable to run on Linux unless their associated system files were compatible with the kernel. (Ex. 1002 ¶190.) Indeed, Gélinas discloses that “[o]nce a service [is] running in a vserver, it is talking directly to the kernel.” (Ex. 1007, 29; *see also id.*, 16-18.)

**m. Limitation 1[a][v]: “the containers of application software excluding a kernel”**

Gélinas contrasts its approach, in which the containers are “sharing the same kernel,” with a virtual machine approach in which each machine “is running its own kernel.” (Ex. 1007, 26.) Thus, the containers exclude a kernel. (Ex. 1002 ¶191.)

- n. **Limitation 1[a][vi]: “wherein some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server”**

Gélinas discloses that a container may be created by “copy[ing] some parts of the current server.” (Ex. 1007, 15-16, 29 (“A vserver is normally created from the root server.”).) In this case, system files resident on the server are copied to the container using the command “cp -ax /sbin /bin /etc /usr /var /dev /lib” (or similar). (Ex. 1007, 15; *see* Ex. 1001, 2:55-3:16 (admitting that at least /bin, /usr, /etc, and /lib directories hold system files on Linux).) Because processes in a container cannot access files outside the container, a POSITA would understand that the copies in the container are utilized in place of the system files which remain on the underlying server. (Ex. 1002 ¶192 (explaining contents of copied directories).)

- o. **Limitation 1[a][vii]: “wherein said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server”**

System files in Gélinas’ containers are copies of associated local system files that remain resident on the underlying server for the reasons explained above. (*Supra* §VI.C.1.n; Ex. 1002 ¶193.)

- p. **Limitation 1[a][viii]: “wherein the application software cannot be shared between the plurality of secure containers of application software”**

Gélinas discloses five ways in which containers are isolated from each other. (Ex. 1007, 2-3.) This isolation prevents applications in one container from

interfering with applications in other containers. (*Id.*, 1-3.) Although an administrator can configure containers to share files, such configuration is optional. (Ex. 1007, 31; *see also id.* 9-10.) Unless an administrator specifically allows it, application software cannot be shared between containers. (*Id.*; Ex. 1002 ¶194.)

**q. Limitation 1[a][ix]: “wherein each of the containers has a unique root file system that is different from an operating system’s root file system”**

Gélinas discloses that “[t]he vserver is trapped into a sub-directory of the main server and can’t escape. This is done by the standard chroot() system call found on all Unix and Linux boxes.” (Ex. 1007, 2.) “As such, [containers] can only see what is under their / directory” and not the operating system’s full “/” directory. (Ex.1007, 31.) A container’s “/” directory is its root, and that root is different from the operating system’s root. (*Id.*; Ex. 1002 ¶195.)

Accordingly, Gélinas renders claim 1 obvious. (Ex. 1002 ¶¶173-196.)

**2. Claim 2: “wherein each container has an execution file associated therewith for starting the one or more applications”**

Gélinas discloses the further limitation of claim 2 by explaining that “[y]ou can setup a script” for “[e]xecuting tasks at vserver start/stop time.” (Ex. 1007, 20-21.) The script is called “/etc/vservers/XX.sh where XX is the name of the virtual server.” (*Id.*) Based on the “.sh” extension, a POSITA would understand that this is a shell script. (Ex. 1002 ¶197.) Shell scripts were commonly used to start

applications in Linux. (*Id.*) Thus, a POSITA would have found it obvious to use XX.sh as the execution file for starting each container’s applications. (*Id.*)

**3. Claim 4: “pre-identifying applications and system files required for association with the one or more containers prior to said storing step”**

Gélinas discloses the further limitation of claim 4 because its containers may be copied or moved “from one physical server to another.” (Ex. 1007, 30.) In this situation, applications and system files in the container would be identified on the source server before they were copied or moved to the destination. (Ex. 1002 ¶198.) For example, the applications and system files would be located in the directory “etc/vservers/XX” (where XX is the name of the container to be copied). (Ex. 1007, 30.) The command used to copy this directory—such as “rsync” (*id.*)—would identify the applications and system files in that directory when copying them. (Ex. 1002 ¶198.) Such identification and copying would occur before the container was stored on the destination server. (*Id.*)

**4. Claim 6: “assigning a unique associated identity to each of a plurality of the containers, wherein the identity includes at least one of IP address, host name, and MAC address”**

Gélinas discloses the further limitation of claim 6 because each container “is assigned a host name and an IP number.” (Ex. 1007, 2, 16-18; Ex. 1002 ¶199.)

