

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

CIRBA INC. (d/b/a DENSIFY)
and CIRBA IP, INC.,

Plaintiffs,

v.

VMWARE, INC.,

Defendant.

Civil Action No. 1:19-cv-00742-LPS

JURY TRIAL DEMANDED

FIRST AMENDED COMPLAINT

Plaintiffs Cirba, Inc. (d/b/a Densify) and Cirba IP, Inc. (collectively, “Densify”) file this First Amended Complaint against Defendant VMware, Inc. (“VMware”), and respectfully allege as follows:

NATURE OF THE ACTION

1. Densify is a quintessential start-up success story. Through innovation and years of hard work, Densify has earned recognition for having set the standard in the industry for cloud and virtual infrastructure optimization. Global 5000 organizations use Densify’s software to reduce costs, operate with less infrastructure, and achieve better application performance.

2. Densify is an industry leader. For the past decade, Densify has won praise as a “Best Cloud Management Solution of the Year,” “EMA Top 3,” “Leader in Cloud Cost Monitoring,” “Best of VMworld 2017 Gold Winner,” “Top 10 Cloud Solution Provider,” “Editors’ Choice,” “Hot Product,” “Vendor to Watch,” “Companies to Watch,” “Top 10 Virtualization Vendors to Watch,” “Cool Vendor,” among other recognitions.

3. For years, Densify's products have led the industry in optimization, competing effectively based on its innovations and foundational patent protection. Its products have saved customers millions of dollars in hardware and software licensing costs, increasing efficiency and decreasing risk.

4. Densify's technology offerings were so compelling that VMware took Densify's ideas and intellectual property, brazenly infringing the patents with products that worked the same way, looked strikingly similar, and even were advertised using the "Densify" trademarks.

5. Rather than innovating on its own, VMware has systematically copied Densify's technology. VMware has used Densify's technology as a blueprint to close in on Densify's lead as one of the best optimization solutions in the market. VMware's infringement has gradually increased over time, copying more and more of the features of Densify's product—the most significant to date being the recent releases of its flagship product, vROps, and its movement into the hybrid cloud space. With these recent developments, VMware is imminently threatening Densify's virtual infrastructure optimization IP, and VMware recently has indicated it intends to leverage its infringing technology to move into Densify's cloud infrastructure optimization business.

6. Densify has no recourse but to file this action to stop VMware's misuse of its intellectual property. Densify has invested years and millions of dollars to develop and commercialize products embodying its intellectual property. VMware is a multi-billion dollar global player that dominates the virtual infrastructure market; 99% of *Fortune 1000* companies reportedly are VMware customers. If VMware's unauthorized use of Densify's intellectual property is not stopped, VMware can use its market power and dominant position to outspend Densify and swamp Densify's marketing and sales of products embodying Densify's patented

technology. Densify cannot fairly compete against a behemoth like VMware unless its intellectual property is respected.

7. Densify, which has not licensed its patents to competitors, has chosen to compete in the marketplace and is entitled to do so based on the exclusivity afforded by its patents. The disclosure of innovation in patents is not intended to facilitate unauthorized use, but rather to incentivize public disclosure for the benefit of all, in return for the promise to inventors of exclusive rights for a limited period of time. Meaningful consequences are needed to protect Densify's exclusive rights.

8. In the end, this case is about ensuring a level playing field so smaller competitors like Densify can compete fairly based on their hard work and protected innovations against larger companies like VMware.

9. Accordingly, Densify brings this action under the patent laws, 35 U.S.C. § 1 *et seq.*, the Trademark Act of 1946 (the Lanham Act), 15 U.S.C. § 1051 *et seq.*, Delaware law, and common law. Densify seeks to stop VMware's infringement of its patents and other intellectual property so that it may compete on the strength of the breakthrough products it worked so hard to create.

