UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

VMware, Inc.

Petitioner,

v.

Cirba IP Inc.

Patent Owner.

Case PGR2021-00098
U.S. Patent 10,951,459

DECLARATION OF DR. VIJAY MADISETTI IN SUPPORT OF PATENT OWNER'S PRELIMINARY RESPONSE

TABLE OF CONTENTS

I, Vijay Madisetti, declare as follows:

1. I have been retained by the Counsel for Patent Owner Cirba IP, Inc. ("Cirba IP") to serve as an expert in connection with *VMware v. Cirba IP Inc.*, PGR2021-00098.

I. SCOPE AND SUMMARY OF OPINIONS

2. The scope of this Declaration is limited to supporting Cirba IP's Patent Owner's Response to VMware's Petition for post grant review of U.S. Patent No. 10,951,459 ("the '459 patent") ("Petition"). Ex. 1001 ('459 patent). I may offer additional opinions based on my further review of the materials in this case, including opinions and/or testimony of other expert witnesses.

3. Specifically, I have been asked to analyze the Petition and Dr. Zadok's declaration in support thereof, as well as the '459 patent and its file history. I further have been asked to provide opinions about how a person of ordinary skill in the art ("POSA") would have understood and interpreted the '459 patent, including its claims and specification; the disclosures of applications incorporated by reference in the '459 patent and a POSA's knowledge of the state of the art. I understand that this Declaration is being submitted in connection with Patent Owner's Response.

4. In preparing this Declaration, I have reviewed the Petition; the declaration of Dr. Erez Zadok that is Exhibit 1006 to the Petition, and other exhibits

to the Petition. I provide my opinions and conclusions regarding these documents below.

5. For purposes of this Declaration, I was asked to consider claims 1-63 of the ’459 patent, which I understand are the challenged claims. *See* Petition at 4. I understand that claims 1, 32, and 63 are the independent claims.

6. This Declaration and exhibits, in addition to my education and experience in the field, form the basis of my opinion. The citations to references in this Declaration are representative. All of the information I am providing should be considered together; that is, information from the main body of this Declaration may be relevant to my exhibits, information from one exhibit may be relevant to another exhibit, and information from my exhibits may be relevant to the Declaration’s main body included herein. I also base my opinions on my knowledge, experience, and training in the fields of technology related to the ’459 patent.

7. I am not offering any conclusions as to the ultimate determinations that I understand the Patent Trial and Appeal Board (“Board”) will make in this proceeding. I am solely providing my opinion on the technical interpretations and implements of the documents I understand are involved in this proceeding.

8. I understand that at this phase of the proceeding, Petitioner may be granted an additional opportunity to file a paper in support of the Petition. If it does, I may offer additional opinions and submit additional declarations in response.

II. MATERIALS CONSIDERED

9. In forming the opinions in this Declaration, I reviewed the Petition (Paper 2) and its accompanying exhibits listed in the "Exhibit List" on page *iii* of the Petition, including the Dr. Zadok's declaration (Ex. 1006), the '459 patent, and the '459 patent's file history. A list of the materials I considered is included as Exhibit 2002.

III. INTRODUCTION, QUALIFICATIONS, & ASSIGNMENT

10. I received my Bachelor of Technology (Honors) in Electronics and Electrical Communication Engineering at the Indian Institute of Technology (IIT) in Kharagpur, India in 1984. I obtained my Ph.D. in Electrical Engineering and Computer Science at the University of California, Berkeley, in 1989. I received the Demetri Angelakos Outstanding Graduate Student Award from the University of California, Berkeley and the IEEE/ACM Ira M. Kay Memorial Paper Prize in 1989.

11. I am a tenured Professor in Electrical and Computer Engineering at the Georgia Institute of Technology ("Georgia Tech"). I am knowledgeable and familiar with information technology infrastructures, designing and evaluating virtualized

environments, virtual machine architecture and optimization, cloud technology, as well as wireless communications, microprocessor architecture, hardware, RF, cellular networks, ASIC design, computer engineering, embedded systems, digital signal processing, and associated software and firmware design for wireless and telecommunications terminals and base stations.

12. I have been working in the areas of cloud computing, workload modeling, virtualization, and benchmarking for over ten years. In the area of characterization, modeling and generation of workloads for cloud computing applications, I have created a synthetic workload generator for virtual environments that accepts benchmark and workload model specifications. I also developed a cloud workload specification language called, GT-CWSL, to provide a structured way for specification of application workloads. These approaches allow accurate characterization of virtualization and cloud performance and have been published in "Synthetic Workload Generation for Cloud Computing Applications," Journal of Software Engineering and Applications, 2011, Vol. 4, pp. 396-410.

13. In other work related to virtualization, performance modeling, and data integration, I have developed cloud-based information integration and informatics framework, called CHISTAR, that allows integration of distributed and heterogeneous data into a common virtual environment (on commodity hardware

running hypervisors), and allows easier analysis and analytics to be performed. This research was presented as a cover feature in the Special Issue on Computing in Healthcare, IEEE Computer Magazine, in 2015 ("Healthcare Data Integration and Informatics in the Cloud").

14. Further, in my work on rapid prototyping of cloud-based virtual services and systems, I proposed a new methodology using loosely coupled virtual components that are not restricted by architecture or programming styles. This approach, shown below, creates virtual components that are then integrated and deployed effectively to realize improvements in deployment efficiency and time to deployment.

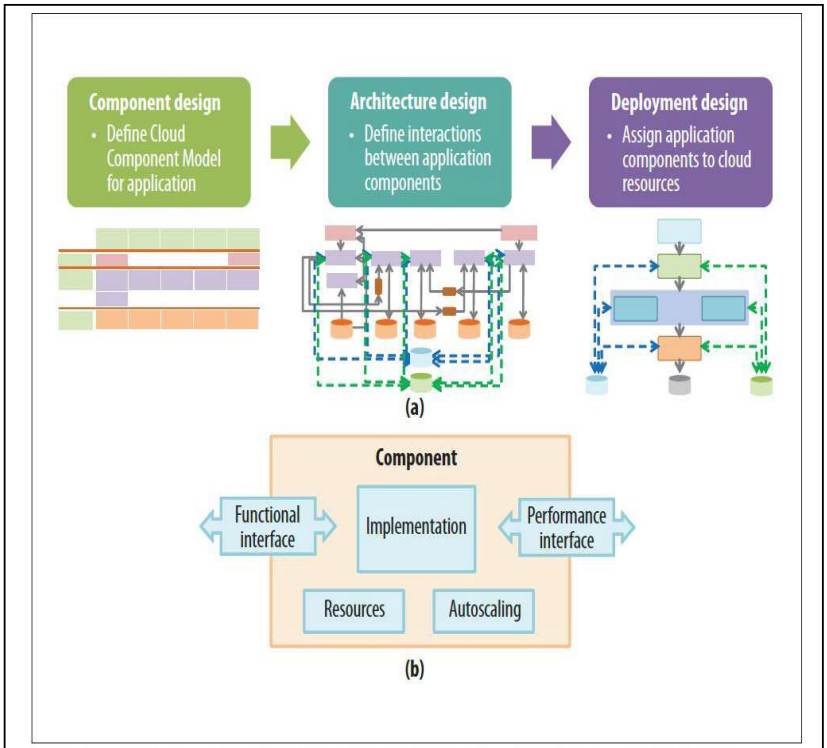

Figure 2. Cloud Component Model (CCM)-based development: (a) design methodology, which comprises component, architecture, and deployment design; and (b) CCM component architecture. CCM's component-based approach is suitable for both Web-based and mobile applications.

15. Generally, the use of our CCM-based virtualization and cloud component methodology improved throughput and response times. These results were published in the flagship journal, IEEE Computer Magazine, in November 2013. A. Bahga, V. Madisetti, "Rapid Prototyping of Advanced Cloud-Based Systems," IEEE Computer, vol. 46, no. 11, Nov. 2013, pp. 76-83.

16. Yet another relevant publication 1997 IEEE paper that concerned an open signal processing system design and implementation environment, BEEHIVE, that allows application developers to rapidly compose and debug functional specifications in a networked, distributed computing environment, and then later

migrate the application (transparently) onto an embedded, distributed, computing hardware/software platform, with the capability to reconfigure (adaptively) the resources assigned to the application to meet the dynamic real-time requirements of the implementation. 1997 IEEE International Conference vol. 1, 21-24 April 1997, pp. 663 - 666.