**5. Claim 8: “wherein the one or more applications and associated system files are retrieved from a computer system having a plurality of secure containers”**

Gélinas discloses the further limitation of claim 8 because it states that a physical server stores a plurality of containers on disk. (Ex. 1007, 9-10.) When applications and system files are executed, a POSITA would understand they are retrieved from the disk that stores the containers, which is part of a computer system (the physical server). (*Id.*, Ex. 1002 ¶200.)

**6. Claim 9: “wherein server information related to hardware resource usage including at least one of CPU memory, network bandwidth, and disk allocation is associated with at least some of the containers prior to the applications within the containers being executed”**

Gélinas discloses the further limitation of claim 9 because a container may have a flag that “unifies the processes in a vservers from a scheduler view point” and thus “prevents a vservers from taking [too] much CPU resources.” (Ex. 1007, 17-18, 24.) A flag may also be set to limit a container to a maximum number of processes (*id.*, 23) or a particular set of networking capabilities (*id.*, 17-18). Thus, server information (e.g., a flag on the server) related to hardware resource usage (e.g., CPU or networking) may be associated with a container. (Ex. 1002 ¶201.)

Further, the flag information is stored in the “XX.conf” file for a particular container. (Ex. 1007, 16-17.) The XX.conf file is created before the container (and any applications within it) is started. (*See id.* (“ONBOOT” parameter in “XX.conf”

controls how containers are started); *id.* at 30 (move “XX.conf” to a new server before starting container there); Ex. 1002 ¶202.)

**7. Claim 10: “wherein in operation when an application residing within a container is executed, said application has no access to system files or applications in other containers or to system files within the operating system during execution thereof”**

Gélinas discloses the further limitation of claim 10 because, as discussed above, containers can only see their own root directory and not what is outside of that directory. (Ex. 1007, 31; *supra* §VI.C.1.q). Because each container is limited to its own files, applications within the container have no access to system files or applications in other containers or within the underlying operating system. (*Id.*; Ex. 1002 ¶203; *see also* Ex. 1007, 2-3 (explaining isolation).)

**8. Claim 13: “associating with a plurality of containers a stored history of when processes related to applications within the container are executed for at least one of, tracking statistics, resource allocation, and for monitoring the status of the application”**

Gélinas discloses the further limitation of claim 13 because its system can produce a report listing “the number of process in each active security context as well as the name of the vserver associated with this context.” (Ex. 1007, 14.) A POSITA would have understood that such a report would be useful for tracking statistics. (*See id.* (command called “/usr/sbin/vserver-stat”); Ex. 1002 ¶204.) Further, Gélinas discloses another command that reports “the running status” of

every container, useful for monitoring the status of applications in the containers. (Ex. 1007, 20; Ex. 1002 ¶204.)

## 9. Claim 14

Claim 14 depends from claim 1. Gélinas discloses each further limitation of claim 14. (Ex. 1002 ¶¶205-210.)

### a. Limitation 14[a][i]: “creating containers prior to said step of storing containers in memory, wherein containers are created by:”

Gélinas discloses creating containers from system files on the underlying server. (Ex. 1007 at 15-16; *supra* §VI.C.1.n.) Further, Gélinas discloses that containers may be copied from a first server to a second server. (Ex. 1007, 30.) Thus, containers are created on the first server before being copied to the second server and stored in the second server’s memory. (Ex. 1002 ¶206.)

### b. Limitation 14[a][ii]: “a) running an instance of a service on a server”

Gélinas discloses that a container on a physical server “runs normal services such as telnet, mail servers, web servers, SQL servers.” (Ex. 1007, 11.)

### c. Limitation 14[a][iii]: “b) determining which files are being used”

Gélinas discloses that files being used by a container “XX” are stored in a directory called “/vservers/XX.” (Ex. 1007, 15-16, 30; Ex. 1002 ¶208.) These files are transferred when a container is copied from a first server to a second server. (Ex. 1007, 30.) Gélinas discloses that the files may be transferred using an “rsync”

command, with the “/vservers/XX” directory supplied to the command as a parameter. (*Id.*; Ex. 1002 ¶208.) Thus, Gélinas discloses determining which files are being used (e.g., the files in “/vservers/XX”) when a command is executed to copy a container. (Ex. 1002 ¶208; *supra*, §VI.C.3.)