THE PARTIES

10. Cirba, Inc. (d/b/a Densify) is a Canadian privately-owned corporation headquartered in Markham, Ontario. Densify's principal place of business is located at 400 – 179 Enterprise Boulevard Markham, Ontario L6G 0E7 Canada. In addition to Markham, Densify has operations in New York, London, and Melbourne. Cirba, Inc. is the exclusive licensee of the patents asserted in this case.

11. Cirba IP, Inc. is Canadian privately-owned corporation headquartered in Markham, Ontario. Cirba IP is located at 400 – 179 Enterprise Boulevard Markham, Ontario L6G 0E7

Canada. Cirba IP, Inc. is a wholly owned subsidiary of Cirba, Inc. Cirba IP, Inc. owns all right, title, and interest in and to the patents asserted in this case, and has exclusively licensed them to Cirba, Inc.

12. VMware, Inc. is a publicly traded corporation organized and existing under the laws of the State of Delaware. VMware's headquarters are located at 3401 Hillview Ave, Palo Alto, California. VMware also has offices in Colorado, Georgia, Massachusetts, Texas, Virginia, and Washington, among other offices located in twenty countries around the world.

JURISDICTION AND VENUE

13. This is a civil action asserting claims of patent infringement of U.S. Patent Nos. 8,209,687 ("the '687 patent") and 9,654,367 ("the '367 patent") (collectively, the "Asserted Patents"), unfair competition under the Lanham Act, 15 U.S.C. § 1125(a), deceptive trade practices under Delaware Code Title 6 § 2532, and common law trademark infringement.

14. This Court has subject matter jurisdiction over this case under 28 U.S.C. §§ 1331 and 1338, 15 U.S.C. § 1121, and 28 U.S.C. § 1367(a).

15. Venue is proper in this Court under 28 U.S.C. §§ 1391 and 1400(b). VMware has offered and sold and continues to offer and sell its infringing products and services in this district. On information and belief, VMware sells and offers to sell the infringing products and services to developers, partners, or customers in this district, such as the University of Delaware, ADP, Alliant, Rent-a-Center, Cardinal Health, and the Make-A-Wish Foundation of America. VMware has committed acts of patent infringement in this district, is incorporated in this district, and does business in this district.

16. This Court has personal jurisdiction over VMware. VMware is incorporated in Delaware and has purposely availed itself of the privileges and benefits of the laws of the State of Delaware. VMware has continuous and systematic business contacts with the State of Delaware.

VMware, directly and/or through subsidiaries or intermediaries, conducts its business extensively throughout Delaware, by shipping, distributing, offering for sale, selling, and advertising its products and services in the State of Delaware and in this district. VMware, directly and/or through subsidiaries or intermediaries, has purposefully and voluntarily placed its infringing products and services into this district and into the stream of commerce with the intention and expectation that they will be purchased and used by consumers in this district.

BACKGROUND

A. Densify And Its Technology

17. Densify was founded in 1999. It is a software company that provides machine learning analytics products related to on-premise information technology (“IT”) infrastructure and the public cloud. It is a company built on innovation of products that save its customers millions of dollars in software licensing and computing infrastructure costs by making their utilization of servers more efficient. Densify has grown to over 180 employees and counts among its customers many of the world’s most prominent and sophisticated companies. Densify has been recognized for its innovations; it has won numerous industry awards and has been granted patents on its technology by the U.S. Patent and Trademark Office (“PTO”). These accomplishments have attracted significant investment in Densify, which is used to innovate new products that make its customers’ businesses run more efficiently, reliably, and profitably.

18. Many companies deploy their own IT infrastructure, including computing, storage and networking equipment, on premises. Generally, computing environments are designed to run at least one workload that performs business functions and consumes compute resources, e.g., resources related to central processing units (“CPUs”), memory, disk, network, and other hardware. The workloads run on computing systems such as servers that supply the computing resources. Each computing system has a finite capacity of resources.

19. Depending on the size of the company, the amount of equipment and, importantly, the number of physical servers that companies deploy can be substantial, which is necessary in order to keep pace with today's high processing demands. There are inefficiencies associated with these large infrastructure environments, requiring IT organizations to manage large numbers of physical servers, each operating only at a fraction of their capacity.