17. In addition to research and commercialization projects in virtualization and cloud computing and advanced cloud applications and data analytics, I have also taught courses and written books in these areas. These books are widely used by dozens of universities all over the world: A Bahga, V. Madisetti, *Cloud Computing: A Hands-On Approach* (2013); A Bahga, V. Madisetti, *Internet of Things: A Hands-On Approach* (2014); A Bahga, V. Madisetti, *Big Data Science & Analytics: A Hands-On Approach* (2016); A Bahga, V. Madisetti, *Blockchain Applications: A Hands-On Approach* (2017).

18. I have been an active consultant to industry and various research laboratories (including Massachusetts Institute of Technology Lincoln Labs and Johns Hopkins University Applied Physics Laboratory). I also have founded three companies in the areas of embedded software, military chipsets involving imaging technology, and wireless communications, and I have supervised the Ph.D. dissertations of over twenty engineers in the areas of computer engineering, signal

processing, communications, rapid prototyping, and system-level design methodology, of which five have resulted in thesis prizes or paper awards.

19. I have been elected a Fellow of the IEEE, for contributions to embedded computing systems. The Fellow is the highest grade of membership of the IEEE, a world professional body consisting of over 300,000 electrical and electronics engineers, with only one-tenth of one percent (0.1%) of the IEEE membership being elected to the Fellow grade each year. Election to Fellow is based upon votes cast by existing Fellows in IEEE.

20. I have authored more than sixty refereed journal publications and around forty peer reviewed conference publications, and I have submitted approximately thirty invention disclosures and provisional patent applications over the past ten years. Nine are issued US patents.

21. I have testified as an expert witness before. Over the past six years, I've testified as an expert in more than 20 proceedings.

22. I am being compensated by Cirba IP at the rate of $500 per hour for my work in this proceeding, including time spent testifying. I am also being reimbursed for reasonable fees and expenses, including hotel and travel expenses, incurred as a result of my work in this case. I have not received any additional compensation for my work in this investigation. My compensation is not tied in any way to the

substance of my testimony or the outcome of this proceeding. I am available to offer these opinions at deposition and at hearings if called upon to do so.

23. A copy of my CV is included as Exhibit 2003.

IV. UNDERSTANDING OF THE LAW

24. I have been informed by Cirba IP's Counsel about the legal principles that apply to the issues upon which Dr. Zadok provided an opinion in his declaration, which is Exhibit 1006 to the Petition. I have applied these legal standards to the facts, circumstances, and materials that I considered, along with my education, training, qualifications, and experience in the field, in reaching the conclusions and opinions expressed in this report. I provide my observations and opinions about the claims and grounds asserted in the Petition based on information as I presently know and understand it, but I reserve the right to amend and/or modify my positions in the future based on additional knowledge and understanding.

A. Claim Construction

25. I have been informed that in post grant review proceedings, the Board construes claims in accordance with the same claim construction standard used to construe the claim in a civil action under 35 U.S.C. 282(b), including construing the claim in accordance with the ordinary and customary meaning as understood by one of ordinary skill in the art and in light of the patent's specification and prosecution

history. I understand the Board also considers prior claim construction determinations from other proceedings concerning the challenged claims. I am informed that, at this time, a prior claim construction determinations was issued on February 24, 2022 concerning the challenged claims of the '459 patent, in a United States District Court in the District of Delaware litigation captioned *Cirba, Inc. d/b/a Densify et al. v. VMware, Inc.*, DDE-1-19-cv-00742-LPS.

26. I understand that the words of a claim are generally given their ordinary and customary meaning, which is the meaning that a term would have to a person of ordinary skill in the art. I understand the person of ordinary skill in the art is assumed to read claim terms not only in the context of the particular claim in which they appear, but in the context of the entire patent, including the specification. I further understand that courts may consider a patent's prosecution history and certain "extrinsic evidence," such as dictionaries and expert testimony.

27. I further understand that a patentee may choose to be his/her own lexicographer and define a term differently than the term's plain and ordinary meaning in the art and that, under such circumstances, the patentee's own definition should control. However, I understand that if the specification does not show an express intent to give a term a new meaning, the term is given its ordinary meaning.

28. When construing claims, I understand that courts will first look to the words of the claim itself. I also understand that courts interpret the entire claim in context, including words surrounding the terms at issue. Terms are not considered in isolation or "in a vacuum." I further understand that the specification may be the best source of understanding the meaning of a term.

29. I understand that courts may consult dictionaries to determine whether a patent gives a term a special meaning. But I also understand that courts should not construe terms to be broader than their ordinary meaning reflected in dictionary definitions and in the overall context of the intrinsic record. I understand that the ordinary meaning of claim terms cannot be broadened in the absence of support in the intrinsic record that the broad meaning was intended. I further understand that while terms by themselves could have varying degrees of breadth, courts construe claim terms to implement the invention described in the specification.

B. Patentability

30. While I am not an attorney, I have been informed that a patent is presumed valid, and that each claim of a patent (whether an independent or dependent claim) is presumed valid independently of the validity of other claims. I understand that the Patent Office will find a patent claim unpatentable in a post grant review if it concludes that it is more likely than not that the claim is unpatentable. I

understand that this is also referred to as the "preponderance of the evidence standard."

C. Written Description

31. I understand that, according to 35 U.S.C. § 112(a), the specification of a patent must contain a written description of the invention, and of the manner and process of making and using it, in such terms as to enable a person of skill in the art to which the patent pertains to make and use the invention. I further understand that the written description requirement concerns whether a person of ordinary skill in the art can recognize what was claimed corresponds to what was described, and not whether the invention works, or how to make it work.

32. I further understand that the "written description" aspect of this requirement means that the patent disclosure must reasonably convey to a POSA that the patent's inventor had possession of the claimed subject matter as of the filing date of the patent. I further understand that, because the specification is directed toward a POSA, it need not include information that was well known in the art in order to satisfy the written description requirement. I further understand that the specification need not include word-for-word support of the claims, nor any specific examples of the claimed inventions. It is my understanding that the specification

can describe the full scope of the claimed invention without examples explicitly covering the full scope of the claimed invention.

V. TECHNOLOGY

A. Background

1. Sprawl and Consolidation

33. As explained in the '459 patent, sprawl is the proliferation of servers and/or applications in a distributed network that can occur when servers and/or applications are incrementally added to the network over time. *See* Ex. 1001 ('459 patent) at 1:47-2:4. Such proliferation is inefficient where the network accumulates far more processing capacity than is required, leaving servers in the network underutilized. *Id*. Having underutilized servers is costly because such servers may nonetheless require maintenance and attention, consume power, and generate heat. *Id*. Such servers may also incur redundant software licensing costs. *Id*.

34. Sprawl is also a concern for virtual machines. *See* Ex. 2010 (Newman) at 15-16; Ex. 2011 (Brodkin) at 1-2; Ex. 2012 (Provazza) at 2. The industry has recognized that virtual machine sprawl not only can lead to wasteful use of physical and software resources, but virtual machine sprawl can be especially difficult to control because virtual machines are relatively easy to generate. *See* Ex. 2010 (Newman) at 15-16; Ex. 2011 (Brodkin) at 1-2; Ex. 2012 (Provazza) at 2. Thus, while virtualization is recognized as one technique for consolidating systems to

eliminate server sprawl, virtualization can lead to other types of sprawl. *See* Ex. 2010 (Newman) at 15-16; Ex. 2011 (Brodkin) at 1-2; Ex. 2012 (Provazza) at 2.

35. To reverse the effects of sprawl, certain consolidation techniques have been employed. For instance, virtualization is a consolidation technique whereby multiple "virtualized" systems run on a single physical machine, thereby reducing the need for separate physical machines. Ex. 1001 ('459 patent) at 2:14-19. Operating system-level application stacking moves different applications running on one or more systems to a single consolidated system that runs all the applications so fewer total systems are required. *Id*. at 2:30-36. Physical consolidation moves physical systems at the operating system level to hardware platforms that can support multiple systems. *Id*. at 2:37-42.

2. Virtualization

36. In computing, "virtualization" generally refers to the act of creating a "virtual" (rather than physical) version of something. Ex. 2012 at 1. Virtualization is accomplished through software that makes computing environments independent of physical infrastructure, turning them into "virtual environments." *Id.* A virtual environment is "a computer that is running in a virtual machine environment, which is the combination of virtual machine monitor and hardware platform. For example, VMware running on an x86 computer is a virtual environment." Ex. 2014. It refers

to the dividing-up of resources (e.g., CPU, network, storage and memory) of a computer (e.g., a server or a host) into multiple "virtual" servers. The virtual servers are called "virtual machines" (VMs). VMs are then used to run software applications like a regular physical server. With virtualization, one physical server can run many applications as each application runs in a VM sharing the resources of the physical machine.[1] The figure below shows a depiction of the architecture of a virtualization technology:

[1] *See* Ex. 2015 (Hammersley) at 2 ("This hardware abstraction allows multiple 'virtual' environments to function side by side on top of a single physical machine.").

