Filed on behalf of Veeam Software Corporation

By: Lori A. Gordon
         Michael Q. Lee
         Sterne, Kessler, Goldstein & Fox PLLC
         1100 New York Avenue, NW
         Washington, D.C.
         Tel: (202) 371-2600
         Fax: (202) 371-2540

UNITED STATES PATENT AND TRADEMARK OFFICE

_____

BEFORE THE PATENT TRIAL
AND APPEAL BOARD

_____

**PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 7,480,822**

Table of Contents

Veeam Software Corporation ("Veeam" or "Petitioner") requests that the United States Patent Office institute an *inter partes* review of claims 1, 2, 6, 7, and 12-15 (collectively, the "challenged claims" or "claims under review") of United States Patent No. 7,480,822 to Arbon *et al.* ("the '822 patent".) According to Office records, the '822 patent is assigned to Symantec Corporation ("Symantec" or "Patent Owner".) A copy of the '822 patent is provided as VSC 1001.

I.   Mandatory Notice (37 C.F.R. § 42.8(a)(1))

REAL PARTY IN INTEREST: The real party-in-interest of Petitioner is Veeam Software Corporation ("Veeam").

RELATED MATTERS: U.S. Patent No. 7,480,822 is involved in the following current proceeding that may affect or be affected by a decision in this proceeding: *Symantec Corporation v. Veeam Software Corporation*, No. 3:12-cv-5443-SI (N.D. C.A.)(consolidated with 3:12-cv-00700-SI)(N.D.C.A.).

SERVICE INFORMATION: Please address all correspondence to the lead counsel as shown below. Petitioner consents to electronic service by email at the email address provided below.

LEAD AND BACKUP COUNSEL:

Lead Counsel	Backup Counsel
Lori A. Gordon, Reg. No. 50,633	Michael Q. Lee, Reg. No. 35,239
STERNE, KESSLER, GOLDSTEIN & FOX	STERNE, KESSLER, GOLDSTEIN & FOX
1100 New York Avenue, N.W.	1100 New York Avenue, N.W.
Washington, D.C. 20005-3932	Washington, D.C. 20005-3932
Tel.: 202-772-8862	Tel.: 202-772-8674
Fax: 202-371-2540	Fax: 202-371-2540
lgordon-PTAB@skgf.com	mlee-PTAB@skgf.com

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

It is certified by the undersigned and the Petitioner, Veeam, that the '822 Patent is available for *inter partes* review. Because the '822 Patent has an effective filing date of July 13, 2005, 35 U.S.C. § 311(c)'s timing requirements do not apply. *See* AIA Technical Corrections Bill, H.R. 6621, 112th Cong. § 1(d)(1) (2013) (enacted).

It is certified by the Petitioner, Veeam, that Veeam is not estopped from requesting an *inter partes* review of the claims under review of the '822 patent.

III. Identification of Challenge (37 C.F.R. § 42.104(b))

A. Overview

The '822 Patent relates to "restoring access to data and other operation states of multiple primary computing systems, even if the primary computing systems are no longer available." ('822 Patent, 1:55-58.) As part of the disclosed process, the running states of a plurality of primary computing systems are captured and subsequently restored onto a single computing system. ('822 Patent, 7:30-43.) Each of these "captured running states" interface with a common virtualization component on the single computing system. ('822 Patent, 6:4-8.) The '822 Patent implies that the common virtualization component operates the captured running states as virtual machines running on the single computing system. ('822 Patent, 6:17-21) (Shenoy Declaration, ¶ 23)[1]. After the running states are captured, the captured running states are booted in an identified order on the single computing system. ('822 Patent, 7:1-10)

The text of the two challenged independent claims 1 and 15 are reproduced below. For ease of discussion, reference labels have been assigned to each claim limitation. For example, 1[A] refers to the limitation "an act of accessing a captured running state of each of the primary computing systems on the single computing system" of claim 1. Claims 1 and 15 contain nearly identical claim

---

[1] The Shenoy Declaration is provided as VSC 1008. Citations to the Shenoy Declaration will be given as (Shenoy Declaration, ¶ #).

limitations (claim 1 is a method, while claim 15 is a computer program product having instructions). For example, claim limitation 1[P] is equivalent to claim limitation 20[P]; claim limitation 1[A] is equivalent to claim limitation 20[A]; etc. Accordingly, both claims are discussed concurrently throughout the petition.

	Claim 1
[P]	A method for restoring access to running states of a plurality of primary computing systems onto a single computing system, the method comprising the following:
[A]	an act of accessing a captured running state of each of the primary computing systems on the single computing system,
[B]	wherein the captured running states each include, or are altered to include, at least one device driver that is configured to interface with a common virtualization component that runs on the single computing system,
[C]	wherein the common virtualization component is configured to at least indirectly interface with hardware on the single computing system that operates at least part of the hardware using a different interface than the at least one device driver is configured to interface with;
[D]	an act of identifying a boot order for each of the primary computing systems; and
[E]	an act of booting the captured running states for each of the primary computing systems in the identified boot order on the single computing system.

	Claim 15
[P]	A computer program product comprising one or more computer-readable storage media having thereon computer-executable instructions that, when executed by one or more processors of a single computing system, cause the single computing system to perform a method for restoring access to running states of a plurality of primary computing systems onto the single computing system, the method comprising the following:
[A]	an act of accessing a captured running state of each of the primary computing systems on the single computing system,
[B]	wherein the captured running states each include, or are altered to include, at least one device driver that is configured to interface with a common virtualization component that runs on the single computing system,

[C]	wherein the common virtualization component is configured to at least indirectly interface with hardware on the single computing system, at least part of the hardware using a different interface than the at least one device driver is configured to interface with;
[D]	an act of identifying a boot order for each of the primary computing systems; and
[E]	an act of booting the captured running states for each of the primary computing systems in the identified boot order on the single computing system.

B. Claim Construction

Except as set forth in this section, all claim terms should be given their broadest reasonable interpretation.

1. "[captured] running state"

Nearly all of the challenged claims cite either a "running state" or a "captured running state" associated with a computing system. The '822 Patent provides many examples of a "running state":

> The running state may be an entire volume or set of volumes that includes application data, the executable instructions for application programs, the executable instructions for the operating system, and configuration information for the application programs and the operating system. The running state may also include compressed or encrypted forms of the volume. The running state may also include just a portion of the volume, just specific files or folders, or transformations of instructions or data on the volume. ('822 Patent, 4:56-64)

Therefore, the broadest reasonable interpretation of "running state" is "data or instructions associated with a computing system."

The meaning of "captured running states" is set forth by the claims and the specification. For example, the first method step of claim 1 recites "accessing a captured running state *of each of the primary computing systems on the single computing system*." The specification of the '822 patent states that "[e]ach of the running states is captured via the use of capture channel 351. The capture channel 351 includes the mechanisms used to *acquire the running state from the respective primary computing systems, and provide the running state to the recovery computing system* 350." ('822 Patent, 5:1-5; emphasis added.) Based on these disclosures a "captured running state" is interpreted as "a running state of a primary computing system that is accessed on a single computing system." Further, the captured running states are configured to interface with a common virtualization component ('822 Patent, 6:4-8.) Therefore, the captured running state is executed within a virtual machine that is interfacing with the common virtualization component. (Shenoy Declaration, ¶ 28.) The captured running states are accessible as data stored on virtual disks of the virtual machines running on the single computing system. (Shenoy Declaration, ¶ 28.) Since a virtual machine is essentially defined by its data (i.e., captured running state) any changes or

interactions performed with the virtual machine are considered to also be performed with the captured running state. (Shenoy Declaration, ¶ 29.)

C.   All Applied References are Prior Art to the '822 Patent

The '822 Patent was filed on July 13, 2005 and does not claim the benefit of an earlier filed application. Therefore, the earliest possible priority date is July 13, 2005. This petition cites six references, each of which is prior art:

1.   Microsoft Virtual Server Administrator Guide ("VS Guide," provided as VSC 1002) – VS Guide is § 102(a) prior art to the '822 Patent. VS Guide is a help file included in the Microsoft Virtual Server 2005 product and was made available to the public upon the release of the product in September, 2004.

2.   Microsoft Virtual Server Migration Toolkit Whitepaper ("VSMT," provided as VSC 1003) – VSMT published in October 2004, and is therefore prior art under at least 35 U.S.C. § 102(a).

3.   VMware P2V Assistant User's Manual ("P2V Manual," provided as VSC VSC 1004) P2V Manual was provided to the public as a .pdf on the VMware website as early as June 17, 2004[2], thus making it prior art under at least 35 U.S.C. § 102(b).

---

[2] Date is documented on pg. 22 of the P2V Manual and is corroborated via Internet Archive: http://web.archive.org/web/20040701021053/http://www.vmware.com/support/p2v2/

4. ESX Server 2.1 Administration Guide ("ESX Guide," provided as VSC 1005) – ESX Guide published on March 31, 2004, and is therefore prior art under at least 35 U.S.C. § 102(b).

5. Rob's Guide to Using VMWARE ("Rob's Guide," provided as VSC 1006) – Rob's Guide published in July, 2004, and is therefore prior art under at least 35 U.S.C. § 102(a).

6. U.S. Patent No. 8,209,680 to Le et al. ("Le," provided as VSC 1007) – Le was filed on June 30, 2003 and claims priority to a provisional application filed on April 11, 2003, and is therefore prior art under at least 35 U.S.C. § 102(e).

None of the art cited in this petition appears on the face of the '822 Patent.

D. Grounds of Challenge

This petition presents four grounds of unpatentability:

- **Ground 1:** The combination of VS Guide and VSMT render claims 1, 2, 6, 7, and 12-15 unpatentable under 35 U.S.C. § 103(a.)

- **Ground 2**: The combination of P2V Manual and ESX Guide render claims 1, 2, 6, 7, and 12-15 unpatentable under 35 U.S.C. § 103(a.)

- **Ground 3**: The combination of Rob's Guide and ESX Guide render claims 1, 2, 6, 7, and 12-15 unpatentable under 35 U.S.C. § 103(a.)

- **Ground 4**: The combination of Le and ESX Guide render claims 1, 2, 6, 7, and 12-15 unpatentable under 35 U.S.C. § 103(a.)

E. Ground 1: Claims 1, 2, 6, 7, and 12-15 of the '822 Patent are Obvious over VS Guide in view of VSMT.

VS Guide is an administrator guide provided as a help file to customers that purchase the Microsoft Virtual Server 2005 software product. (Shenoy Declaration, ¶ 33.) VS Guide explains to users how to "create and run one or more virtual machines, each with its own operating system, on a single physical computer." (VS Guide, pg. 1.) VS Guide also describes the computer architecture for running multiple virtual machines on a single physical computer using the Microsoft Virtual Server software: "[s]tarting from the bottom of the logical stack, the host operating system manages the physical computer. The Virtual Machine Monitor (VMM) virtualization layer manages virtual machines, providing the software infrastructure for hardware emulation. Each virtual machine consists of a set of virtualized devices." (VS Guide, pg. 124.) VS Guide further describes how to identify a boot order for the virtual machines.