20. Enter virtualization. In computing, "virtualization" refers to the act of creating a "virtual" version of something. Virtualization relies on software to simulate hardware functionality and create a virtual computer system, which allows companies to run more than one virtual computer system, or virtual machine ("VM"), on a single physical server. In other words, virtualization offers greater efficiency and economies of scale.

21. In computing, a virtual machine or VM is essentially an emulation of a computer system. It functions as and represents a real computer machine but does so only logically as defined by software (virtualization software is referred to as a "hypervisor"). VMs originated from the desire to run multiple operating systems on the same piece of hardware and get the benefit of sharing computing resources. VMs are used to run software applications, and with virtualization, one physical server can run many applications as each application runs in a VM sharing the resources of the physical machine. Different application needs on the same physical hardware can create conflicts in complex systems—each piece of software competing for limited processing and storage capabilities of a physical server. Simplified, VMs allow time and capacity sharing among competing applications and their operating systems. If one physical machine can host multiple VMs, the utilization of its resource capacity becomes more efficient, and requires fewer physical resources, and software licenses, and hence costs less.

22. In virtualized environments, the physical server and related software often are called a “host.” VMs running on a host often are called “guests.” Virtualization is accomplished through software that makes logical computing environments independent of physical infrastructure. In this context, virtualization refers to dividing up resources of a host into multiple VMs (i.e., multiple virtual servers). With virtualization, one host can run many applications because each run on a VM sharing resources of the host.

23. Virtualization also allows for multiple hosts to be configured as one larger logical entity; the resulting group of hosts can be referred to as a “cluster.” A cluster offers an advantage of managing several hosts as one larger resource pool. Each host can be interchangeable in the pool, which enables VMs to move between them. The moves could happen, for example, if a host becomes too busy and another has more available resources.

24. Decisions must be made regarding which VMs to run on which hosts. Running too few VMs or applications on a host means that more servers or capacity must be purchased. Running too many VMs or applications can create risk by over-utilizing hardware, i.e., it creates resource contention where VMs compete for the same resources. The inefficiencies and risks can be reduced by optimizing workload placements and making sure the amount of resources assigned to each VM is commensurate with host constraints.

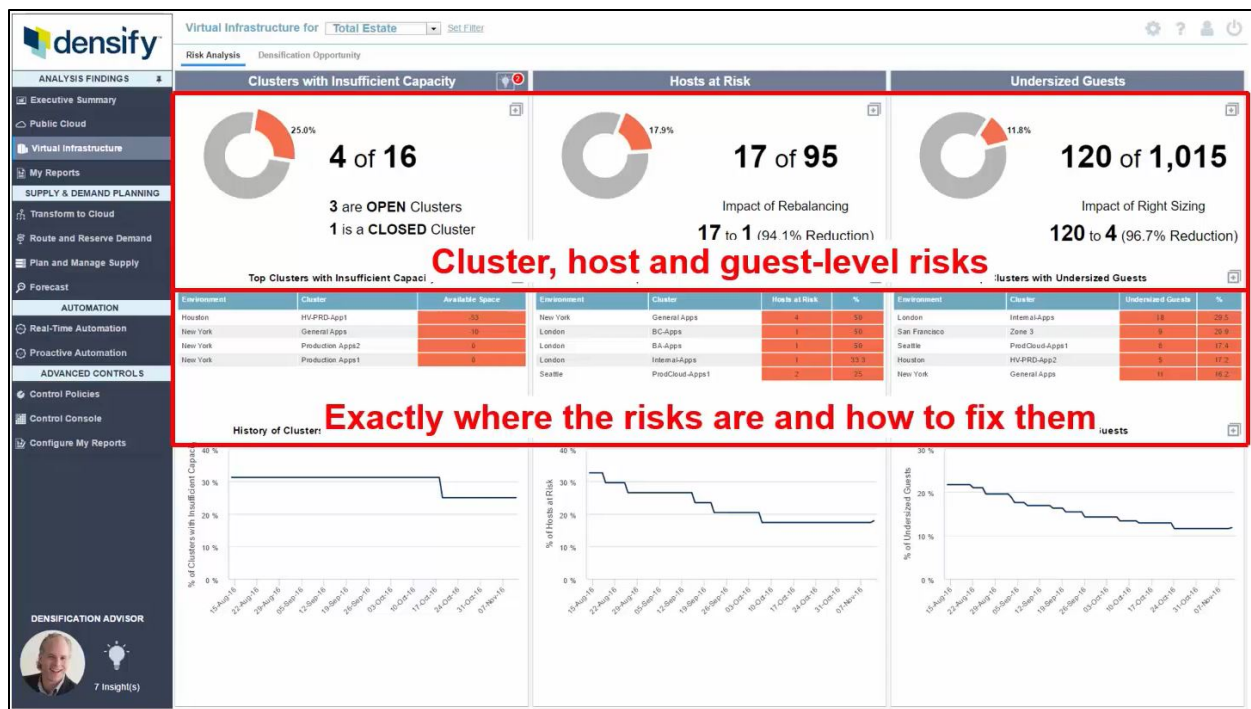
25. Even with promise of greater efficiency and scale, companies were not able to optimize a complex virtualized infrastructure on a day-to-day basis. For example, in virtualized environments, determining optimal placements of VMs in short times, while honoring complex operating constraints, was not practical.