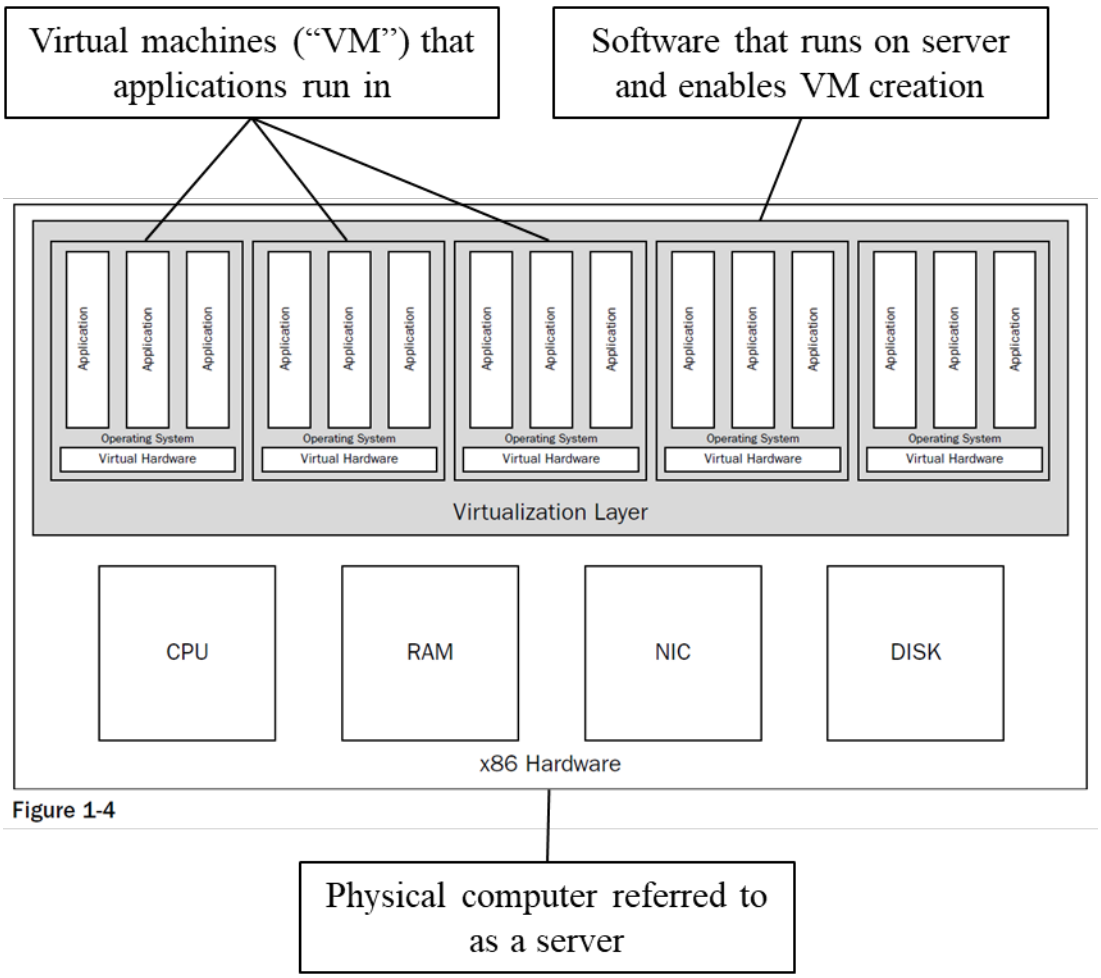

Ex. 2015 (Hammersley) at 5 (annotated).

37. Virtualization also allows for multiple physical servers to be configured as one larger logical entity; the resulting group of physical servers is called a cluster. A cluster may offer an advantage of managing several physical servers as one larger resource pool. An additional benefit to a cluster is that VMs may be dynamically moved to different physical servers within a cluster. This VM may move because other VMs running on the same physical server may be utilizing all its resources

(e.g., CPU, network, storage, memory) and another physical server may have more available resources, as shown in the picture below. Ex. 2016 at 19.

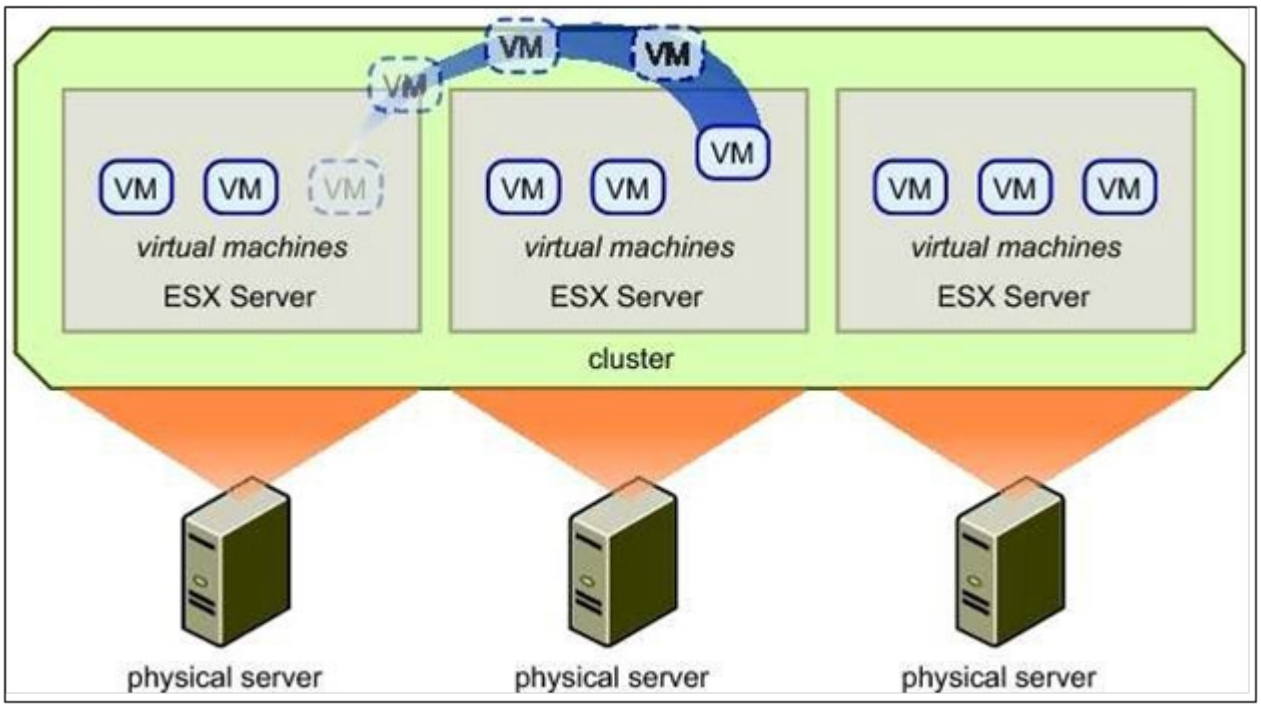

38. The core software for virtualization is referred to as a "hypervisor," which enables the dividing-up of resources among VMs. The hypervisor also coordinates the sharing of resources among VMs. Generally, the physical server and related software can be called a "host." The VM running on the host can be called a "guest."

39. In the late 2000s and early-to-mid 2010s, following the increasing popularity of hypervisors and virtualization, a number of vendors—including VMware, Microsoft, and Red Hat—offered competing hypervisor products in the marketplace in what was colloquially referenced within the industry as the

"hypervisor wars." *See* Ex. 2017 (Defense Industry Daily); Ex. 2018 (Steele); Ex. 2019 (Kerner); Ex. 2020 (Venezia); Ex. 2021 (Dubie). It is generally agreed that VMware emerged victorious in these "hypervisor wars," due at least in part to its DRS technology. *Id*.

B. General Concepts Related to Virtualization

1. Initial Placement

40. As discussed above, virtualization allows for the division of computing resources across one or more physical servers through the use of virtual machines. An important consideration in virtualization and virtualized environments is the initial placement of the created virtual machines across the virtualized environment, which will bear on the future resource utilization and efficiency of the virtualized environment as explained further below.

2. Workload Placement and Optimizations

41. Workload placement refers to the provisioning and allocation of resources in virtualized environments, as well as the process of provisioning and allocating resources in virtualized environments so that the efficiency of the virtual machines and other components of the virtualized environments are optimized. Allocated resources that ultimately impact workload placement and optimization include CPU memory, CPU utilization, processing power, and others. Resources may be allocated to virtual machines if they are compatible with the virtual machine.