**d. Limitation 14[a][iv]: “c) copying applications and associated system files to memory without overwriting the associated system files so as to provide a second instance of the applications and associated system files”**

As explained above, Gélinas discloses copying a container from a first server to a second server. (Ex. 1007, 30.) Because the contents of the copied container are stored on the second server, they do not overwrite the contents on the first server. (Ex. 1002 ¶209.)

For the reasons above, Gélinas renders claim 14 obvious. (*Id.* ¶¶205-210.)

## **VII. SECONDARY CONSIDERATIONS OF NONOBVIOUSNESS**

Secondary considerations should be considered but do not control the obviousness conclusion. *Newell Cos. v. Kenney Mfg. Co.*, 864 F.2d 757, 768 (Fed. Cir. 1988). Where, as here, a strong *prima facie* obviousness showing exists, any evidence of secondary considerations may not dislodge the obviousness conclusion. *Leapfrog Enters. v. Fisher-Price, Inc.*, 485 F.3d 1157, 1162 (Fed. Cir. 2007). Petitioner is aware of no evidence supporting a claim for secondary considerations.

## **VIII. DISCRETIONARY DENIAL IS UNWARRANTED**

Pursuant to and in reliance on Acting Director Coke M. Stewart’s March 26, 2025, Memorandum regarding Interim Processes for PTAB Workload Management, Petitioner understands that discretionary denial issues, if any, will be raised in a separate brief to be filed by Patent Owner. If Patent Owner files such a brief, Petitioner intends to respond in an opposition brief consistent with Acting Director Coke M. Stewart’s March 26, 2025, Memorandum regarding Interim Processes for PTAB Workload Management. Accordingly, Petitioner will not address discretionary denial issues in this Petition other than to note that (a) Petitioner is filing concurrently with this Petition a Motion for Joinder with Amazon’s IPR2025-00563 for the ’814 patent, and, if joined, Petitioner would be taking an understudy role and the Board’s finite resources would not be impacted; (b) Microsoft filed on April 18, 2025, along with its Petition for Inter Partes Review of the ’814 patent that is “substantively identical to” Amazon’s IPR2025-00563, a Motion for Joinder to the same Amazon’s IPR2025-00563, and Microsoft likewise stated there that it will, if joined, accept an understudy role. *See Microsoft Corp. v. VirtaMove, Corp.*, IPR2025-00851, Paper 3 (PTAB Apr. 18, 2025).

## **IX. MANDATORY NOTICES**

### **A. Real Parties-In-Interest (37 C.F.R. §42.8(b)(1))**

Petitioner Oracle Corporation (“Oracle”) is the Real Party-in-Interest. No other parties exercised or could have exercised control over this Petition. No other

parties funded or directed this Petition. *See* Office Patent Trial Practice Guide, 77 Fed. Reg. 48759-60.

## **B. Related Matters**

### **1. United States Patent & Trademark Office**

The application from which U.S. Patent No. 7,519,814 issued claims priority to two provisional applications: No. 60/502,619, filed September 15, 2003 and No. 60/512,103, filed October 20, 2003.

The following U.S. patent applications claim the benefit of priority to U.S. Patent 7,519,814:

(i) U.S. Patent Application 11/432,843 (U.S. Patent No. 7,757,291), filed May 12, 2006;

(ii) U.S. Patent Application 11/380,285 (U.S. Patent No. 7,774,762), filed April 26, 2006;

(iii) U.S. Patent Application 12/075,842 filed March 13, 2008.

### **2. USPTO Patent Trial and Appeal Board**

Concurrently with the present petition, Petitioner is also filing IPR2025-00964, IPR2025-00965, and IPR2025-01002, challenging U.S. Patent No. 7,519,814 (the “814 patent”).

Petitioner is also filing IPR2025-00966, IPR2025-00981, and IPR2025-00982, challenging U.S. Patent No. 7,784,058 (the “058 patent”), which is also

asserted in *VirtaMove, Corp. v. Oracle Corporation*, Case No. 7:24-cv-00339-ADA, listed below.