26. Densify’s technology addresses these needs. With Densify’s analytics, managing virtualized infrastructures to an optimal state can be done with automation. In fact, intelligent automation has become critical as environments scale and become more complex.

27. Densify enables automated infrastructure optimization for virtualized environments through its proprietary software. Densify offers products that optimize virtual environments and are referred to herein as “Densify’s Optimization Engine.” Densify’s Optimization Engine is predictive analytics software that optimizes public cloud, bare metal cloud, and on-premise virtualized environments, enabling customers to operate with less cloud cost, less infrastructure and better performing applications.


28. Densify’s Optimization Engine decreases risk in a VM infrastructure. Its predictive analytics anticipate capacity risk, place workloads, and allocate resources to avoid capacity shortfalls, meet compliance and other key operating policies, which results in reducing unnecessary movement of VMs and avoiding application performance issues. The figure below is an example of how Densify’s Optimization Engine displays risks in a customer’s environment across clusters, hosts, and guests.¹ The display shows, for example, how: (1) specific placement and allocation recommendations can address risks; (2) the customer can automate and execute those recommendations; and (3) to track the progress over time.

¹ Screenshot from Densify Video, Cloud Optimization Done For You, available at <https://www.densify.com/resources/video-optimize-on-premises-virtual-infrastructure-bare-metal-cloud> (Ex. 1). A full version of the video from which the screenshot (Ex. 1) was taken has been provided as Ex. 31.