For example, resources may be allocated to a virtual machine with compatible hardware, operating system, and software for supporting the resources. *See* Ex. 2022 (Rosenblum) at 3, 4; Ex. 2010 (Newman) at 40, 48.

3. Load Balancing

42. Load balancing distributes VMs across multiple physical servers to achieve specific goals. The most common goal of load balancing is to place VMs so that the workload demands of applications are balanced across the physical servers in a cluster.

4. Monitoring

43. Physical servers provide resources for applications to run or execute. Some of the common resources may include CPU, memory, storage and network. The consumption of these resources in a virtualized environment may be monitored by other services, which collect and analyze usage data of VMs. For example, a monitoring service may show compute usage of a VM in periodic segments over a 24-hour period. Monitoring services may further provide various statistics based on the data collected. Further, users may define various actions that many be taken based on this collected data.

5. Consolidation

44. Consolidation is used to combat the effects of server sprawl by running multiple systems on fewer physical systems. *See* Ex. 1001 ('459 patent) at 2:5-42;

Ex. 2010 (Newman) at 16-17. Virtualization is one type of consolidation strategy that that enables multiple virtualized systems to run on fewer physical systems. *See*, *e.g.*, Ex. 1001 ('459 patent) at 2:13-19; Ex. 2010 (Newman) at 16. Using virtualization, a number of "virtual machines" ("VM") can be consolidated to run on a single physical machine, thereby saving costs and utilizing resources more efficiently. *Id*.; *see also id*. at 22. For example, the operating systems and applications from multiple servers can be consolidated to run on a single physical machine. *Id*. at 22-23.

45. When consolidating systems through virtualization, several factors, including complex systems configurations, diverse business requirements, dynamic workloads, and the heterogeneous nature of distributed systems, can cause incompatibilities between systems that can limit consolidation combinations. Ex. 1001 ('459 patent) at 2:51-60; Ex. 2010 (Newman) at 40. For example, an application that is supported on one operating system (or one operating system version) might not be supported on another. Ex. 2010 (Newman) at 40. Additionally, applications that are migrated from multiple servers to a single server during consolidation may not be compatible and may not behave the same way after consolidation as they did prior to consolidation. *Id*.

VI. THE ’459 PATENT

46. The ’459 patent discloses methods and systems for determining the compatibility of computer systems. As the information technology industry trended away from large, centralized computer systems (such as mainframes) toward smaller, “distributed systems” of multiple computers operating in parallel, significant practical challenges emerged. Ex. 1001 at 1:35-58. For example, while distributed approaches allowed computer resources to be deployed in relatively low-cost increments and with greater flexibility, they led to a proliferation of servers (known as “sprawl”), resulting in an overabundance of processing capacity and posing complex and costly system management challenges. *Id*. at 1:47-2:4.

47. Organizations sought to reduce costs and increase efficiency by consolidating, or “combining,” distributed computing systems. *Id*. at 2:5-42. Choosing a consolidation strategy, however, was often difficult, error-prone, and time consuming due to the virtually infinite number of possible consolidation permutations in large enterprise computing environments, which included suboptimal and incompatible system combinations. *Id*. at 2:43-60. For example, consolidation strategies could employ various techniques of combining computing resources, including “virtualization,” *id*. at 2:14-19, operating system-level (or “OS-level”) “application stacking,” *id*. at 2:20-29, “database stacking,” *id*. at 2:30-36, and

physical consolidation, *id*. at 2:37-42, among others. Choosing a strategy or combination of strategies involved striking a balance between cost reduction and maintaining or enhancing the functionality and reliability of the consolidated system. *Id*. at 2:43-50. Strategy selections were further complicated by various considerations of incompatibility among the computer systems to be consolidated, such as complex systems configurations, diverse business requirements, dynamic workloads, and the heterogenous nature of distributed systems. *Id*. at 2:51-60; *see also id*. at 39:43-52.

48. The '459 patent addresses these problems and discloses novel and inventive solutions for consolidating multiple "source" computer systems onto fewer "target" computer systems. *Id*. at 5:57-64. "Sources" or "source systems" are systems *from* which applications and/or data are to be moved, and "targets," "target servers" or "target systems" are systems *to* which such applications and/or data are to be moved. *Id*. Within the '459 patent's solutions, sources and targets may be physical, virtual, or hypothetical computer systems (that do not yet exist) for planning future activities and performing "what if" consolidation analyses. *Id*. at 5:65-6:10. In addition to computer systems, "systems" may also encompass any entity that is capable of being analyzed for compatibility, such as applications, database instances, servers, and computers. *Id*. at 6:17-33.

49. Among other things, the novel and inventive solutions disclosed in the'459 patent include a "multi-dimensional" compatibility analysis that evaluates the compatibility of systems against other systems based on complex factors that impact whether those systems can be combined. The multi-dimensional compatibility analysis evaluates the compatibility of a "transfer set" of multiple sources (e.g., "N" sources) to be transferred to a target. *Id*. at 29:37-45, 30:36-44. Referring to annotated Figure 3 below, a "transfer" (Transfer 1 and Transfer 2 depicted in **23**) is the movement of a source (S1 and S2) onto a target (S3). *Id*. at 6:52-54. A "transfer set" (**23**) is one or more transfers onto a common target (S3). *Id*.

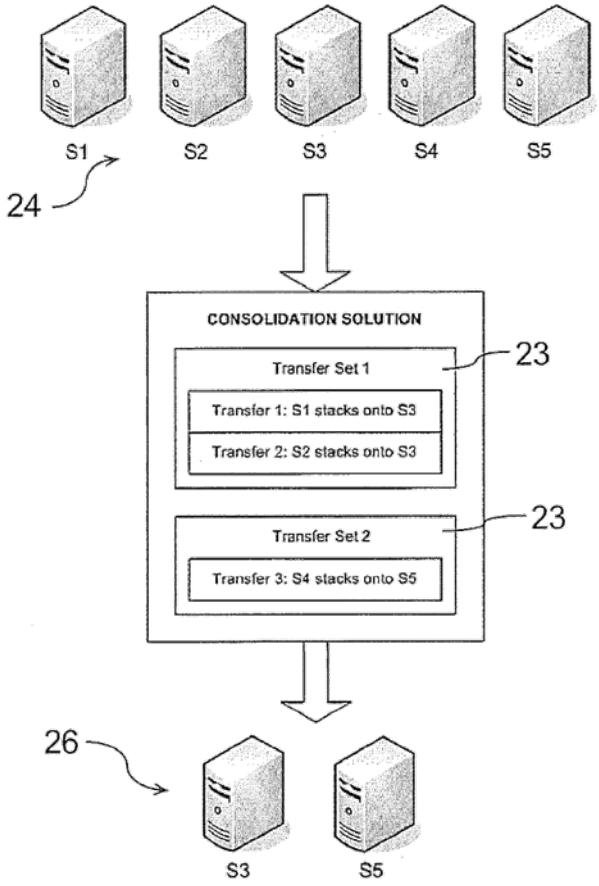

50. The multi-dimensional compatibility analysis may evaluate the compatibility of each source in a transfer set against the target ("N-to-1"). *Id*. at 30:36-42. It may also evaluate the compatibilities of multiple sources in a transfer set amongst each other ("N-by-N"). *Id*. In other words, an N-to-1 compatibility analysis assesses each source system in a transfer set against the target ("intercompatibility"), and an N-by-N compatibility analysis evaluates each source system a transfer set against each of the other source systems ("intracompatibility"). *Id*. at 30:49-52.

51. Factors that impact whether systems can be combined include technical, business, and workload parameters. *Id*. at 5:40-53, 6:66-7:7, 10:21-23, 15:24-29. Examples of technical parameters include the operating system, application settings, and hardware devices of a system. *Id*. at 5:40-42. Examples of business parameters include data segregation requirements, hardware lease agreements, and software licensing agreements, as well as factors related to regulatory compliance and other business objectives. *Id*. at 5:43-48, 10:65-7:1, 39:43-48. Examples of workload parameters include utilization and capacity of system resources, such as processors, memory, throughput, and network bandwidth. *Id*. at 5:49-53.

52. System compatibility with respect to these parameters may be evaluated during the compatibility analysis using "differential rules" and "workload stacking" algorithms. *Id*. at 6:66-7:7. With reference to annotated Figure 1 below, parameters or characteristics of each system, including technical, business, and workload parameters are first obtained through data collection (**18**), and differential rules are then used to evaluate those parameters and determine the compatibility of systems as they relate to technical, business, and workload constraints (**10**). *Id*. at 5:18-24, 5:29-39, 10:21-23, 11:6-9, 15:24-39; 29:48-51; Fig. 1.