Other *inter partes* review proceedings filed against the '814 or '058 patents include:

(i) *Google LLC v. VirtaMove, Corp.*, Case No. IPR2025-00487 (PTAB, filed January 31, 2025);

(ii) *Google LLC v. VirtaMove, Corp.*, Case No. IPR2025-00488 (PTAB, filed January 31, 2025);

(iii) *Google LLC v. VirtaMove, Corp.*, Case No. IPR2025-00489 (PTAB, filed January 31, 2025);

(iv) *Google LLC v. VirtaMove, Corp.*, Case No. IPR2025-00490 (PTAB, filed January 30, 2025);

(v) *Amazon.com, Inc. v. VirtaMove, Corp.*, Case No. IPR2025-00561 (PTAB, filed January 31, 2025);

(vi) *Amazon.com, Inc. v. VirtaMove, Corp.*, Case No. IPR2025-00563 (PTAB, filed January 31, 2025);

(vii) *Amazon.com, Inc. v. VirtaMove, Corp.*, Case No. IPR2025-00566 (PTAB, filed January 31, 2025);

(viii) *International Business Machines Corp. v. VirtaMove, Corp.*, Case No. IPR2025-00591 (PTAB, filed February 6, 2025);

(ix) *International Business Machines Corp. v. VirtaMove, Corp.*, Case No. IPR2025-00599 (PTAB, filed February 7, 2025);

(x) *Microsoft Corp. v. VirtaMove, Corp.*, Case No. IPR2025-00849 (PTAB, filed April 18, 2025);

(xi) *Microsoft Corp. v. VirtaMove, Corp.*, Case No. IPR2025-00850 (PTAB, filed April 18, 2025);

(xii) *Microsoft Corp. v. VirtaMove, Corp.*, Case No. IPR2025-00851 (PTAB, filed April 18, 2025);

(xiii) *Microsoft Corp. v. VirtaMove, Corp.*, Case No. IPR2025-00852 (PTAB, filed April 18, 2025);

(xiv) *Microsoft Corp. v. VirtaMove, Corp.*, Case No. IPR2025-00853 (PTAB, filed April 18, 2025);

(xv) *Microsoft Corp. v. VirtaMove, Corp.*, Case No. IPR2025-00854 (PTAB, filed April 18, 2025);

(xvi) *Microsoft Corp. v. VirtaMove, Corp.*, Case No. IPR2025-00855 (PTAB, filed April 18, 2025).

**3. U.S. District Court for the Eastern District of Texas**

(i) *VirtaMove, Corp. v. Hewlett Packard Enterprise Company*, Case No. 2:24-cv-00093;

(ii) *VirtaMove, Corp. v. International Business Machines Corp.*, Case No. 2:24-cv-00064.

**4. U.S. District Court for the Western District of Texas**

(i) *VirtaMove, Corp. v. Google LLC*, Case No. 7:24-cv-00033  
(transferred to Northern District of California per Order dated May 7, 2025, *see* Ex. 1028);

(ii) *VirtaMove, Corp. v. Amazon.com, Inc. et al.*, Case No. 7:24-cv-00030  
(pending transfer to Northern District of California per Order dated February 19, 2025, *see* Docket Entry No. 94);

(iii) *VirtaMove, Corp. v. Microsoft Corp.*, Case No. 7:24-cv-00338;

(iv) *VirtaMove, Corp. v. Oracle Corp.*, Case No. 7:24-cv-00339.

**5. U.S. District Court for the Northern District of California**

(i) *Red Hat, Inc. v. VirtaMove, Corp.*, Case No. 5:24-cv-04740;

(ii) *VirtaMove, Corp. v. Google LLC*, Case No. 5:25-cv-00860.

**C. Counsel, Service, and Fee Information**

<b>Lead Counsel</b>	<b>Back-Up Counsel</b>
<p>Bas de Blank  Registration No. 74,930  (M2BPTABDocket@orrick.com)</p> <p>Postal &amp; Hand-Delivery Address:  Orrick, Herrington &amp; Sutcliffe LLP  1000 Marsh Road  Menlo Park, CA 94025  Tel: 650-614-7400  Fax: 650-614-7401</p> <p><i>Attorney for Petitioner Oracle Corporation</i></p>	<p>Parth Sagdeo  Registration No. 71,275  (PTABDocketP2S7@orrick.com)</p> <p>Jared Bobrow  <i>Pro Hac Vice</i> to be submitted  (PTABDocketJ3B3@orrick.com)</p> <p>Postal &amp; Hand-Delivery Address:  Orrick, Herrington &amp; Sutcliffe LLP  1000 Marsh Road  Menlo Park, CA 94025  Tel: 650-614-7400  Fax: 650-614-7401</p> <p>Shane Anderson  <i>Pro Hac Vice</i> to be submitted  (PTABDocket9SA@orrick.com)</p> <p>Postal &amp; Hand-Delivery Address:  Orrick, Herrington &amp; Sutcliffe LLP  355 S. Grand Ave.  Ste. 2700  Los Angeles, CA 90071  Tel: 213-629-2020  Fax: 213-612-2499</p> <p><i>Attorneys for Petitioner Oracle Corporation</i></p>