VS Guide discloses substantially all of the limitations of the challenged claims. But, to the extent that VS Guide does not explicitly disclose that the "running states" are "captured," the related Virtual Server Migration Toolkit whitepaper (VSMT) provides these teachings. For example, VSMT describes that the toolkit provides "a reliable and repeatable process for migrating an operating system and installed applications from a physical server to a virtual machine *running on Virtual Server 2005*." (VSMT, pg. 5; emphasis added.)

FIG. 1 from the Shenoy Declaration below illustrates how the various claimed elements of the '822 patent match with the elements disclosed by both VS Guide and VSMT.
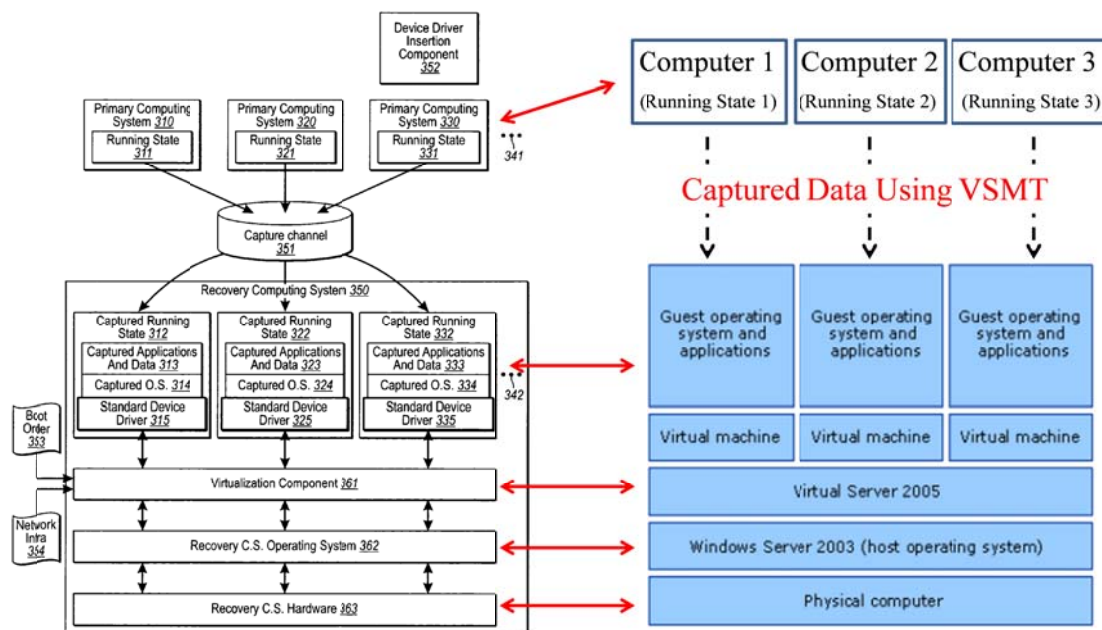
FIG. 1 - (Shenoy Declaration, pg. 12)

The diagram on the left has been reproduced from Figure 3 of the '822 Patent while the diagram on the right has been reproduced from the figure on page 124 of the VS Guide and annotated by Dr. Shenoy. Dr. Shenoy has annotated the figure from VS Guide to illustrate that VSMT discloses capturing the running states of primary computing systems as illustrated in the top portion of the right-hand diagram in FIG. 1. (Shenoy Declaration, ¶ 36.) The claimed captured running states are the guest operating systems and applications, the claimed common virtualization component is the virtual server 2005 software, the operating system of the recovery computing system is the same as the windows server 2003 host

- 10 -

operating system, and the hardware of the recovery computing system is the same as the physical computer upon which Virtual Server 2005 is running. (Shenoy Declaration, ¶ 37.)

A person of ordinary skill in the art would have combined the teachings of VS Guide and VSMT. VS Guide teaches how to run and manage multiple virtual machines on a single physical computer using Virtual Server 2005. VSMT teaches how to take the data of multiple physical machines (i.e., running states) and migrate the data to virtual machines running on a single computer operating under Virtual Server 2005. (Shenoy Declaration, ¶ 38.) Both products were manufactured by Microsoft and were specifically designed to work with one another. In fact, VSMT even explicitly states that "[u]sing Virtual Server 2005 and VSMT, you can take advantage of server consolidation opportunities. . ." (VSMT, pg. 5.) Thus, it would have been obvious to one having ordinary skill in the art to combine VS Guide with VSMT as doing so is nothing more than combining known elements using known techniques to yield a predictable result (e.g., server consolidation). (Shenoy Declaration, ¶ 38.) Accordingly, as set forth in this section, VS Guide in view of VSMT teach each and every limitation of claims 1, 2, 6, 7, and 12-15.

1. ***Independent Claims 1 and 15 are Obvious Over VS Guide and VSMT***

a) **VS Guide and VSMT disclose a method for restoring access to running states of a plurality of primary computing systems onto a single computing system and a computer program product for performing the same (1[P] & [15[P])**

VS Guide describes having multiple virtual machines running on a single physical computing system: "[s]etting up multiple virtual machines allows you to configure different combinations of operating systems and applications on one physical computer." (VS Guide, pg. 17.) Each virtual machine restores access to a running state by executing the running state as a captured running state. (Shenoy Declaration, ¶ 40.)

VSMT describes capturing "running states" of a physical computer by converting physical computers into Virtual Server 2005 virtual machines: "[t]he VSMT is designed to help solve problems associated with migrating legacy operating systems and line of business applications to new hardware by creating a reliable and repeatable process for migrating an operating system and installed applications from a physical server to a virtual machine running on Virtual Server 2005." (VSMT, pg. 5.) The operating system and applications (i.e., running state) of the physical server are converted to run on a virtual machine (i.e., restored.) executing on a single host computer running Virtual Server 2005. (Shenoy Declaration, ¶ 40.)

Both VS Guide and VSMT describe computer software products created by Microsoft. Thus, the combination discloses the computer program product of claim 15[P] and the method of claim 1[P].

b) VS Guide and VSMT disclose accessing a captured running state of each of the primary computing systems on the single computing system (1[A] & 15[A])

VS Guide describes allowing users to access a particular virtual machine, thus accessing its associated data: "[s]tore the virtual machine folder in a location to which all of the users have access. By default, virtual machines are stored in the \Documents and Settings\All Users\Documents\Shared Virtual Machines folder, which is configured so that all users can access it." (VS Guide, pg. 66.) The virtual machine data is the "captured running state" since it was migrated from a physical machine into the virtual machine running on the single computing system (See FIG. 1 and description from claim elements 1[P] and 15[P] above). (Shenoy Declaration, ¶ 40.) VS Guide also describes how one can access a virtual machine remotely via a website: "[f]rom the Administration Website, you can also manage the state and configuration of virtual machines. In addition, you can access guest operating systems by using Remote Control view. . . ." (VS Guide, pg. 46.) Thus, VS Guide discloses "an act of accessing a captured running state of each of the primary computing systems on the single computing system."

c) VS Guide and VSMT disclose that the captured running states each include, or are altered to include, at least one device driver

configured to interface with a common virtualization component running on the single computing system (1[B] & 15[B])

VS Guide describes how the Virtual Server 2005 software includes a virtualization layer (i.e., common virtualization component) for interfacing with each virtual machine: "[s]tarting from the bottom of the logical stack, the host operating system manages the physical computer. The Virtual Machine Monitor (VMM) virtualization layer manages virtual machines, providing the software infrastructure for hardware emulation." (VS Guide, pg. 124.) The virtualization layer is a part of the Virtual Server 2005 software. (Shenoy Declaration, ¶ 42.)

VS Guide also describes the installation of Virtual Machine Additions, which include device drivers for a given virtual machine. For example, one of the many drivers provided is a video driver: "[v]irtual Machine Additions is not installed by default when you install Virtual Server, although you must install it in order to use a guest operating system effectively. . .The following features are included in Virtual Machine Additions. . .video drivers" (VS Guide, pg. 22.)

The virtualization layer acts as an interface between the virtual machines and the hardware of the host computer system. (Shenoy Declaration, ¶ 42.) Additionally, the installed drivers (e.g. "optimized video drivers") allow virtual machines to communicate with the computer's hardware via the virtualization layer. (Shenoy Declaration, ¶ 42.) Thus, VS Guide discloses "wherein the captured running states each include, or are altered to include, at least one device

driver that is configured to interface with a common virtualization component that runs on the single computing system."

> **d)** **VS Guide and VSMT disclose that the common virtualization component is configured to at least indirectly interface with hardware on the single computing system using a different interface than the interface used by the device driver (1[C] & 15[C])**

The configuration in this limitation is required for hardware virtualization in virtualization software such as Microsoft Virtual Server 2005. (Shenoy Declaration, ¶ 43.) For example, VS Guide describes how the virtualization layer (VMM) is responsible for emulating (i.e. simulating) various hardware components: "the guest operating system and applications run on the virtual machine as if they were running on physical hardware, rather than emulated hardware. When a guest operating system is running, the special-purpose VMM kernel manages the CPU and hardware during virtual machine operations. . . ." (VS Guide, pg. 124.) While the VMM interacts with the hardware of the computing system, the drivers included in the virtual machine interface with the VMM (Virtualization component.) (Shenoy Declaration, ¶¶ 43-44.) Thus, VS Guide discloses "wherein the common virtualization component is configured to at least indirectly interface with hardware on the single computing system that operates at least part of the hardware using a different interface than the at least one device driver is configured to interface with."

e) **VS Guide and VSMT disclose identifying a boot order for each of the primary computing systems (1[D] & 15[D])**

The VS Guide describes identifying a boot order for multiple virtual machines in the form of an option for staggering bootup of virtual machines: "[s]econds to delay automatically starting a virtual machine: The number of seconds to wait after the Virtual Server service is started before this virtual machine is started automatically when Virtual Server starts. This setting is useful when multiple virtual machines are configured to start automatically because you can stagger the startup." (VS Guide, pg. 58.)