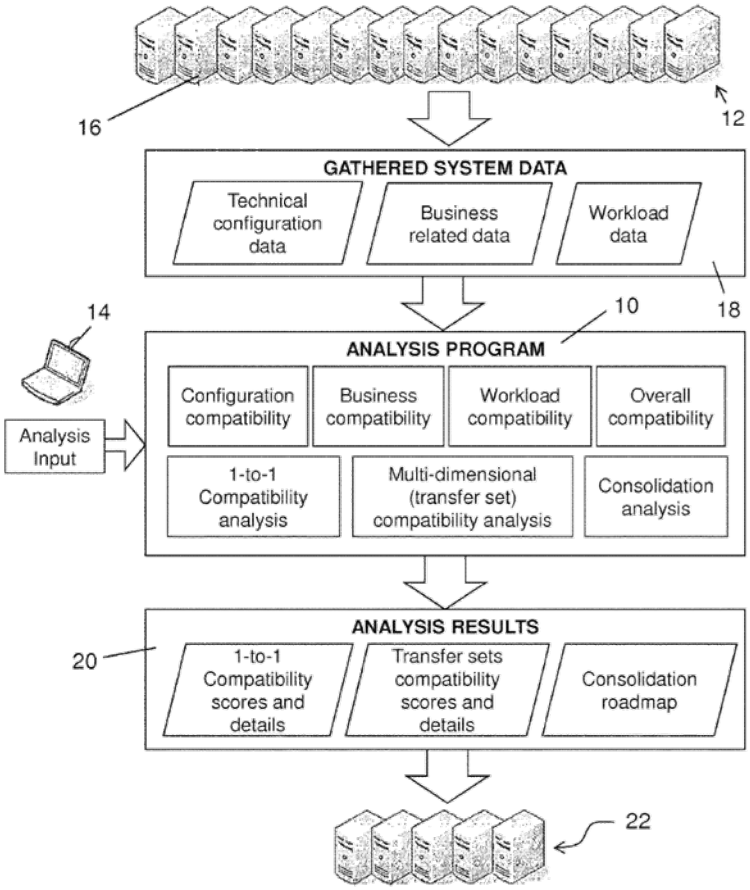

53. The differential rules may be grouped into "rule sets" representing the criteria that allow or prevent systems from being combined for a given consolidation strategy, while meeting requisite compatibility constraints. *Id*. at 7:32-38, 10:21-27. Constraints may include technical constraints, business constraints, and workload constraints. *Id*. at 10:21-23, 15:24-29, 29:48-51. Constraints can be expressed in rule sets, *id*. at 15:27-29, and a rule set can be used to evaluate system compatibility with respect to a constraint or category of constraints, such as technical constraints (*see*, *e.g.*, Fig. 16) or business constraints (*see*, *e.g.*, Fig. 12). *Id*. at 10:9-11, 10:21-

25 (“[T]he differential rule sets 28 are used to evaluate the compatibility of systems as they relate to technical and business related constraints.”), 8:12-13 (“Rule sets 28 examine system compatibility with respect to technical configuration and business-related factors.”). For example, individual business constraints or combinations of business constraints may be expressed in rule sets. *See*, *e.g.*, Fig. 12 (annotated below).

Name	Description
Solaris Kernel Settings	Detailed Analysis of Solaris Kernel Settings
Solaris Name Service Settings	Detailed Analysis of Solaris Name Service Settings
Solaris NVRAM Settings	Detailed Analysis of Solaris EEPROM Settings
UNIX Application Stacking	Detailed analysis rules for OS-level stacking of UNIX apps onto existing or new systems
UNIX Forward Consolidation	Detailed analysis rules for forward consolidation of UNIX applications onto existing infrastructure
Availability Target	Prevent combining servers with differing availability (uptime) targets
Backup Window	Prevent combining servers with differing backup windows
Building	Prevent consolidating servers between buildings
Business Service	Prevent consolidating servers across business services
Combined Constraints	Combined business constraint ruleset
Department	Prevent consolidating servers across departments
Operational Environment	Prevent consolidating servers between environments
Operational Role	Prevent consolidating servers that have differing functional roles
Ownership	Prevent combining servers that belong to different owners
Primary Application	Prevent combining servers that are servicing different applications
Service Plan	Prevent consolidating servers that have different service plans
Validation	analysis
Oracle Data Stacking	Assess data consolidation potential between Oracle instances
Oracle Instance Stacking	Assess affinity for instance-level stacking
SQL Server 2000 (WMI)	Map out data-level consolidation opportunities in MS SQL 2000 environments
Windows File and Print Consolidation	Analyze compatibility of file and print servers
Hardware Models	Map out manufacturers and models
Macro Analysis	Combined Macro Analysis ruleset
OS Map	Map out OS Types, Versions and kernel patch levels
OS Map (with Kernel Bits)	Map out high-level OS affinity regions

Business Service	Prevent consolidating servers across business services

Combined Constraints	Combined business constraint ruleset

Ownership	Prevent combining servers that belong to different owners

Figure 12

54. The ’459 patent discloses improvements over existing approaches to system consolidation. For example, the multi-dimensional compatibility analysis evaluates the compatibility of business constraints *between* or *among* candidate systems for consolidation, and not just against arbitrary criteria (such as lease status or financial savings targets). *Id*. at 6:61-65, 15:24-29. The ’459 patent’s multi-

dimensional analysis also enables the placement of sources on targets to be done in accordance with all three types of constraints (i.e., technical, business, and workload) using a multi-dimensional assessment (e.g., N-to-1 or N-by-N) of system compatibility. *Id*. at 30:36-42; 39:16-31. Prior methods did not evaluate multi-dimensional compatibility with respect to all three types of constraints in a combinatorial manner, thus they produced different results.

A. The Claims of the '459 Patent

55. As demonstrated by representative claim 1 below, the claims of the '459 patent cover embodiments of the '459 patent's multi-dimensional compatibility analysis. Claim 1 recites[2]:

> 1. **(Preamble)** A system for determining placement of source computer systems on target computer systems, the system configured to execute operations causing the system to:
>
> **(1[a])** collect data for a collection of computer systems, the collection of computer systems comprising a plurality of source systems and a plurality of target systems;

[2] Here I include the claim element annotations "Preamble," "1[a1]", "1[a2]," and "1[b]," used in the Petition.

(1[b]) determine a placement of at least one source system from the collection of computer systems on at least one target system from the collection of computer systems by employing the following operations:

(1[c]) evaluate compatibility between a specific source system from the plurality of source systems and a specific target system from the plurality of target systems by evaluating one or more rules that operate against attributes or data relating to the source and target systems being evaluated;

(1[d]) evaluate compatibility between the specific source system from the plurality of source systems and one or more other source systems either already placed on the specific target system, or being evaluated for placement onto the specific target system, to determine if the specific source system can be placed with those other source systems on the specific target system, by evaluating one or more rules that operate against attributes or data relating to the source systems;

(1[e]) evaluate compatibility between the specific source system from the plurality of source systems and one or more other

source systems either already placed on the specific target system, or being evaluated for placement onto the specific target system, to determine if the specific source system can be placed with those other source systems on the specific target system, by evaluating one or more rules that operate against attributes or data relating to the source systems;

(1[f]) evaluate compatibility between the specific source system and the specific any one of the plurality of target system by evaluating the impact on resource utilization of the specific target system of placing the specific source system on the specific target system, in combination with the one or more other source systems, either already placed on the specific target system, or being evaluated for placement onto the specific target system; and

(1[g]) issue instructions to place the at least one source system on the at least one target system in accordance with the determined placement.

Ex. 1001 at 40:2-39 (claim 1) (annotated). Independent claims 32 and 63 are different in that they are directed to a computer implemented method and a non-

transitory computer readable medium, respectively, but they recite similar features as claim 1, and thus I treat claim 1 as representative for purposes of this Declaration. I understand that Petitioner also treated claim 1 as representative for its analysis. Petition at 21-22 ("Specifically, it lacks written description support for the following limitations, which appear in substantially identical form in each of the '459 patent's independent claims (1, 32, and 63).")

56. For example, claim 1 recites, among other things, "evaluate compatibility between the specific source system from the plurality of source systems and one or more other source systems either *already placed* on the specific target system"" Ex. 1001 at 40:18-21 (claim 1 (emphasis added)).

57. Claim 1 further recites "evaluate compatibility … by evaluating the impact on resource utilization of the specific target system of placing the specific source system on the specific target system, in combination with the one or more other source systems, either *already placed* on the specific target system." *Id.* at 40:28-32 (claim 1 (emphasis added)).