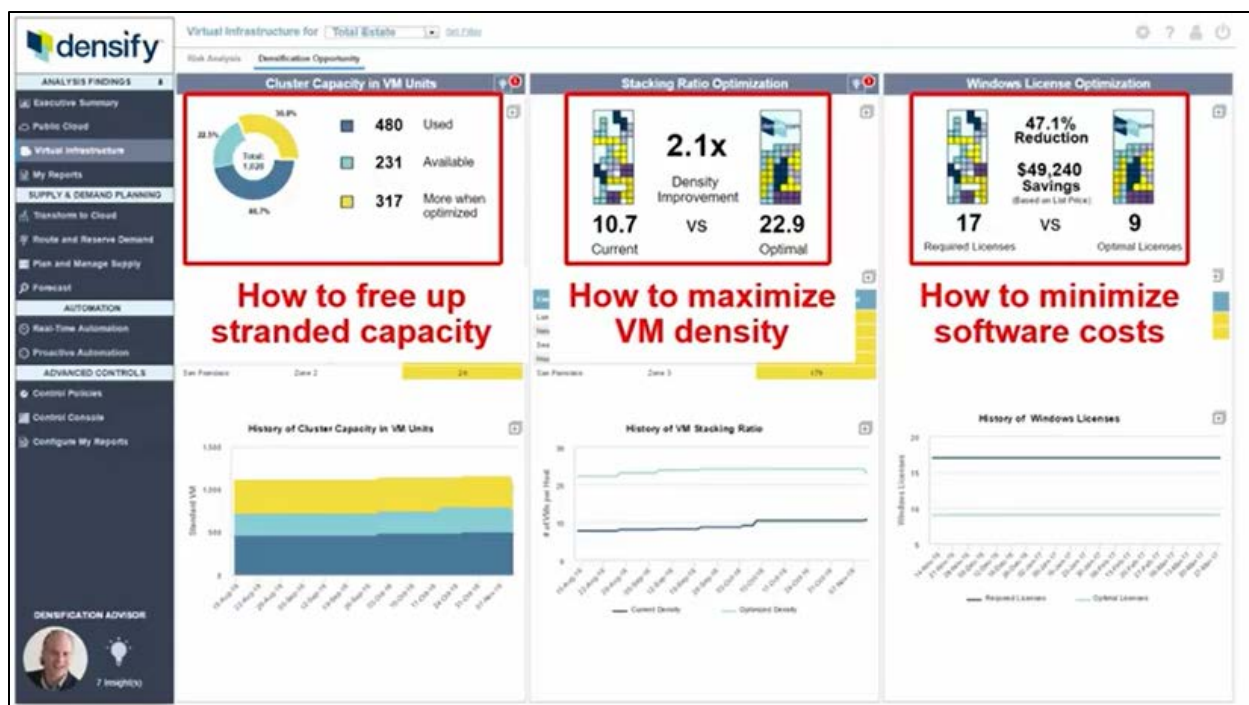
29. The Densify Optimization Engine reduces infrastructure costs. Its intelligent workload stacking and resource allocation increases optimization by identifying savings of an average of 33% on hardware and 55% on software licensing. The picture below illustrates, by way of analogy, the effect of inefficient stacking, creating operational risk and stranded capacity, and then how Densify’s Optimization Engine organizes the workload to reduce risks and inefficiencies.²

² Screenshot from Densify Video, Cloud Optimization Done For You, available at <https://www.densify.com/resources/video-optimize-on-premises-virtual-infrastructure-bare-metal-cloud> (Ex. 2). A full version of the video from which the screenshot (Ex. 2) was taken has been provided as Ex. 31.



30. Densify’s Optimization Engine automates workload placement and resource optimization. The figure below illustrates how Densify’s Optimization Engine provides customer control over making stranded capacity available, how to increase the density of VMs, and how to minimize software costs.³

³ Screenshot from Densify Video, Cloud Optimization Done For You, available at <https://www.densify.com/resources/video-optimize-on-premises-virtual-infrastructure-bare-metal-cloud> (Ex. 3). A full version of the video from which the screenshot (Ex. 3) was taken has been provided as Ex. 31.



31. Through its innovative product offerings, Densify has attracted a substantial customer following among the *Fortune 5000*. Customers pay millions of dollars for Densify's Optimization Engine, which in turn save those companies many more millions of dollars. Densify has invested heavily in research and development to create innovative, award winning products that have been in high demand from some of the world's leading companies. Densify's customer following and awards are the result of its patented innovations, including those found in the patents asserted in this case.

B. VMware And Its Technology

32. VMware is a company focused on platform virtualization and cloud computing software and services. Founded in 1998, it entered the server market in 2001. It began as a leader in virtualization software and has expanded its offerings over the years. It is known for its innovation of the hypervisor—the layer of software residing between the operating system and the system hardware that enables virtualization.