In other words, a user may set different delays for different virtual machines to stagger their start up (e.g., a boot order.) By staggering delays, a user can manipulate the order in which virtual machines are booted (Shenoy Declaration, ¶ 45.) The order to start the virtual machines may be any order identified by the user, however, it would be obvious that the order would be identified based on how the primary computing systems are booted. (Shenoy Declaration, ¶ 46.) As described above, the virtual machines (i.e., captured running states) were initially migrated to the single computing system from the plurality of primary computing systems as described by VSMT. This migration of the computers' data to a single computer still allows one to "operate [] server applications in their native environment. . . ." (VSMT, pg. 5.) Thus, in order to maintain their native environment, one would typically operate the virtual machines in the same way

that their associated primary computing systems were operated, including what order they should be booted. (Shenoy Declaration, ¶ 46.) The '822 Patent even acknowledges well known booting dependencies between various services running on different computing systems existed at the time the '822 patent was filed. ('822 Patent, 6:50-53.) For at least these reasons, it would be obvious to one having ordinary skill the art to identify the boot order based on a boot order of the primary computing systems. Accordingly, VS Guide in view of VSMT discloses "an act of identifying a boot order for each of the primary computing systems."

     f)   **VS Guide And VSMT disclose booting the captured running states in the identified boot order on the single computing system (1[E] & 15[E])**

Each of the virtual machines automatically boots based on the selected order when the Virtual Server software is started: "[s]econds to delay automatically starting a virtual machine: The number of seconds to wait after the Virtual Server service is started before this virtual machine is started automatically when Virtual Server starts." (VS Guide, pg. 58.) Thus, VS Guide discloses "an act of booting the captured running states for each of the primary computing systems in the identified boot order on the single computing system."

     2.   *Dependent claim 2 is obvious over VS Guide in view of VSMT*

Claim 2 limits claim 1 by adding "an act of capturing the running states of each of the primary computing systems." VSMT discloses capturing and

- 17 -

migrating the state of a physical computer to a virtual machine: "[t]he migration process can be summarized in three phases: physical machine capture, virtual machine creation and deployment to virtual machine." (VSMT, pg. 12.) That is, the running state of the physical machine is captured and transferred to the virtual machine running on a single machine. Thus, the combination of VS Guide and VSMT renders claim 2 obvious.

3. *Dependent claim 6 is obvious over VS Guide in view of VSMT*

Claim 6 limits claim 1 by stating that "the at least one device driver are each the same on each of the captured running states." VS Guide discloses including at least one device driver in each of the captured running states as previously discussed for independent claim 1. The drivers allow the virtual machines to interact with the virtualization layer, and are installed from the same program— "Virtual Machine Additions." (VS Guide, pg. 22) As such, it would be obvious that the drivers installed for each virtual machine are the same since they are all installed from the same software package – Virtual Machine Additions. (Shenoy Declaration, ¶ 49.) For example, the virtualization layer creates a virtualized version of the host computer's sound card, and every virtual machine would include a same device driver provided by Virtual Machine Additions for the virtualized sound card. (Shenoy Declaration, ¶ 49.) Thus, the combination of VS Guide and VSMT renders claim 6 obvious.

4. *Dependent claim 7 is obvious over VS Guide in view of VSMT*

Claim 7 limits claim 1 by stating that "the act of identifying a boot order for each of the primary computing systems comprises: an act of the single computing system receiving an identification of the boot order from a source external to the single computing system." VS Guide discloses how users may control the management of the virtual machines via a website:

> "[t]he Administration Website is a browser-based tool for configuring and managing Virtual Server 2005 and its associated virtual machines. . .From the Administration Website you can perform most management tasks for Virtual Server except for moving, copying, and deleting resource and configuration files and changing their file system security settings. . .From the Administration Website, you can also manage the state and configuration of virtual machines. . . ." (VS Guide, pg. 46.)

Users may use the website to remotely change the configuration of the virtual machines, for example, to modify their boot order. (Shenoy Declaration, ¶ 50.) Further, the boot order is set by a user, which is also a source external to the single computing system. Thus, the combination of VS Guide and VSMT renders claim 7 obvious.

5. *Dependent claim 12 is obvious over VS Guide in view of VSMT*

Claim 12 limits claim 1 by stating that "each of the running states is an entire volume of the corresponding primary computing system." As described

previously, VSMT teaches the running states of the primary computing systems are captured using the virtual server migration toolkit. VSMT teaches taking an *image* of a computer system, and further teaches sending the image to the single computing system as a captured running state: "VSMT is used to create images of physical computers and deploy them in virtual machines running on Virtual Server 2005." (VSMT, pg. 10.) The image is taken of the physical computer's disk and written into a virtual disk: "[a]t this point the disk image is written from the ADS image into the VHD of the virtual machine." (VSMT, pg. 14.)

An image of a disk would include all of the disk's data, and thus the image would incorporate an entire volume of the computing system. (Shenoy Declaration, ¶ 51.) Thus, the combination of VS Guide and VSMT renders claim 12 obvious.

6. *Dependent claim 13 is obvious over VS Guide in view of VSMT*

Claim 13 limits claim 1 by adding "an act of the common virtualization component emulating the network between at least two of the plurality of primary computing systems when the capture running states of the at least two primary computing systems attempt to communicate over the network." VS Guide describes creating a virtual network between virtual machines running on a host computer: "[i]n this scenario, virtual machines can read, monitor, and capture the network traffic of other virtual machines attached to this virtual network. This is

the same situation that exists when physical computers are attached to the same network hub." (VS Guide, pg. 125.)
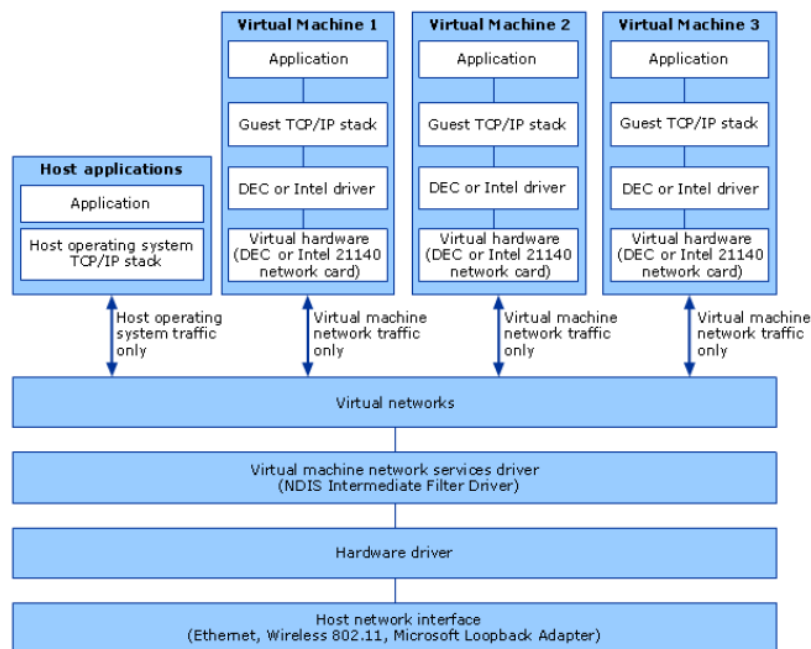
FIG. 2 – (VS Guide, pg. 125)

FIG. 2 above illustrates how multiple virtual machines communicate along a virtual network by emulating the capabilities of the host network card. (Shenoy Declaration, ¶¶ 52-53.) Thus, the combination of VS Guide and VSMT renders claim 13 obvious.

7.   ***Dependent claim 14 is unpatentable over VS Guide in view of VSMT***

Claim 14 limits claim 1 by stating that "the common virtualization component is configured to interface with hardware on the single computing system using an operating system of the single computing system." VS Guide describes how a virtual machine monitor runs on the host operating system: "[t]he

Virtual Machine Monitor (VIM) component of Virtual Server 2005 is a kernel-mode driver *running in the host operating system*. It has limited and specific functionality, acting as a firewall between the host operating system and virtual machines. It has access to the physical computer processor and manages resources between the two environments. . . ." (VS Guide, pg. 126; emphasis added.) The virtual machine monitor performs the same functions as the claimed common virtualization component. (Shenoy Declaration, ¶ 54.) Furthermore, as the virtual machine monitor runs on the operating system, it would access the hardware via the operating system. (Shenoy Declaration, ¶ 54.) FIG. 1 from the Shenoy Declaration (reproduced above) also illustrates how the Virtual Server 2005 software (which includes the virtual machine monitor) runs on top of the host operating system. Thus, the combination of VS Guide and VSMT renders claim 14 obvious.

F.   **Ground 2: Claims 1, 2, 6, 7, and 12-15 of the '822 Patent are Obvious over P2V Manual in view of ESX Guide.**

P2V Manual is a user manual for the VMware product "P2V Assistant" while ESX Guide is a user manual for the VMware product, ESX Server 2.1. (Shenoy Declaration, ¶ 55.) P2V Manual describes "migrat[ing] disk data from a physical (source) machine to a target virtual machine. . ." (P2V Manual, pg. 7.) Specifically, P2V Manual describes migrating to ESX virtual machines: "[P2V Assistant is a] useful mechanism for creating an ESX Server compatible, .dsk

formatted target disk from a physical source disk using the P2V Assistant." (P2V Manual, pg. 14.) ESX Guide discloses that "ESX Server contains many features to help you manage your virtual machines' resources" (ESX Guide, pg. 29.)

P2V Manual discloses the migration of data from physical machines to virtual machines running on machine single computing system within the VMware virtualization environment, and thus discloses substantially all of the limitations of claims 1, 2, 6, 7, and 12-15. But, to the extent that P2V Manual does not explicitly disclose booting the virtual machines in an identified order, these teachings can be found in the ESX Guide.

FIG. 3 from the Shenoy Declaration (reproduced below) illustrates how the various claimed elements of the '822 patent map to the elements disclosed by both P2V Manual and ESX Guide.
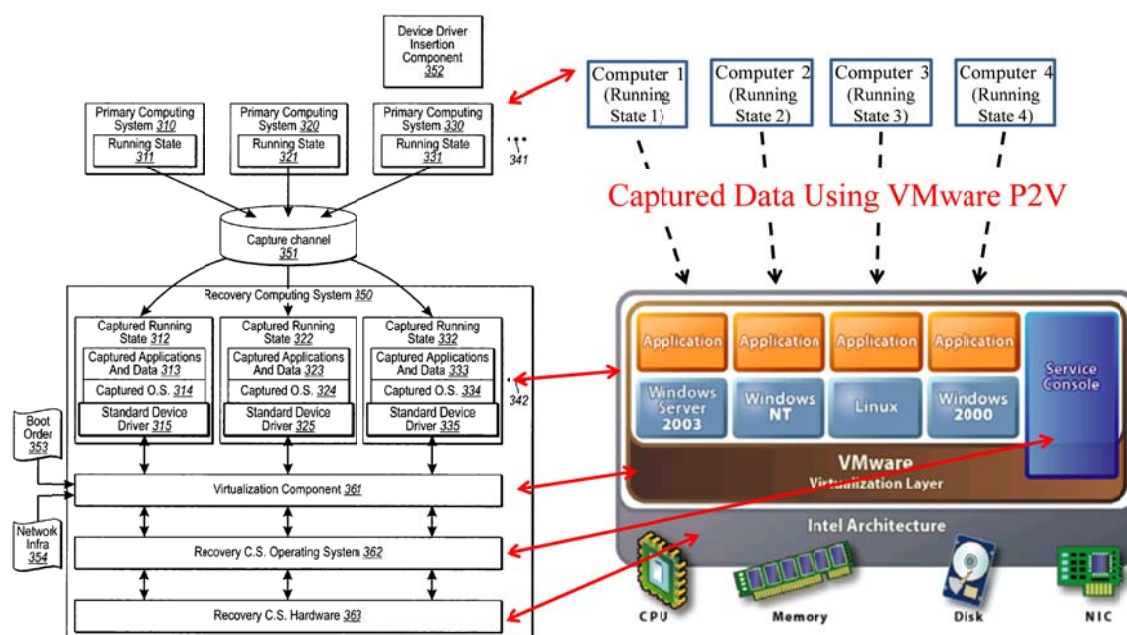

FIG. 3 – (Shenoy Declaration, pg. 28.)