58. Claim 1 further recites: "*issue instructions to place* the at least one source system on the at least one target system in accordance with the determined placement." *Id.* at 40:37-39 (claim 1 (emphasis added)).

B. The Prosecution History of the ’459 Patent

59. I understand that the claims of the ’459 patent are presumed valid.

60. I have reviewed the prosecution history of the ’459 patent, including related U.S. patent applications incorporated by reference in the ’459 patent. These applications include U.S. Pat. App. No. 14,/341,471 (the “’471 application”) (Ex. 1008), U.S. Pat. App. No. 11/738,936 (the “’936 application”) (Ex. 1009), 11/535,355 (the “’355 application”) (Ex. 1010), 11/535,308 (the “’308 application”) (Ex. 1011), and U.S. Provisional Application No. 60/745,322 (the “’322 application”) (Ex. 1012).

61. U.S. Patent Application No. 16/687,966 (“the ’966 application”), which issued as the ’459 patent, was filed on November 19, 2019. Ex. 1001 (’459 patent) at (22).

62. As originally filed, the ’966 application included one claim, reproduced below:

> A computer-implemented method for determining a consolidation solution for a plurality of computer systems comprising:
>
> a) obtaining a data set comprising a compatibility score for each pair of said plurality of computer systems, each said compatibility score having been computed using parameters of said plurality of computer systems to indicate a comparable compatibility of respective parameters of one of said plurality of computer systems with respect to another of said plurality of computer systems;

b) using said compatibility scores to determine one or more candidate transfer sets from said plurality of computer systems and, each candidate transfer set indicating one or more source computer systems having compatible parameters that enable the source computer system(s) to be transferred to a target computer system;

c) selecting a desired one of said one or more candidate transfer sets according to a predetermined selection criterion and adding said desired one of said one or more candidate transfer sets to a first candidate consolidation solution;

d) repeating b) and c) by, at each iteration, removing computer systems included in said desired one of said one or more candidate transfer sets and using said compatibility scores corresponding to remaining computer systems to generate at least one additional desired candidate transfer set to be added to said first candidate consolidation solution; and

e) providing said first candidate consolidation solution as a desired consolidation solution when no additional candidate transfer sets can be determined in b).

Ex. 1002 ('966 application) at 63.

63. Dr. Zadok contends that this claim "reflected the analysis program described above and in the rest of the specification." Ex. 1006 (Zadok Decl.), ¶ 55. I disagree that the specification of the '966 application is so limited. Dr. Zadok goes on to argue that the '966 application "did *not* refer to a source system that was 'placed,' let alone 'already placed,' on a target system," and "nor did it refer to 'issuing instructions' to place a source system onto a target system according to a determined placement." *Id.*, ¶ 56. I disagree. As explained below, *see supra*, §§ IX. and X, the as-filed claim may not use the specific claim terms referenced by Dr. Zadok, but it does use concepts that describe the referenced claim terms. For

example, the as-filed claim references both compatibility and consolidation analyses, and also claims "selecting a desired one of said one or more candidate transfer sets according to a predetermined selection criterion and adding said desired one of said one or more candidate transfer sets to a first candidate consolidation solution," "generat[ing] at least one additional desired candidate transfer set to be added to said first candidate consolidation solution," and "providing said first candidate consolidation solution as a desired consolidation solution." Ex. 1002 ('966 application) at 63.

64. On July 2, 2020, Patent Owner filed a preliminary amendment, canceling the previously-filed claim and adding 57 new claims. Ex. 1003 ('966 application, July 2, 2020 Preliminary Amendment). On October 5, 2020, the examiner issued a Notice of Allowability for each of the amended claims. Ex. 1004 ('966 application, October 5, 2020 Notice of Allowability). The Notice of Allowability included a statement from the examiner that limitations including "e.g., "…evaluate compatibility between any of the two or more of the plurality of source systems…in combination with any other source systems…determine a placement of at least one source system of the collection of system…by employing the evaluating operations…" were not taught by the prior art. *Id.* at 2.

65. On December 10, 2020, Patent Owner filed a voluntary amendment. Ex. 1005 ('966 application, December 10, 2020 Voluntary Amendment). Patent Owner provided remarks explaining that "[t]he claims have been amended to rearrange the "determine/determining" operation and to clarify the protection being sought." *Id.* at 14. Four claims—claims 4, 25, 32, and 53—were canceled and new claims 59-68 were added. *Id.* Patent Owner also stated "no new subject matter has been added by way of these amendments." *Id.*

66. The '459 patent issued on March 16, 2021. Ex. 1001 ('459 patent) at (45).

VII. LEVEL OF ORDINARY SKILL

67. I understand that Petitioner sets forth a level of ordinary skill for a person of ordinary skill in the art in the field of the '459 patent on pages 18-19 of the Petition as follows:

> A person of ordinary skill in the art ("POSITA") for the technology of the '459 patent would have at least a bachelor's degree in computer science, computer engineering, electrical engineering, or a related degree and 2–3 years of experience related to with administering remote computing environments. This experience could include, for example, work in cloud

> computing, virtualization, network communications, client-server communication analysis, load balancing, stress testing, and/or system consolidation. A POSITA would have been familiar with designing, implementing, monitoring, evaluating, optimizing, and managing physical and virtual computing environments. More education would be acceptable in place of work experience and vice versa. (Declaration of Dr. Erez Zadok (Ex. 1006), ¶ 59.)

68. I meet this requirement for the definition of a person of skill in the art because I have a Ph.D. in Electrical Engineering and Computer Science and over ten years of experience working in the areas of cloud computing, workload modeling, virtualization, and benchmarking, which relate to administering distributed computing environments. As such, I am qualified to give an opinion on what a person of skill in the art would have understood based on the disclosures in the references discussed in this proceeding.

69. I further understand that Dr. Zadok's understanding is that the '459 patent's priority date is its November 19, 2019 filing date. Ex. 1006 (Zadok Decl.), ¶ 29. For the purposes of my opinions in this Declaration, I have been asked to assume the '459 patent is entitled to the earliest priority date listed on the face of

the ’459 patent, April 21, 2006. Therefore, when I refer to one of ordinary skill in the art in my opinion below regarding the ’459 patent, I am referring to one of ordinary skill in the art as of April 21, 2006. However, my opinions are not affected and would not change if November 19, 2019 (or any other date between April 21, 2006 and November 19, 2019) is determined to be the effective filing date of the ’459 patent.

VIII. CLAIM CONSTRUCTION

70. The Petition included a claim construction analysis for the terms “source system” and “target system.” The Petition says that these terms should be given the same construction here as those adopted by the Board in a different proceeding instituting *inter partes* review of U.S. Patent No. 10,523,492. *See* Petition at 19-20 (citing *VMware, Inc. v. Cirba IP Inc.*, Case No. IPR2021-00008, Paper No. 10, Decision Granting Institution of *Inter Partes* Review, at 8 (PTAB May 7, 2021). Dr. Zadok also opines that “a ‘source system’ is ‘a system from which applications and/or data are to be moved,’” and “a ‘target system,’ is stated as ‘a system to which [a source system’s] applications and/or data are to be moved.’” Ex. 1006 (Zadok Decl.), ¶ 60.

71. The Petition did not include a claim construction analysis for any other claim terms. Dr. Zadok also opines that “[f]or all other terms, if any, I have been

instructed to perform my analysis with the understanding that the terms of the Challenged Claims should be afforded their plain and ordinary meaning." *Id.*, ¶ 61.

72. The Board determined that it was "not necessary to provide an express interpretation of the terms "source system" and "target system" at this stage of the proceeding. Paper No. 12 at 11. The Board further observed that, "although Petitioner makes arguments based on the meaning of the term 'placement' in the claims, Petitioner never conducts a claim construction analysis for the terms 'placed' or 'placement.'" *Id.* at 22.

73. Whether the Board ultimately determines that any claim constructions are or are not necessary to decide the issues raised in the Petition, it is my opinion that the '459 patent's specification is instructive as to the meaning and/or scope of certain claim terms, as understood by a POSA.

74. Specifically, the '459 patent expressly defines "source system" and "target system," which are terms used throughout the independent claims, while also providing additional guidance as to the meaning and scope of each term. Additionally, the independent claims of the '459 patent use the phrases "placed" and "placement," which have a plain and ordinary meaning informed by the surrounding claim language, as well as the definitions and guidance provided by the '459 patent's specification.

A. Claim Construction Analysis

1. "source system"

75. The '459 patent's specification expressly defines "source system" to mean "a system from which applications and/or data are to be moved." Ex. 1001 ('459 Patent) at 5:57-58.