Petitioner consents to service by electronic mail at the following addresses:

M2BPTABDocket@orrick.com, PTABDocketJ3B3@orrick.com,

PTABDocket9SA@orrick.com; PTABDocketP2S7@orrick.com, and oracle-

virtamove\_ohs@orrick.com. Pursuant to 37 C.F.R. §42.10(b), Petitioner attaches a Power of Attorney.

**D. Payment of Fees (37 C.F.R. §42.103)**

The fee set forth in 37 C.F.R. §42.15(a) for this petition has been paid. Any additional fees due in connection with this petition may be charged to the deposit account of Orrick, Herrington, & Sutcliffe LLP: 15-0665.

**E. Grounds for Standing (37 C.F.R. §42.104(a))**

Oracle certifies that the '814 patent is available for IPR and that Oracle is not barred or estopped from requesting IPR. This petition is being filed within one year of service of the district court Complaint.

**X. CONCLUSION**

Petitioner respectfully requests cancellation of the challenged claims.

Respectfully submitted,

ORRICK, HERRINGTON & SUTCLIFFE LLP

Dated: May 16, 2025

/Bas de Blank/

---

Bas de Blank  
Orrick, Herrington & Sutcliffe LLP  
Lead Counsel for Petitioner  
Registration No. 74,930  
1000 Marsh Road  
Menlo Park, CA 94025  
Tel: 650-614-7400  
Fax: 650-614-7401  
Email: m2bptabdocket@orrick.com

*Attorney for Petitioner Oracle Corporation*

**CERTIFICATE OF COMPLIANCE**

The undersigned certifies that the foregoing PETITION FOR *INTER PARTES* REVIEW complies with the type volume limitation in 37 C.F.R. § 42.24(a). According to the utilized word-processing system's word count, the petition—excluding the caption, table of contents, table of exhibits, mandatory notices, certificate of word count, and certificate of service—contains 13,772 words.

Dated: May 16, 2025

By: */Bas de Blank/*

---

Bas de Blank  
Orrick, Herrington & Sutcliffe LLP  
Lead Counsel for Petitioner  
Registration No. 74,930  
1000 Marsh Road  
Menlo Park, CA 94025  
Tel: 650-614-7400  
Fax: 650-614-7401  
Email: m2bptabdocket@orrick.com

*Attorney for Petitioner Oracle  
Corporation*

**CERTIFICATE OF SERVICE**

The undersigned hereby confirms that the foregoing Petition for *Inter Partes* Review and associated documents and exhibits were caused to be served on May 16, 2025 via overnight courier upon the following counsel of record for Patent Owner:

Allen, Dyer, Doppelt + Gilchrist, PA  
1135 East State Road 434, Suite 3001  
Winter Springs, FL 32708

Copies of this Petition and accompanying documents and exhibits were also served via electronic mail on Patent Owner’s counsel of record for related PTAB proceedings and in the related district court litigation – Russ, August & Kabat:

Reza Mirzaie (rmirzaie@raklaw.com)  
Marc A. Fenster (mfenster@raklaw.com)  
Neil Rubin (nrubin@raklaw.com)  
James A. Milkey (jmilkey@raklaw.com)  
Qi (Peter) Tong (ptong@raklaw.com)  
Jacob R. Buczko (jbuczko@raklaw.com)  
Christian W. Conkle (cconkle@raklaw.com)  
Jefferson Cummings (jcummings@raklaw.com)  
Daniel B Kolko (dkolko@raklaw.com)  
Jonathan Ma (jma@raklaw.com)  
Mackenzie Paladino (mpaladino@raklaw.com)  
James S. Tsuei (jtsuei@raklaw.com)

*/Karen Johnson/*

\_\_\_\_\_  
Karen Johnson