The diagram on the left has been reproduced from Figure 3 of the '822 Patent while the diagram on the right has been reproduced from the figure on page 22 of the ESX Guide and annotated by Dr. Shenoy. Dr. Shenoy has annotated the figure from ESX Guide to illustrate that P2V Manual discloses capturing the running states of primary computing systems as illustrated in the top portion of the right-hand diagram in FIG. 3. (Shenoy Declaration, ¶ 57.) The claimed captured running states are the same as the applications and their respective operating systems, the claimed common virtualization component is the same as the VMware virtualization layer of ESX Guide, the operating system of the recovery computing system is the same as the service console of ESX Guide, and the hardware of the recovery computing system is the same as the Intel architecture layer with its hardware components shown graphically. (Shenoy Declaration, ¶ 58.)

It would have been obvious to one having ordinary skill in the art to combine the teachings of P2V Manual and ESX Guide. P2V Manual discloses how to clone data from a physical computer into a virtual disk running with ESX Server. (P2V Manual, pg. 83.) In fact, P2V Manual explicitly suggests that users "[r]efer to the *VMware ESX Server User's Manual* for complete and up to date instructions." (P2V Manual, pg. 83.) ESX Guide is the user manual for the most recent version of the ESX Server product at the time of P2V Manual. (Shenoy Declaration, ¶ 59.) One of ordinary skill in the art would have been motivated to

look to combine P2V Manual and ESX Guide to migrate data to virtual machines running under ESX Server, and subsequently configure those virtual machines using ESX Server. (Shenoy Declaration, ¶ 59.)

As set forth in this section, P2V Manual in view of ESX Guide teach each and every limitation of claims 1, 2, 6, 7, and 12-15.

> ### 1. *Independent Claims 1 and 15 are Obvious Over P2V Manual and ESX Guide*
>
> **a) P2V Manual and ESX Guide disclose a method for restoring access to running states of a plurality of primary computing systems onto a single computing system and a computer program product for preforming the same (1[P] & [15[P])**

P2V Manual discloses how VMware's P2V Assistant software can "migrate disk data from a physical (source) machine to a target virtual machine, without making any changes to the source machine. The method is similar to backing up a physical machine and restoring it on another machine with different hardware — virtual hardware." (P2V Manual, pg. 7.) Additionally, "P2V enables: . . . Migration of legacy servers." (P2V Manual, pg. 8.) The migrated disk data from the various servers are the running states of primary computing systems that are moved to virtual machines running (i.e. restored) on a single computer. (Shenoy Declaration, ¶ 61-62.)

Both P2V Manual and ESX Guide describe computer software products created by VMware. Thus, the combination discloses the computer program product of claim 15[P] and the method of claim 1[P].

b) **P2V Manual and ESX Guide disclose accessing a captured running state of each of the primary computing systems on the single computing system (1[A] & 15[A])**

P2V Manual discloses that the migrated data (i.e., captured running state) is stored on a virtual disk through a cloning process: "where the cloned disk is a virtual disk that is an exact copy of the source physical disk. This involves copying the data on a physical (source) machine's hard disk and transferring that data to a target virtual disk (the new cloned disk)." (P2V Manual, pg. 8.) The virtual disk is configured to be booted by a virtual machine: "[t]his adjustment is performed on the target virtual disk after the cloning and enables the target virtual disk to function as a bootable system disk in a virtual machine." (P2V Manual, pg. 9.) Further, the P2V manual discusses accessing the cloned virtual disk after it has been captured to reconfigure it to work in a virtual environment: "[t]his reconfiguration alters elements that are essential to working in a virtual environment." (P2V Manual, pg. 67.) (Shenoy Declaration, ¶ 64.) Indeed, "[v]irtual disks, in order to function with virtual machines, require some reconfiguring." (P2V Manual, pg. 67.) For example, the P2V Assistant can be used to "[p]re-install the VMware SVGA driver for smooth graphics and mouse" (P2V

Manual, pg. 74.)   The reconfiguration or altering of the virtual disk (which contains the captured running state) necessarily involves accessing the virtual disk. Thus, P2V Manual discloses "an act of accessing a captured running state of each of the primary computing systems on the single computing system."

        c)   **P2V Manual and ESX Guide disclose that the captured running states include, or are altered to include, at least one device driver configured to interface with a common virtualization component (1[B] & 15[B])**

As discussed previously, P2V Manual discloses configuring the virtual disks to be bootable within virtual machines. This configuration involves the use of drivers: "the cloned disk is not a bootable disk until a P2V Assistant performs a system reconfiguration. The system reconfiguration process configures selected drivers so that the operating system can boot in a virtual machine." (P2V Manual, pg. 12.)   More specifically, P2V Manual also discloses the installation of a graphics card driver as part of a software package called VMware tools: "[t]he P2V Assistant creates a temporary VMware SVGA graphics card driver during the reconfiguration. Whether you allowed P2V Assistant to install a temporary VMware SVGA driver or not, you *need* to install the SVGA driver during the installation of VMware Tools." (P2V Manual, pg. 89; emphasis added.)

ESX Guide describes including a software package called VMware Tools in a guest operating system (part of the captured running state on a virtual machine): "VMware Tools is a software package installed in the guest operating system that

gives you device drivers specific to VMware virtual devices where necessary, and it also includes several communication channels between the virtual machine and the ESX Server virtualization layer." (ESX Guide, pg. 33.)   The VMware virtualization layer of ESX Server is used to virtualize the computer's hardware resources for each of the running virtual machines, and is the claimed common virtualization component. (Shenoy Declaration, ¶ 65.)

Indeed, the '822 Patent explicitly states that VMware is an example of the claimed virtualization component: "the virtualization component 361 interfaces with the recovery computing system hardware 363 using the recovery computing system's operating system 362.   An example of the virtualization component is VMWARE®. . . ." ('822 Patent, 6:17-21.)   Thus, P2V Manual and ESX Guide disclose "wherein the captured running states each include, or are altered to include, at least one device driver that is configured to interface with a common virtualization component that runs on the single computing system."

d)   P2V Manual and ESX Guide disclose that the common virtualization component is configured to at least indirectly interface with hardware on the single computing system using a different interface than the interface used by the device driver (1[C] & 15[C])

The configuration in this limitation is required for hardware virtualization in virtualization software such as VMware ESX Server. (Shenoy Declaration, ¶ 67.) P2V Manual does not explicitly describe the virtualization environment used to run

the virtual machines, however, this is shown in the related ESX Guide. For example, ESX Guide teaches that "the VMware virtualization layer. . . provides the idealized hardware environment and virtualization of underlying physical resources." (ESX Guide, pg. 22.)

The ESX Server 2.1 product uses VMware architecture (See ESX Guide, pg. 22.)   Furthermore, the "virtualization of the underlying physical resources" provided by the VMware virtualization layer would require a different interface than that used between the virtualization layer and the virtual machine's device drivers. (Shenoy Declaration, ¶ 68.) Thus, P2V Manual and ESX Guide disclose "wherein the common virtualization component is configured to at least indirectly interface with hardware on the single computing system that operates at least part of the hardware using a different interface than the at least one device driver is configured to interface with."

    e)   **P2V Manual and ESX Guide disclose identifying a boot order for each of the primary computing systems (1[D] & 15[D])**

P2V Manual does not explicitly disclose the identification of a boot order of the primary computing systems. However, ESX Guide describes that users of ESX Server 2.1 can set the order in which the virtual machines boot: "[u]sing the system-wide Virtual Machine Startup and Shutdown option, you can:. . .Determine the global order in which virtual machines start and stop." (ESX Guide, pg. 257.)

FIG. 4 from ESX Guide is reproduced below and illustrates a screenshot for setting the boot order for virtual machines.
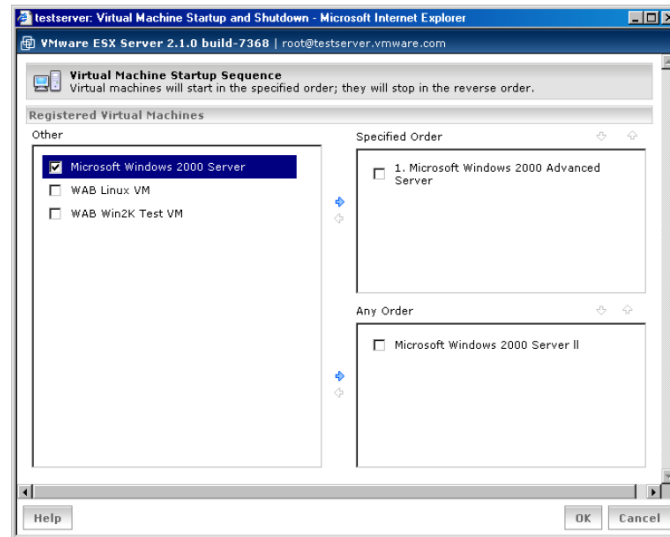
FIG. 4 – (ESX Guide, pg. 262)

The order to boot virtual machines may be any order identified by the user, however, it would be obvious that the order would be identified based on how the primary computing systems are booted. (Shenoy Declaration, ¶ 70.) Virtual machines (which execute the captured running states) are created by migrating the disk data from physical machines (i.e., primary computing systems.) P2V Manual also states that "[t]he conversion of the physical machine into a virtual machine includes all the identifying elements of the physical machine. That means the virtual machine has the same name, identification, and network connectivity of the physical machine." (P2V Manual, pg. 7.) The '822 Patent even acknowledges that well known booting dependencies between various services running on different

computing systems existed at the time the '822 Patent was filed. ('822 Patent, 6:50-53.)