76. Also relevant to the meaning of "source system" as a phrase is the '459 patent's guidance as to the meaning of "system." In the '459 patent, source systems and target systems may be physical, virtual, and even hypothetical computer systems (that do not yet exist) for planning future activities and performing "what if" consolidation analyses. *Id*. at 5:65-6:10. In addition to computer systems, the '459 patent indicates that "systems" may also encompass any entity that is capable of being analyzed for compatibility, such as applications, database instances, servers, and computers. *Id*. at 6:17-33.

77. Further relevant to the meaning of "source system" as a phrase is the '459 patent's guidance as to "source systems" not only being systems from which applications and/or data "are to be moved," but also systems from which they "are moved." This movement is accomplished through, for example, consolidation, which "can be considered to include one or more 'transfers,'" where "[t]he actual transfer describes the movement of a single source entity onto a target" and "[t]he

transfer type (or consolidation strategy) describes how a source entity *is* transferred onto a target, e.g. virtualization, OS stacking etc." *Id.* at 6:35-41. The '459 patent also indicates that its compatibility analyses "can be performed on an analysis with pre-existing source-target transfers," *Id.* at 34:44-46, and that "the compatibility analysis can evaluate the incremental effect of adding other source systems … to the specified transfer sets," such that a source system "can be individually assessed against [a] transfer set." *Id.* at 30:18-26. As used in the '459 patent, a "transfer set" is "one or more transfers that involve a common target, wherein the set specifies one or more source entities, the target and a transfer type." *Id.* at 6:42-45. As an example of a transfer set, the '459 patent identifies a transfer set T1, comprised of source systems S1, S2, and S3 stacked onto S4, and represented as follows:

T1 (S1, S2, S3 stacked onto S4)

Id. at 29:59-30:7.

78. Further relevant to the meaning of "source system" as a phrase is the '459 patent's guidance that its compatibility analyses may "be used for purposes other than consolidation such as capacity planning, regulatory compliance, change, inventory, optimization, administration etc. and any other purposes where compatibility of systems is useful for analyzing systems 16." *Id.* at 39:44-53.

79. Thus, it is my opinion that "source system" would have been understood by a POSA to mean "a physical, virtual, or hypothetical system from which applications and/or data are moved or are to be moved."

2. "target system"

80. The '459 patent's specification expressly defines "target system" to mean "a system to which such applications and/or data are to be moved." Ex. 1001 at 5:59-60. The phrase "such applications and/or data" refers to applications and/or data from a source system. *Id*. at 5:57-58 Thus, "target system" would have been understood to mean "a system to which the applications and/or data from a source system are to be moved."

81. As with "source system," the meaning of "target system" as a phrase is informed by the '459 patent's guidance as to the meaning of "system." As explained above, the '459 patent indicates that source systems and target systems may be physical, virtual, and hypothetical computer systems, or any entity that is capable of being analyzed for compatibility, such as applications, database instances, servers, and computers. *Id*. at 5:65-6:10, 6:17-33.

82. Also as with "source system," further relevant to the meaning of "target system" as a phrase is the '459 patent's guidance as to "target systems" not only being systems to which applications and/or data "are to be moved," but also systems

to which they "are moved." This movement is accomplished through, for example consolidation, *id.* at 6:35-41, and also "can be performed on an analysis with pre-existing source-target transfers," *id.* at 34:44-46. *See also id.* at 6:42-45, 29:59-30:7, 30:18-26.

83. Also, as with "source system," further relevant to the meaning of "target system" as a phrase is the '459 patent's guidance that its compatibility analyses may "be used for purposes other than consolidation such as capacity planning, regulatory compliance, change, inventory, optimization, administration etc. and any other purposes where compatibility of systems is useful for analyzing systems 16." *Id.* at 39:44-53.

84. Thus, it is my opinion that "target system" would have been understood by a POSA to mean "a physical, virtual, or hypothetical system to which applications and/or data are moved or are to be moved."

3. "place" and "placement"

85. The terms "place" and "placement" are used only in the claims of the '459 patent, and not otherwise used or defined in the specification.

86. I understand that, where the specification does not expressly redefine or restrict the common understanding of claim terms, the plain and ordinary meaning of those claim terms should apply.

87. Relevant to the meaning of "place" and "placement as phrases is the '459 patent's guidance as to the meaning of the terms beyond the act of "moving." In the '459 patent, actions contemplated with respect to "place" and "placement" include: "consolidating" (*id.* at 2:5-7); "stacking" (*id.* at 2:20-22, 17:56-60, Fig. 3); "moving" (*id.*); "forward consolidat[ing]" (*id.* at 6:6-10); "transferring (*id.* at 6:35-41, 6:42-45, 13:36-38, 34:64-35:6, Fig. 3, Fig. 27); "migrating" (*id.* at 24:65-67); "arranging" (*id.* at 27:9-12); "packing" (*id.* at 34:64-35:6); "fitting" (*id.*).

88. Further relevant to the meaning of "place" and "placement" as phrases is the common use within the virtualization and information technology infrastructure industry of these terms consistent with the guidance of the '459 patent.

89. Further relevant to the meaning of "place" and "placement" as phrases is the common understanding of the meaning of those terms as indicated by dictionaries defining those terms as "to put in or as if in a particular place," "to put in a particular state," and "to distribute in an orderly manner." *See* Ex. 2007 (*Place*, Dictionary by Merriam-Webster, https://www.merriam-webster.com/dictionary/place (last visited March 4, 2022)); Ex. 2008 (Merriam-Webster Dictionary (2004)) at 549. Similarly, the term "move" is not commonly

understood to be coextensive with the terms "place" and "placement." *See* Ex. 2009 (Roget's II The New Thesaurus (3d ed. 2003)) at 736.

90. Thus, it is my opinion that "place" and "placement" would have been understood by a POSA, consistent with their plain and ordinary meaning as used in the '459 patent, to not be limited to the action of "moving." It is also my opinion that a POSA would have understood the plain and ordinary meaning of these terms not to be limited to the applications and/or data within a source system, and to also include placing an entire source system (and its contents) on a target system.

IX. THE ZADOK DECLARATION

91. I have reviewed the Declaration of Dr. Erez Zadok, Ph.D. (Ex. 1006) submitted by Petitioner in this proceeding. I have been asked torespond only to certain portions of Dr. Zadok's declaration, as set forth below. That I have not addressed any particular aspect of Dr. Zadok's declaration does not mean (and should not be interpreted to mean) that I agree with it. I reserve the right to provide my opinions and respond to any part of Dr. Zadok's declaration, as appropriate, during the trial.

A. Dr. Zadok's Opinions With Respect to the "Already Placed" Limitation

92. Dr. Zadok's declaration includes opinions that none of the compatibility analyses in the '459 patent (1-to-1, N-to-1, and N-by-N) "discloses or

describes the concept recited … in the claims of evaluating a specific source system against 'one or more other source systems [] already placed on the specific target system." *See* Ex. 1006 (Zadok Decl.) ¶¶ 65-66.

93. However, Dr. Zadok's opinion for each of the three compatibility analyses consists of little more than a single paragraph with a limited excerpt from the '459 patent specification and then making conclusory statements about the purported lack of written description without further analysis. *See id.* at ¶ 67 (citing 14 lines for 1-to-1 compatibility analysis); *id.* at ¶ 68 (citing 9 lines for N-to-1 compatibility analysis); *id.* at 69 (citing 7 lines for N-by-N compatibility analysis).

94. For his analysis of the '459 patent's consolidation analysis, Dr. Zadok's opinions expand from one paragraph to three, but the depth of his analysis fares no better, again citing limited excerpts from the '459 patent specification and then making conclusory statements about the purported lack of written description without further analysis. *See id.* at ¶¶ 70-72 (citing 16 lines for consolidation analysis).

B. Dr. Zadok's Opinions With Respect to "source system" Limitation

95. Dr. Zadok opines that a "source system" in the '459 patent "is defined in the context of potential future transfers that may occur after moving applications and/or data from a source system to a target system. *See id.* at ¶ 74. Dr. Zadok goes

on to provide a host of further opinions concerning the purported meaning of "source system," all without citation to the '459 patent specification (or any other disclosure from the '459 patent priority chain). *See id.* For the reasons discussed above in my claim construction analysis with respect to "source system," (*see supra* § VIII.A.1), as well as below in my analysis of the Petition, (*see infra* § X.A). I disagree with Dr. Zadok's opinions concerning the scope of "source system."

96. Dr. Zadok also offers an opinion based on an example of a source system that is a physical computer to conclude that "[t]he '459 patent's claims … contradict the source system definition by suggesting that a source system in its entirety … is transferred or placed on a target system instead, and not that the source system's application and/or data alone are moved." *Id.* at ¶¶ 74-75. I disagree with Dr. Zadok's opinion that the '459 claims contradict the definition of "source system" in the '459 patent. I further note that a district court recently rejected this same argument by Petitioner in a related proceeding. Ex. 2005 (2/24/2022 claim construction Memorandum Opinion) at 14.

97. Dr. Zadok's focus on "moving … entire physical source system[s]," *see id.* at ¶¶ 75, 77, is misplaced and ignores the disclosures of the '459 patent specification that source systems and target systems may be physical, virtual, and hypothetical computer systems, or any entity that is capable of being analyzed for

compatibility, such as applications, database instances, servers, and computers. Ex. 1001 ('459 patent) at 5:65-6:10, 6:17-33.

98. Dr. Zadok also opines, again with minimal citation, that the '459 patent "focuses on the consolidation of systems." Ex. 1006 (Zadok Decl.) ¶¶ 76. I disagree. The '459 patent does disclose consolidation analysis, but overwhelmingly focuses on compatibility analyses. Indeed, the title of the '459 patent is "Method and System for Determining ***Compatibility*** of Computer Systems." Ex. 1001 ('459 patent) at (12) (emphasis added).

C. Dr. Zadok's Opinions With Respect to "Issue Instructions" Limitation

99. Dr. Zadok's declaration includes opinions that "nowhere does the '459 patent explain how to issue instructions to carry out the results of its analyses" and "stops at providing analysis results without any description of how to carry out those results." *See* Ex. 1006 (Zadok Decl.) ¶ 79. Dr. Zadok goes on to opine that "[t]here is nothing in Figure 7 (or the rest of the patent) that indicates the program actually carries out the results, such as by issuing instructions," *id.* at ¶ 81, and that "[t]he examples of the compatibility and consolidation analyses all stop at choosing or displaying results—they do not disclose executing on the results," *id.* at ¶ 82. I disagree, for reasons explained below. *See infra*, § X.B.

X. THE GROUNDS OF UNPATENTABILITY ASSERTED IN THE PETITION FAIL TO DEMONSTRATE THE '459 PATENT'S CHALLENGED CLAIMS LACK ADEQUATE WRITTEN DESCRIPTION

A. The Pre-AIA '936 Application Discloses Evaluating A Source System With "Other Source Systems … Already Placed On The Specific Target System."

100. The '936 application describes computer systems that may be "hypothetical systems … defined and included in the analysis to evaluate various types of 'what if' consolidation scenarios." Ex. 1009 at [0082]. The '936 application further describes these "hypothetical source systems can be used to simulate the case where a new application is to be introduced into the environment 12 and 'forward consolidated' onto existing target systems 16." *Id.* Fig. 12 of the '936 application, which is described as "a table showing example rule sets," includes a rule set for "UNIX Forward Consolidation," whose description reads "[d]etailed analysis rules for forward consolidation of UNIX applications onto existing infrastructure." *Id.* at Fig. 12, [0030], [00117], [00155]. A POSA would understand that the "forward consolidation" use case described in the '936 application may involve the creation of hypothetical source systems to determine their compatibility, evaluated against pre-existing, already placed, source systems on existing target systems.

101. The '936 application describes a 1-to-1 compatibility analysis where the compatibility of every possible source-target pair combination is evaluated on a

1-to-1 basis. *Id.* at [0095]. The '936 application further describes that "[i]n practice, it may be prudent to consolidate systems 16 incrementally and assess the impact of each transfer before proceeding with additional transfers." *Id.* A POSA would understand that the described incremental transfer of source systems onto target systems implicates placing source systems onto target systems on which other source systems have already been placed.

102. The '936 application describes multi-dimensional rule-based compatibility analyses where "a transfer set can include multiple sources (N) to be transferred to the target" and "the analysis may evaluate the compatibility of sources amongst each other (N-by-N) as well as each source against the target (N-to-1)." Ex. 1009 at [00321]. As above with respect to the 1-by-1 compatibility analysis, the '936 application explains the multi-dimensional rule-based compatibility analyses "can evaluate the incremental effect of adding other source systems (specified in the analysis input) to the specified transfer sets." *Id.* at [00319]. The '936 application provides two examples of evaluating source systems against existing transfer sets, a first where "the compatibility of the source systems S5 to S20 can be individually assessed against the transfer set T1," and a second where "the compatibility of the source systems S1 to S4 and S11 to S20 can be assesses with respect to the transfer set T2." *Id.* The '936 application describes transfer set T1 as being comprised of

"[source systems] S1, S2, S3 stacked onto [target system] S4" and transfer set T2 as being comprised of "[source systems] S5, S6, S7, S8, S9 stacked onto [target system] S10." *Id.* at [00316]. A POSA would understand that the "individual assessment" of the identified source systems against transfer sets comprised of already placed source systems describes evaluation of source systems against already placed source systems, and that the source systems of the exemplary transfer sets do not simply disappear or otherwise cease to exist once placed for purposes of the described compatibility analyses.

103. The '936 application describes compatibility analyses that "can be performed on an analysis with pre-existing source-target transfers." *Id.* at [00345]. Systems in this analysis input may be designated as "a source only, as a target only or as either a source or a target." *Id.* These system designations "serve as constraints when defining transfers in the context of a compatibility analysis." *Id.* A POSA would understand that the transfer sets in such an analysis, like the transfer set T1 and transfer set T2 discussed in the above example, may represent pre-existing source-target transfers, where transfers of source systems to a specific target system have already occurred. In such an analysis scenario, the transfer set would include "already placed" source systems.

104. As is the case for the disclosed compatibility analyses, the ’936 application also describes evaluation of source systems against already placed source systems in the context of its disclosed consolidation analyses. The ’936 application discloses an “auto fit algorithm” used by its consolidation analyses to search for the best consolidation solution. *See, e.g.*, i*d.* at [0051], [00111], [00342], [00347], [00361], Fig. 30. The ’936 application describes that, for the consolidation solution auto fit algorithm, “candidate transfer sets are … compiled from the source-target candidates,” and that “[f]or each target, the candidate sources are sorted in descending order and are incrementally stacked onto the target." *Id.* at [00366]. The ’936 application further describes that, “[t]he transfer set score is computed as each source is stacked and if the score is above the minimum allowable score, the source is added to the transfer set 23.” *Id.* As discussed above, a transfer set as described in the ’936 application consists of source systems already placed on a specific target system (e.g., T1 (S1, S2, S3 stacked onto S4)), and would understand this disclosure to describe evaluating source systems against already placed source systems in existing transfer sets for purposes of the disclosed consolidation analyses.

105. Paragraph 345 of the ’936 application describes “analysis with pre-existing source-target transfers.” *Id.* at [00345]. A POSA would understand that this “directed analysis” refers to a use case where the disclosed compatibility

analyses begin with a partial consolidation solution already in place, and then add additional placements using the pre-existing source-target transfers as constraints. A POSA would further understand the pre-existing source-target transfers represent how the compatibility analysis models "already placed" systems, which in the context of the '936 application, may be real or virtual or hypothetical systems. *Id.* at [0082] – [0085].

B. The Pre-AIA '936 Application Discloses "Issuing Instructions"

106. The '936 application discloses consolidation solutions (or "roadmaps") generated by the disclosed consolidation analyses, which provide "one or more transfer sets 23 based on a common pool of source and target entities," and can be represented visually. *Id.* at [0088], [00305]. The '936 application further describes that, once the analysis input is specified and the consolidation analysis is executed, "the chosen transfer sets are presented," including "a consolidation summary" and a list of "the actual transfers proposed by the consolidation analysis." *Id.* at [00381]. An "actual transfer" in the context of the '936 application is described as "the movement of a single source entity onto a target, wherein the specification identifies the source, target and transfer type." *Id.* at [0086]. The '936 application also explains that "[t]he transfer type (or consolidation strategy) describes how a source entity is transferred onto a target, e.g., virtualization, OS stacking, etc." *Id.* A POSA

would understand that the consolidation solution roadmaps described by the ’936 application correspond to the “determined placement” claim term of independent claims 1, 32, and 63 of the ’459 patent. A POSA would further understand that the further actual transfers described by the ’936 application correspond to the “issue instructions” and “issuing instructions” limitations of the ’459 patent.

107. I confirm that the contents of this Declaration are true to the best of my knowledge and belief insofar as it states facts and that it contains my honest opinions on the matters upon which I have been asked to give them.

108. I declare under penalty of perjury under the laws of the United States ofAmerica that the foregoing is true and correct to the best of my knowledge.

Executed on this 4th day of March 2022, in Johns Creek, GA.

Vijay K. Madisetti, PhD