The P2V manual explains that one feature of the P2V Assistant product is to "provide[] a smooth migration path for existing systems." (P2V Manual, pg. 8.) To ensure such a smooth migration from physical machines to virtual machines, it would be obvious to one having ordinary skill the art to identify the boot order of the virtual machines based on a boot order of the physical machines. (Shenoy Declaration, ¶ 70.) Thus, P2V Manual and ESX Guide disclose "an act of identifying a boot order for each of the primary computing systems."

f) **P2V Manual and ESX Guide disclose booting the captured running states in the identified boot order on the single computing system (1[E] & 15[E])**

P2V Manual does not explicitly disclose booting the captured running states in an identified order. However, ESX Guide teaches that each of the virtual machines will boot in the identified order when the ESX Server software is started: "[u]sing the system-wide Virtual Machine Startup and Shutdown option, you can: Configure your server to determine if virtual machines start up or shut down when the system starts or shuts down." (ESX Guide, pg. 257.) The virtual machines boot in the order specified by the user on the single computing system as described for the previous claim element. (Shenoy Declaration, ¶ 71.) Thus, ESX Guide

discloses "an act of booting the captured running states for each of the primary computing systems in the identified boot order on the single computing system."

2. *Dependent Claim 2 is obvious over P2V Manual and ESX Guide*

Claim 2 limits claim 1 by adding "an act of capturing the running states of each of the primary computing systems." P2V Manual discloses "[t]he process of creating a cloned disk, where the cloned disk is a virtual disk that is an exact copy of the source physical disk. This involves copying the data on a physical (source) machine's hard disk and transferring that data to a target virtual disk (the new cloned disk)." (P2V Manual, pg. 8.) The copying of the data from the physical disk to the virtual disk captures the running state from the physical computer. (Shenoy Declaration, ¶ 72.) It would be obvious that the running states of multiple physical computers could be captured by merely repeating this copying technique. (Shenoy Declaration, ¶ 73.) Thus, the combination of P2V Manual and ESX Guide renders claim 2 obvious.

3. *Dependent Claim 6 is obvious over P2V Manual and ESX Guide*

Claim 6 limits claim 1 by stating that "the at least one device driver are each the same on each of the captured running states." ESX Guide discloses including at least one device driver in each of the captured running states as previously discussed for independent claim 1. The drivers allow the virtual machines to interact with the emulated hardware created by the virtualization layer, and are

installed from the same software package—"VMware Tools." (ESX Guide, pg. 33)

It would be obvious that the drivers installed for each virtual machine are the same

since they are all installed from the same software package – VMware Tools.

(Shenoy Declaration, ¶ 74.)   For example, the virtualization layer creates a

virtualized version of the host computer's sound card, and every virtual machine

would include a same device driver provided by VMware Tools for the virtualized

sound card. (Shenoy Declaration, ¶ 74.)   Thus, the combination of P2V Manual

and ESX Guide renders claim 6 obvious.

4.   *Dependent Claim 7 is obvious over P2V Manual and ESX Guide*

Claim 7 limits claim 1 by stating that "the act of identifying a boot order for

each of the primary computing systems comprises: an act of the single computing

system receiving an identification of the boot order from a source external to the

single computing system." ESX Guide describes how users may use a web-based

management tool to configure and control the various virtual machines: "VMware

ESX Server provides the VMware Management Interface, a Web-based

management tool that allows you to. . .Control (power on, suspend, resume, reset

and power off) the virtual machines on the server." (ESX Guide, pg. 87.)

Additionally, FIG. 4 from ESX Guide (reproduced above) illustrates that the

startup order controls are accessed via Microsoft Internet Explorer (i.e., over the

internet.), and thus is accessible remotely. The identification of a boot order is

made by a user, which is also an external source from the single computing system. (Shenoy Declaration, ¶ 75.) Thus, the combination of P2V Manual and ESX Guide renders claim 7 obvious.

5. *Dependent Claim 12 is obvious over P2V Manual and ESX Guide*

Claim 12 limits claim 1 by stating that "each of the running states is an entire volume of the corresponding primary computing system." As discussed previously, P2V Manual discloses "[t]he process of creating a cloned disk, where the cloned disk is a virtual disk that is an exact copy of the source physical disk. This involves copying the data on a physical (source) machine's hard disk and transferring that data to a target virtual disk (the new cloned disk)." (P2V Manual, pg. 8.) When cloning a physical disk, the P2V Assistant software verifies the existence of a volume on the disk as illustrated in FIG. 5 below. Cloning a physical disk to a virtual disk would also capture the entire volume of the physical disk. (Shenoy Declaration, ¶ 76.) Thus, the combination of P2V Manual and ESX Guide renders claim 12 obvious.
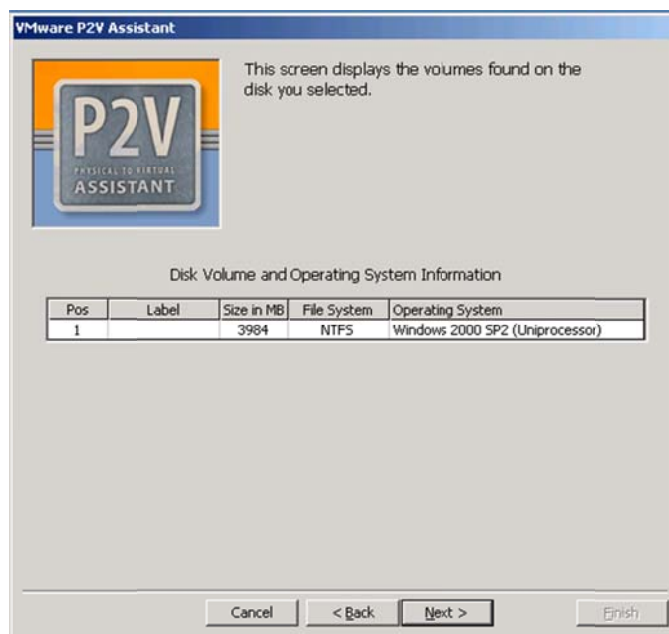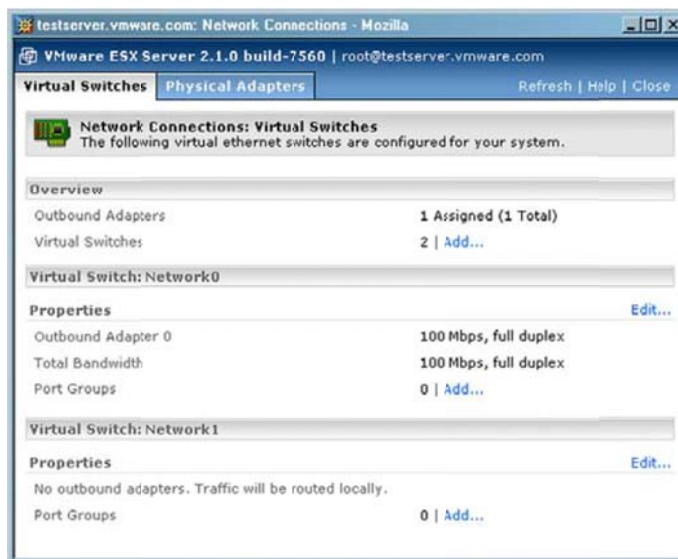
FIG. 5 (P2V Manual, pg. 52)

6. *Dependent Claim 13 is obvious over P2V Manual and ESX Guide*

Claim 13 limits claim 1 by adding "an act of the common virtualization component emulating the network between at least two of the plurality of primary computing systems when the capture running states of the at least two primary computing systems attempt to communicate over the network." ESX Guide describes how "ESX Server manages both the allocation of resources and the secure isolation of traffic meant for different virtual machines even when they are connected to the same physical network card." (ESX Guide, pg. 24.) Additionally, the network connections for the virtual machines may be configured as shown in FIG. 6 from ESX Guide on the next page. The "Virtual Switches" tab allows the user to add virtual switches, which allow the various virtual machines to communicate with one another across their own virtual network. (Shenoy

Declaration, ¶¶ 77 - 78.) Thus, the combination of ESX Guide and ESX Installation renders claim 13 obvious.

FIG. 6 – (ESX Guide, pg. 219)

7. *Dependent Claim 14 is obvious over P2V Manual and ESX Guide*

Claim 14 limits claim 1 by stating that "the common virtualization component is configured to interface with hardware on the single computing system using an operating system of the single computing system." ESX Guide discloses the use of a service console that "is also installed as a first component and is used to bootstrap the ESX Server installation and configuration, as well as to boot the system and initiate execution of the virtualization layer and resource manager." (ESX Guide, pg. 27.) "It is a customized version of Linux based on the Red Hat 7.2 distribution." (ESX Guide, pg. 194.) Linux is a well-known type of operating system. (Shenoy Declaration, ¶ 79.) The service console initiates the virtualization layer and interacts with hardware on the host computer: "[y]ou can

configure the server processor and disk resources for the service console." (ESX Guide, pg. 243.) (Shenoy Declaration, ¶ 80.) Thus, the combination of P2V Manual and ESX Guide renders claim 14 obvious.

G.   Ground 3: Claims 1, 2, 6, 7, and 12-15 of the '822 Patent are Obvious over Rob's Guide in view of ESX Guide.

Rob's Guide is a book written for users of VMware products providing details on how to use various features of VMware products beyond what is taught in the product manuals. (Rob's Guide, pg. 1.) Rob's Guide devotes an entire chapter to physical to virtual conversion (i.e., capturing running states) for consolidating the resources of the physical machines to virtual machines running on a single computer: "[t]he VMware GSX Server and ESX Server products can be used to consolidate servers to one general hardware platform." (Rob's Guide, pg. 8.) Rob's Guide also details the virtual architecture for running virtual machines concurrently on the same physical machine.

Rob's Guide discloses substantially all of the limitations of the challenged claims. But, to the extent that Rob's Guide does not explicitly disclose identifying a boot order and booting the captured running states (i.e., virtual machines) in the identified boot order, as discussed previously, ESX guide provides these missing teachings.

Rob's Guide explicitly teaches the use of VMware's ESX Server for consolidating multiple virtual machines on a single ESX server computing system:

"[t]he VMware GSX Server and ***ESX Server products*** can be used to consolidate servers to one general hardware platform." (Rob's Guide, pg. 8; emphasis added.) ESX Guide is a user guide for ESX Server 2.1. As both are directed to the same product—ESX Server—it would have been obvious to one having ordinary skill in the art to combine the teachings of Rob's Guide and ESX Guide. (Shenoy Declaration, ¶ 82.) A person of ordinary skill in the art would be motivated to use VMware's ESX server product to perform server consolidation based on the teachings found in Rob's Guide and would look to ESX Guide for further "information on how to use VMware ESX Server." (ESX Guide, pg. 21.)

Accordingly, as set forth in this section, Rob's Guide in view of ESX Guide teach each and every limitation of claims 1, 2, 6, 7, and 12-15.

1. Independent Claims 1 and 15 are Obvious Over Rob's Guide and ESX Guide

a) Rob's Guide and ESX Guide disclose a method for restoring access to running states of a plurality of primary computing systems onto a single computing system and a computer program product for performing the same (1[P] & [15[P])

Rob's guide teaches using physical to virtual conversion to consolidate multiple servers onto one platform: "[t]he VMware GSX Server and ESX Server products can be used to consolidate servers to one general hardware platform." (Rob's Guide, pg. 8.) "With the process of Physical to Virtual conversion, . . . you copy your machines into Virtual Machines on one box, and you have exactly the

same environment and test results as you would have in your production environment." (Rob's Guide, pg. 129.) Copying the data from physical machines to virtual machines running on a single computer restores access to the data (i.e., running state) of each of the physical machines. (Shenoy Declaration, ¶ 84.)

Further, because Rob's Guide discloses using VMware's ESX Server software to perform the physical to virtual conversion, the combination discloses the computer program product of claim 15[P] and the method of claim 1[P].

b)   Rob's Guide and ESX Guide disclose accessing a captured running state of each of the primary computing systems on the single computing system (1[A] & 15[A])

As discussed in the claim limitation above, a running state of a computer system is captured and consolidated on one machine. The destination virtual machine includes the captured running state as it includes the data copied over to the single computer from another computer (See description from previous claim elements 1[P] and 15[P] above). (Shenoy Declaration, ¶ 85.) Rob's Guide describes accessing this "captured running state" to make modifications: "[w]hatever method you choose to copy data and whatever operating system you transfer you will have to make modifications to your destination Virtual Machine. This is described for Windows and NetWare in section 18: 'Modify your restored operating system to work with VMware.'" (Rob's Guide, pg. 133.) Additionally, ESX Guide states "[i]f you need to view a particular virtual machine's desktop,

you can attach the VMware Remote Console and connect to the virtual machine."

(ESX Guide, pg. 94.) Thus, Rob's Guide in view of ESX Guide disclose "an act of

accessing a captured running state of each of the primary computing systems on

the single computing system."

c) Rob's Guide and ESX Guide disclose that the captured running states include, or are altered to include, at least one device driver configured to interface with a common virtualization component (1[B] & 15[B])

Rob's Guide discloses the loading of drivers into recently restored virtual

machines in order to use the VMware emulated hardware: "[o]nce you have

restored an image to a Virtual Machine, you will need to modify the operating

system to load the correct drivers for the hardware that is used in VMWare."

(Rob's Guide, pg. 151.) Additionally, Rob's Guide mentions the use of VMware

tools: "[a] few basic tips can help you to experience better performance in your

Virtual Machines. First of all make sure that you always install the VMware tools.

. . ." (Rob's Guide, pg. 33.) VMware tools are also described in ESX Guide:

"VMware Tools is a software package installed in the guest operating system that

gives you device drivers specific to VMware virtual devices where necessary, and

it also includes several communication channels between the virtual machine and

the ESX Server virtualization layer." (ESX Guide, pg. 33.) The '822 Patent

acknowledges that VMware is an example of the claimed virtualization component

(See '822 Patent, 6:17-21.) Thus, Rob's Guide in view of ESX Guide disclose

"wherein the captured running states each include, or are altered to include, at least one device driver that is configured to interface with a common virtualization component that runs on the single computing system."

> **d) Rob's Guide and ESX Guide disclose that the common virtualization component is configured to at least indirectly interface with hardware on the single computing system using a different interface than the interface used by the device driver (1[C] & 15[C])**

As explained above, Rob's Guide teaches the use of a common virtualization component. But, to the extent that Rob's Guide does not explicitly teach such a common virtualization component that is "configured to at least indirectly interface with hardware on the single computing system using a different interface than the interface used by the device driver," ESX Guide provides the missing teachings. For example, ESX Guide describes that "the VMware virtualization layer. . .provides the idealized hardware environment and virtualization of underlying physical resources." (ESX Guide, pg. 22.)

The ESX Server 2.1 product uses VMware architecture as described in ESX Guide (ESX Guide, pg. 22.) Furthermore, the "virtualization of the underlying physical resources" provided by the VMware virtualization layer would require a different interface than that used between the virtualization layer and the virtual machine's device drivers. (Shenoy Declaration, ¶ 88.) Thus, Rob's Guide in view of ESX Guide discloses "wherein the common virtualization component is

configured to at least indirectly interface with hardware on the single computing system that operates at least part of the hardware using a different interface than the at least one device driver is configured to interface with."

e) **Rob's Guide in view of ESX Guide Discloses Identifying a Boot Order for Each of the Primary Computing Systems (1[D] & 15[D])**

Rob's guide does not explicitly describe identifying a boot order. However, ESX Guide discloses "[u]sing the system-wide Virtual Machine Startup and Shutdown option, you can:. . .Determine the global order in which virtual machines start and stop." (ESX Guide, pg. 257.) FIG. 4 above also shows a screenshot of a user selecting a startup order for various virtual machines.

The order to start the virtual machines may be any order identified by the user, however, it would be obvious that the order would be identified based on how the primary computing systems are booted. (Shenoy Declaration, ¶ 90.) Recall that Rob's Guide described how converting physical machines to virtual machines running on one computer provides "exactly the same environment and test results as you would have in your production environment." (Rob's Guide, pg. 129.) In order to maintain the same environment, it would be obvious to boot the virtual machines in the same identified order as the physical machines (i.e., primary computing systems) (Shenoy Declaration, ¶ 90.) The '822 Patent even acknowledges that well known booting dependencies between various services

running on different computing systems existed at the time the '822 Patent was filed. ('822 Patent, 6:50-53.) Thus, Rob's Guide in view of ESX Guide discloses "an act of identifying a boot order for each of the primary computing systems."

f) **Rob's Guide and ESX Guide disclose booting the captured running states in the identified boot order on the single computing system (1[E] & 15[E])**

ESX Guide teaches that each of the virtual machines will boot when the ESX Server software is started: "[u]sing the system-wide Virtual Machine Startup and Shutdown option, you can: Configure your server to determine if virtual machines start up or shut down when the system starts or shuts down." (ESX Guide, pg. 257.) The virtual machines boot in the order specified by the user on the single computing system as described above for the previous claim element. (Shenoy Declaration, ¶ 91.) Thus, Rob's Guide in view of ESX Guide discloses "an act of booting the captured running states for each of the primary computing systems in the identified boot order on the single computing system."

2. *Dependent Claim 2 is obvious over Rob's Guide and ESX Guide*

Claim 2 limits claim 1 by adding "an act of capturing the running states of each of the primary computing systems." Rob's Guide describes two methods for capturing the running state of a computer and transferring it to a virtual machine. The first method is an imaging technique: "[t]o use this method you create an image file from your pc and store it somewhere. From that location you add the

file into virtual hard disk file and restore the image inside your Virtual Machine."
(Rob's Guide, pg. 131.) The second method describes copying data across a
network: "[w]ith this method a pc sends its data directly into a Virtual Machine via
TCP/IP where it is restored to the destination hard disk." (Rob's Guide, pg. 131.)
Thus, the combination of Rob's Guide and ESX Guide renders claim 2 obvious.

3. Dependent Claim 6 is obvious over Rob's Guide and ESX Guide

Claim 6 limits claim 1 by stating that "the at least one device driver are each
the same on each of the captured running states." Rob's Guide and ESX Guide
both disclose including at least one device driver in each of the captured running
states as previously discussed for independent claim 1. The drivers allow the
virtual machines to interact with specific hardware of the host computer via the
virtualization layer, and are installed from the same software package—"VMware
Tools." (ESX Guide, pg. 33.) As such, It would be obvious that the drivers
installed for each virtual machine are the same since they are all installed from the
same software package – VMware Tools. (Shenoy Declaration, ¶ 93.) For
example, the virtualization layer creates a virtualized version of the host
computer's sound card, and every virtual machine would include a same device
driver provided by VMware Tools for the virtualized sound card. (Shenoy
Declaration, ¶ 93.) Thus, the combination of Rob's Guide and ESX Guide renders
claim 6 obvious.

4. Dependent Claim 7 is obvious over Rob's Guide and ESX Guide

Claim 7 limits claim 1 by stating that "the act of identifying a boot order for each of the primary computing systems comprises: an act of the single computing system receiving an identification of the boot order from a source external to the single computing system." Rob's Guide does not explicitly disclose the boot order. However, ESX guide describes how users may use a web-based management tool to configure and control the various virtual machines: "VMware ESX Server provides the VMware Management Interface, a Web-based management tool that allows you to. . .Control (power on, suspend, resume, reset and power off) the virtual machines on the server." (ESX Guide, pg. 87.) Additionally, FIG. 4 above illustrates that the startup order controls are accessed via Microsoft Internet Explorer (i.e., over the internet.) and thus is accessible remotely. Further, the identification of a boot order is made by a user, which is also a source that is external from the single computing system. (Shenoy Declaration, ¶ 95.) Thus, the combination of Rob's Guide and ESX Guide renders claim 7 obvious.

5. Dependent Claim 12 is obvious over Rob's Guide and ESX Guide

Claim 12 limits claim 1 by stating that "each of the running states is an entire volume of the corresponding primary computing system." Rob's Guide describes how imaging may be used to capture a computer's data: "[w]ith imaging

software, you copy a hard disk to another hard disk (mostly in another PC) directly via the network or via an image file to create a copy of a pc." (Rob's Guide, pg. 130.) Rob's Guide also describes how the physical to virtual conversion may be done via an imaging technique: "create an image file from your pc and store it somewhere. From that location you add the file into virtual hard disk file and restore the image inside your Virtual Machine." (Rob's Guide, pg. 131.)

Creating an image of a computer's disk includes any and all data volumes present on that disk. (Shenoy Declaration, ¶ 96.) Rob's Guide describes copying an entire volume of data on a hard disk when it creates an image of the hard disk to transfer to another location. (Shenoy Declaration, ¶ 96.) Thus, the combination of Rob's Guide and ESX Guide renders claim 12 obvious.

6. Dependent Claim 13 is obvious over Rob's Guide and ESX Guide

Claim 13 limits claim 1 by adding "an act of the common virtualization component emulating the network between at least two of the plurality of primary computing systems when the capture running states of the at least two primary computing systems attempt to communicate over the network." Rob's Guide describes a VMware network adapter for networking between virtual machines: "a host-only adapter is a virtual adapter, that shows up in the list of adapters in Network Settings, but that does not have any relationship with any of the physical adapters in your computer." (Rob's Guide, pg. 87.) FIG. 7 from Rob's Guide

(reproduced below) illustrates a list of network adapters of a host computer that includes a VMware network Adapter.
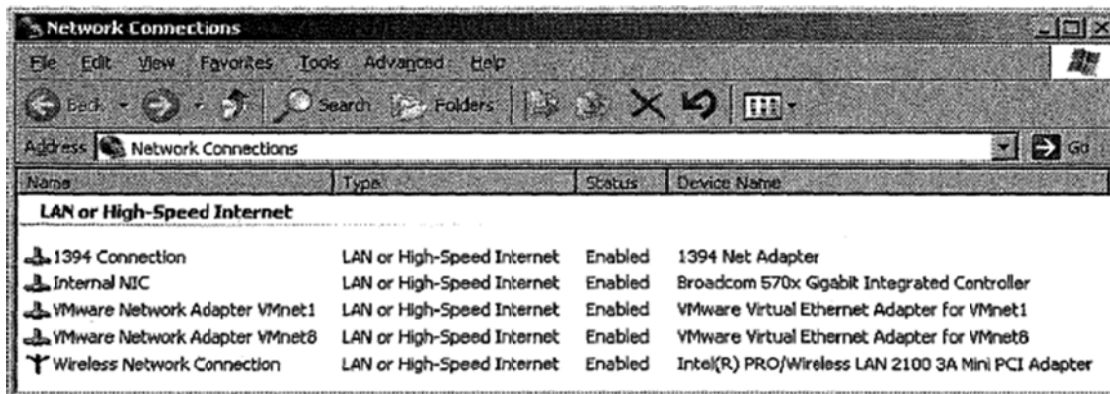

FIG. 7 – (Rob's Guide, pg. 87)

The VMware Network Adapter is provided by the VMware virtualization component for emulating a network connection between the various virtual machines on the host system. (Shenoy Declaration, ¶ 97.) Additionally, ESX Guide describes the use of virtual switches as previously discussed. The virtual switches allow the various virtual machines to communicate with one another across their own virtual network set up by ESX Server's virtualization layer. (Shenoy Declaration, ¶ 98.) Thus, the combination of Rob's Guide and ESX Guide renders claim 13 obvious.

7.   *Dependent Claim 14 is obvious over Rob's Guide and ESX Guide*

Claim 14 limits claim 1 by stating that "the common virtualization component is configured to interface with hardware on the single computing system using an operating system of the single computing system." Rob's Guide

as a whole discloses how to use various VMware software products and explicitly states the use of ESX Server: "[t]he VMware GSX Server and ESX Server products can be used to consolidate servers to one general hardware platform." (Rob's Guide, pg. 8.) The ESX Server software uses a Linux-based operating system (called a service console) that acts as an interface between the virtualization layer and the hardware as disclosed in the ESX Guide. (Shenoy Declaration, ¶ 99.) The ESX Guide discloses that the service console "is also installed as a first component and is used to bootstrap the ESX Server installation and configuration, as well as to boot the system and initiate execution of the virtualization layer and resource manager." (ESX Guide, pg. 27.)

In order to configure the virtualization layer, the service console also interfaces with the hardware of the computer to correctly configure the resources that will be virtualized by the virtualization layer: "[y]ou can configure the server processor and disk resources for the service console." (ESX Guide, pg. 243.) (Shenoy Declaration, ¶ 100.) Thus, the combination of Rob's Guide and ESX Guide renders claim 14 obvious.

H. Ground 4: Claims 1, 2, 6, 7, and 12-15 of the '822 Patent are Obvious over Le in view of ESX Guide.

Le is a patent assigned to VMware, Inc. that covers numerous aspects of virtualization. Le is primarily directed to imaging the contents of a computer's disk for migration to a virtual disk (Le, 86:22-47.) Le also describes the use of

physical to virtual conversion for consolidating multiple computers onto a single computing system. (Le, 14:49-15:3.) And Le details the virtualization architecture to allow multiple virtual machines to run simultaneously on a single computer (Le at Figures 1-2).

Le discloses substantially all of the limitations of the challenged claims. But, to the extent it can be argued that Le does not explicitly disclose identifying a boot order and booting the captured running states (i.e., virtual machines) in the identified boot order, ESX Guide provides such teachings as discussed previously.

A person having ordinary skill in the art would have been motivated to combine Le and ESX Guide. First, Le is a patent owned by VMware that describes VMware virtualization features and mechanisms. (Shenoy Declaration, ¶ 101.) ESX Guide is the manual for a VMware product—ESX Server 2.1, which is capable of performing the methods discussed in Le. (Shenoy Declaration, ¶ 102.) In fact, Le explicitly discusses the use of ESX Server: "[c]ompared with the hosted deployment, a kernel may offer greater performance because it can be co-developed with the VMM and be optimized for the characteristics of a workload consisting of VMMs. The ESX Server product of VMware, Inc., has such a configuration." (Le, 13:22-26.) As such, it would have been obvious to one having ordinary skill in the art to combine the teachings of Le and ESX Guide, since both relate to VMware products and systems, and one would be motivated to seek the

administrative guide to ESX Server (ESX Guide) in order to learn about the configuration options for the virtual machines that have been migrated as described in Le. (Shenoy Declaration, ¶ 102.)

1. *Independent Claims 1 and 15 are Obvious Over Le and ESX Guide*

a) **Le and ESX Guide disclose a method for restoring access to running states of a plurality of primary computing systems onto a single computing system and a computer program product for performing the same (1[P] & [15[P])**

Le describes that "a common need is to convert physical computers to virtual machines." (Le, 44:60-61.) Migrating the data from a physical computer to a virtual machine restores access to the running state of the physical computer on the computer that hosts the corresponding virtual machine. (Shenoy Declaration, ¶ 104.) Indeed, Le further explains that "[a] physical-to-virtual (P2V) conversion of a physical computer is generally defined as migrating the physical computer's disk state, which includes its operating system, software applications and data, into a virtual disk. . . ." (Le, 44:64-67.) Le also describes that multiple computers may be converted in this way: "[s]erver consolidation is a context in which this conversion capability is particularly desirable. The idea behind server consolidation is to take multiple server computers and run them as virtual machines on a single physical computer." (Le, 14:54-57.)

Le mentions using VMware's ESX Server software product. As such, the combination discloses the computer program product of claim 15[P] and the method of claim 1[P].

b)   Le and ESX Guide disclose accessing a captured running state of each of the primary computing systems on the single computing system (1[A] & 15[A])

Le teaches booting the operating system of a captured running state: "[t]he resulting virtual disk can be attached to a virtual machine, which can then boot the migrated operating system and run the migrated software, resulting in a system functionally equivalent to the physical computer." (Le, 45:1-5.) The migrated operating system and software associated with the virtual machine is analogous to the captured running state. (Shenoy Declaration, ¶ 105.) Le further discloses reconfiguring the migrated operating system, which necessarily involves accessing the captured running state (Le, 45:7-9.) (Shenoy Declaration, ¶ 106.) As such, Le discloses "an act of accessing a captured running state of each of the primary computing systems on the single computing system."

c)   Le and ESX Guide disclose that the captured running states include, or are altered to include, at least one device driver configured to interface with a common virtualization component (1[B] & 15[B])

Le describes how drivers are included in the running virtual machines to allow them to interface with virtual devices: "[t]he guest system software includes a guest operating system 520 (which may simply be a copy of a conventional

operating system), and drivers 522 as needed for the various virtual devices 540. . .

." (Le, 12:13-16.) The virtual devices are provided by a virtualization component, referred to as a virtual machine monitor (VMM): "[a] VMM is usually a thin piece of software that runs directly on top of a host, or directly on the hardware, and virtualizes resources of the physical host machine." (Le, 12:39-41.) FIG. 8 below illustrates an annotated version of Figure 1 from Le. The captured running state includes drivers 522. These drivers interface with the emulated devices of the VMM 600. (Shenoy Declaration, ¶ 108.)
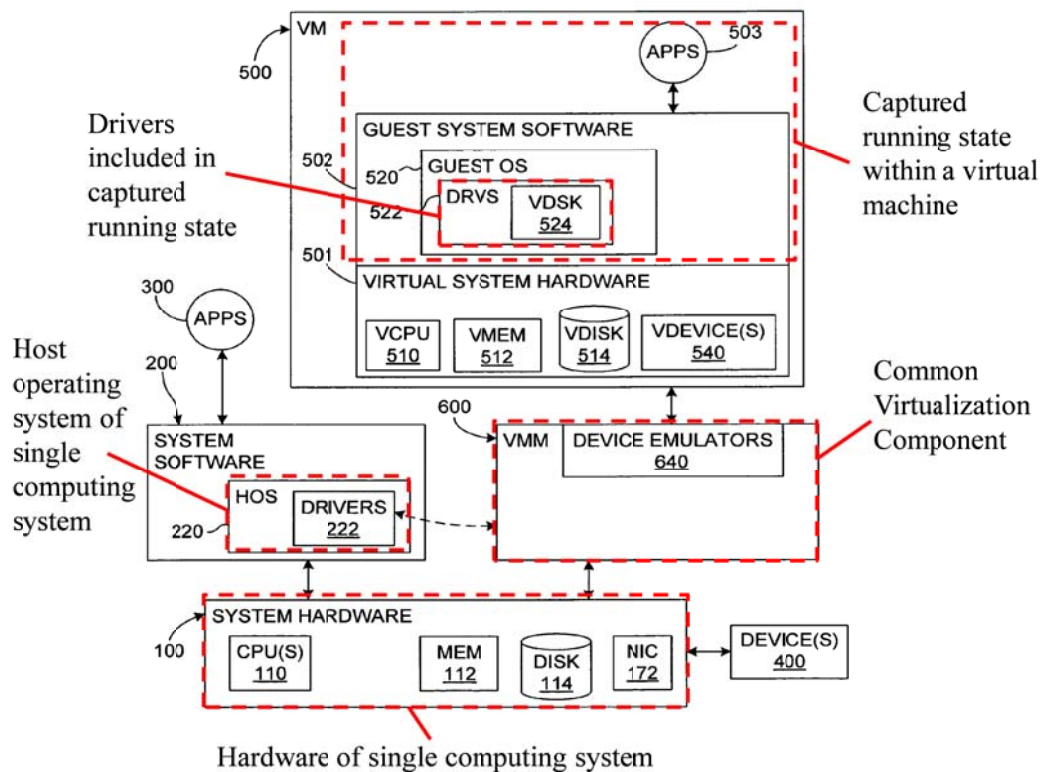
FIG. 8 – (Shenoy Declaration, pg. 59.)

As such, Le discloses "wherein the captured running states each include, or are altered to include, at least one device driver that is configured to interface with a common virtualization component that runs on the single computing system."

d) Le and ESX Guide disclose that the common virtualization component is configured to at least indirectly interface with hardware on the single computing system using a different interface than the interface used by the device driver (1[C] & 15[C])

As just discussed, Le teaches the use of drivers within a virtual machine to interact with a common virtualization component (i.e. the VMM.) Le also describes that the VMM interacts with hardware on the host computer: "the VMM 600 is co-resident at system level with the host operating system 220 such that both the VMM and the host OS can independently modify the state of the host processor." (Le, 12:65-13:1.) The VMM also interacts with the hardware through the host computer's operating system: "the VMM calls into the host OS (symbolized by the dashed, double-ended arrow) via a special one of the drivers 222 and a dedicated one of the user-level applications 300 to have the host OS perform certain I/O operations of behalf of the VM." (Le, 13:1-5.) In either situation where the VMM interfaces with the host computer's hardware, the interface used is different than the driver interface between the virtual machine and the VMM. (Shenoy Declaration, ¶¶ 109-111.) This can also be seen visually in FIG. 8 from above where the double ended arrow between VM 500 and VMM 600

represents the device driver interface while the double ended arrow between VMM 600 and system hardware 100 represents a different interface with the hardware. (Shenoy Declaration, ¶ 110.) As such, Le discloses "wherein the common virtualization component is configured to at least indirectly interface with hardware on the single computing system that operates at least part of the hardware using a different interface than the at least one device driver is configured to interface with."

e) **Le and ESX Guide disclose identifying a boot order for each of the primary computing systems (1[D] & 15[D])**

Le does not explicitly identify a boot order. However, ESX Guide describes that "[u]sing the system-wide Virtual Machine Startup and Shutdown option, you can. . .Determine the global order in which virtual machines start and stop." (ESX Guide, pg. 257.) FIG. 4 from above also shows a screenshot of a user selecting a startup order for various virtual machines.

The order to start the virtual machines may be any order identified by the user, however, it would be obvious that the order would be identified based on how the primary computing systems are booted. (Shenoy Declaration, ¶ 113.) The virtual machines' data (i.e., captured running states) can be migrated from the applications and operating systems of physical servers (i.e., primary computing systems.) In order to maintain their native environment, one would typically operate the virtual machines in the same way that their associated primary

computing systems were operated, including what order they should be booted. (Shenoy Declaration, ¶ 113.) The '822 Patent even acknowledges that well known booting dependencies between various services running on different computing systems existed at the time the '822 Patent was filed. ('822 Patent, 6:50-53.) For at least these reasons, it would be obvious to one having ordinary skill the art to identify the boot order based on a boot order of the primary computing systems. Thus, Le in view of ESX Guide disclose "an act of identifying a boot order for each of the primary computing systems."

f) **Le and ESX Guide disclose booting the captured running states in the identified boot order on the single computing system (1[E] & 15[E])**

ESX Guide teaches that each of the virtual machines will boot when the ESX Server software is started: "[u]sing the system-wide Virtual Machine Startup and Shutdown option, you can: Configure your server to determine if virtual machines start up or shut down when the system starts or shuts down." (ESX Guide, pg. 257.) The virtual machines boot in the order specified by the user on the single computing system as described above for the previous claim element. (Shenoy Declaration, ¶ 114.) Thus, Le in view of ESX Guide discloses "an act of booting the captured running states for each of the primary computing systems in the identified boot order on the single computing system."

2. Dependent Claim 2 is obvious over Le and ESX Guide

Claim 2 limits claim 1 by adding "an act of capturing the running states of each of the primary computing systems." Le describes an act of "migrating the physical computer's disk state, which includes its operating system, software applications and data, into a virtual disk, with minimal user intervention." (Le, 44:65-45:1.) Le further states that "[t]he migration comprises two steps: 1) the disk imaging step copies disk data from the physical disk to the virtual disk; and 2) the reconfiguration step reconfigures the migrated operating system to make it compatible with the virtual machine's hardware configuration." (Le, 45:5-9.) Imaging a computer to migrate to a virtual machine captures the running state of the computer. (Shenoy Declaration, ¶ 115.) Thus, the combination of Le and ESX Guide renders claim 2 obvious.

3. Dependent Claim 6 is obvious over Le and ESX Guide

Claim 6 limits claim 1 by stating that "the at least one device driver are each the same on each of the captured running states." Le discloses including at least one device driver in each of the captured running states as previously discussed for independent claim 1. Furthermore, ESX Guide discloses the installation of drivers from the same software package – "VMware Tools." (ESX Guide, pg. 33.) As such, It would be obvious that the drivers installed for each virtual machine are the same since they are all installed from the same software package – VMware Tools.

(Shenoy Declaration, ¶ 116.) For example, the virtualization component creates a virtualized version of the host computer's sound card, and every virtual machine would include same a device driver provided by VMware Tools for the virtualized sound card. (Shenoy Declaration, ¶ 116.) Thus, the combination of Le and ESX Guide renders claim 6 obvious.

4. *Dependent Claim 7 is obvious over Le and ESX Guide*

Claim 7 limits claim 1 by stating that "the act of identifying a boot order for each of the primary computing systems comprises: an act of the single computing system receiving an identification of the boot order from a source external to the single computing system." Le does not explicitly disclose the boot order, however, ESX guide describes how users may use a web-based management tool to configure and control the various virtual machines: "VMware ESX Server provides the VMware Management Interface, a Web-based management tool that allows you to. . .Control (power on, suspend, resume, reset and power off) the virtual machines on the server." (ESX Guide, pg. 87.) Additionally, FIG. 4 above illustrates that the startup order controls are accessed via Microsoft Internet Explorer (i.e., over the internet) and thus is accessible remotely. Further, the identification of a boot order is made by a user, which is also a source that is external from the single computing system. (Shenoy Declaration, ¶¶ 117-118.) Thus, the combination of Le and ESX Guide renders claim 7 obvious.

5. *Dependent Claim 12 is unpatentable over Le in view of ESX Guide*

Claim 12 limits claim 1 by stating that "each of the running states is an entire volume of the corresponding primary computing system." Le describes how imaging may be used to capture a computer's data: "[t]he migration comprises two steps: 1) the disk imaging step copies disk data from the physical disk to the virtual disk. . . ." (Le, 45:5-7.)

Creating an image of a computer's disk would include any and all data volumes present on the disk. (Shenoy Declaration, ¶ 119.) Le describes copying an entire volume of data on a hard disk when it creates an image of the hard disk to transfer to another location. (Shenoy Declaration, ¶ 119.) Thus, the combination of Le and ESX Guide renders claim 12 obvious.

6. *Dependent Claim 13 is obvious over Le and ESX Guide*

Claim 13 limits claim 1 by adding "an act of the common virtualization component emulating the network between at least two of the plurality of primary computing systems when the capture running states of the at least two primary computing systems attempt to communicate over the network." Le teaches that virtual machines have "the ability to simulate a network of multiple machines using a single physical computer, using a VMM's virtual network capabilities." (Le, 15:28-30.) Additionally, ESX Guide describes that "ESX Server manages both the allocation of resources and the secure isolation of traffic meant for

different virtual machines even when they are connected to the same physical network card." (ESX Guide, pg. 24.) The "Virtual Switches" tab seen in FIG. 5 taken from ESX Guide allows the user to add virtual switches. The virtual switches allow the various virtual machines to communicate with one another across their own virtual network set up by ESX Server's virtualization layer. (Shenoy Declaration, ¶¶ 120-121.) Thus, the combination of Le and ESX Guide renders claim 13 obvious.

7. Dependent Claim 14 is obvious over Le and ESX Guide

Claim 14 limits claim 1 by stating that "the common virtualization component is configured to interface with hardware on the single computing system using an operating system of the single computing system." Le describes how the VMM may interface with the host operating system using special drivers: "the VMM calls into the host OS (symbolized by the dashed, double-ended arrow) via a special one of the drivers 222 and a dedicated one of the user-level applications 300 to have the host OS perform certain I/O operations of behalf of the VM." (Le, 13:1-5.) FIG. 8 reproduced above also illustrates the dashed, doubled-ended arrow between the operating system drivers and the VMM. These drivers 222 are used to interface with the computer's hardware as indicated by the double ended arrow between system software 200 and the hardware 100 in FIG. 7.

(Shenoy Declaration, ¶ 122.) Thus, the combination of Le and ESX Guide renders claim 14 obvious.
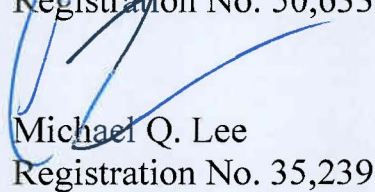
IV. Conclusion

For the grounds specified above, *inter partes* review of claims 1, 2, 6, 7, and 12-15 of U.S. Patent No. 7,480,822 is respectfully requested.

Respectfully submitted,
STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.

Lori A. Gordon
Registration No. 50,633

Michael Q. Lee
Registration No. 35,239

Attorneys for Petitioner

Date: October 22, 2013

1100 New York Avenue, N.W.
Washington, D.C. 20005-3934
(202) 371-2600

V.   APPENDIX A

Exhibit List

VSC 1001	U.S. Patent No. 7,480,822
VSC 1002	Microsoft Virtual Server Administrator Guide
VSC 1003	Microsoft Virtual Server Migration Toolkit Whitepaper
VSC 1004	VMware P2V Assistant User's Manual
VSC 1005	ESX Server 2.1 Administration Guide
VSC 1006	Rob's Guide to Using VMware
VSC 1007	U.S. Patent No. 8,209,680 to Le et al.
VSC 1008	Declaration of Dr. Prashant Shenoy
VSC 1009	*Curriculum Vitae* of Dr. Prashant Shenoy

CERTIFICATION OF SERVICE (37 C.F.R. §§ 42.6(e), 42.105(a))

The undersigned hereby certifies that the above-captioned "Petition for *Inter Partes* Review of U.S. Patent No. 7,480,822," including the accompanying Exhibit List and Exhibits (VSC 1001 – 1009), was served in its entirety on October 22, 2013, upon the following parties via FedEx®:

ALG INTELLECTUAL PROPERTY, LLC
922 W. Baxter Dr.
Suite 100
South Jordan, UT 84095
Patent owner's correspondence address of
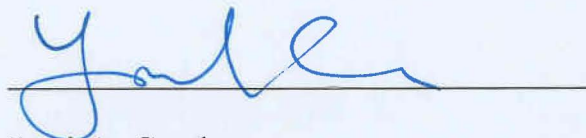*record for U.S. Patent No. 7,480,822*

QUINN EMANUEL URQUHART &
SULLIVAN, LLP

Jennifer A. Kash
50 California Street, 22nd Floor
San Francisco, CA 94111

David A. Nelson, Esq.
500 West Madison Street, Suite 2450
Chicago, IL 60661

Kate Cassidy
51 Madison Avenue, 22nd Floor
New York, NY 10010
Additional addresses known to Petitioner
as likely to affect service

STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.

Date: October 22, 2013

1100 New York Avenue, N.W.
Washington, D.C.20005-3934
(202) 371-2600

Lori A. Gordon
Attorney for Petitioner
Registration No. 50,633