33. VMware's early success was grounded in its innovations relating to server virtualization. But, as analysts have recognized, it has been hesitant to embrace new technologies: "The history of VMware presents a company that disrupted the IT industry with server virtualization, but that was hesitant to embrace new technologies, such as the public cloud. VMware's future success will depend on how it finds a role in the current cloud market and how it approaches containers."⁴

34. "The concept of virtualization has moved from the server to other areas in the data center infrastructure. The history of VMware shows a company trying to keep up with its original innovation."⁵ Over time, VMware has been successful in raising money, with which it acquired companies with innovative technology and developed products using the technology of others.

35. In 2008, "after a disappointing financial performance due to mismanagement of the company's growing scale, the board of directors replaced VMware president and CEO Diane Greene with Paul Maritz, a former Microsoft veteran who headed EMC's cloud computing business unit. This leadership reorganization marked a point in the history of VMware that foreshadowed a shift in the company's direction."⁶

36. Thereafter, VMware began rapidly acquiring innovations from other companies, and, at least in the case of Densify, copying a competitor's technology. The history shows VMware has been taking an increasingly aggressive approach in maintaining its market power as the market shifts and new innovators spring up.

⁴ Korzeniowski, P., "What The History Of VMware Reveals About Its Future Projects," available at <https://searchvmware.techtarget.com/tip/What-the-history-of-VMware-reveals-about-its-future-projects> (Ex. 4).

⁵ *Id.* (Korzeniowski, Ex. 4).

⁶ *Id.* (Korzeniowski, Ex. 4).

37. Today, VMware faces challenges in adapting its business to the public cloud as companies shift away from on premise to cloud-based infrastructures. “VMware Cloud on AWS isn’t as attractive for organizations developing new applications. VMware maintains its position with data center technicians, but it’s not growing with business unit developers who increasingly control significant portions of IT spending. . . . VMware’s virtualization software reshaped the computer industry. Cloud and containers are now having a similar effect. VMware’s response to these trends will determine its impact in the coming years.”⁷

38. This case centers on VMware’s products and services related to its virtualization platform, including but not limited to, vRealize Operations (“vROps”), vRealize Automation (“vRA”), Distributed Power Manager (“DPM”), Distributed Resource Scheduler (“DRS”), Storage DRS (“sDRS”), High Availability (“HA”), and other related products and services, including VMware’s suite of software products and services (e.g., VMware Cloud Foundation, Project Dimension, vSphere, vCenter Server, vCloud Director, Cloud Provider Pod, vRealize Suite, vRealize Suite Lifecycle Manager, vCloud Suite, vRealize Business for Cloud, vRealize Operations for Horizon, vCloud NFV) that include vROps, vRA, DPM, DRS, sDRS, or HA (collectively, the “VMware Accused Products”).

39. According to VMware, vROps is an enterprise software product that “can proactively identify and solve emerging issues with predictive analysis and smart alerts, ensuring optimal performance and availability of system resources - across physical, virtual, and cloud infrastructures.”⁸ VMware states that vROps provides “complete monitoring capability in one

⁷ *Supra* n. 4 (Korzeniowski, Ex. 4).

⁸ vRealize Operations Manager 7.0 Help Guide at 5 (Nov. 15, 2018), available at <https://docs.vmware.com/en/vRealize-Operations-Manager/7.0/vrealize-operations-manager-70-help.pdf> (Ex. 5).

place, across applications, storage, and network devices, with an open and extensible platform supported by third-party management packs.”⁹ In addition, according to VMware, vROps “increases efficiency by streamlining key processes with preinstalled and customizable policies while retaining full control.”¹⁰ Using data collected from system resources (objects), vROps “identifies issues in any monitored system component, often before the customer notices a problem.”¹¹

40. VMware’s core hypervisor is referred to as ESXi. VMware’s DRS is a tool for managing VM workloads and “works on a cluster of ESXi hosts and provides resource management capabilities like load balancing and virtual machine (VM) placement. DRS also enforces user-defined resource allocation policies at the cluster level, while working with system-level constraints.”¹² VMware states that DRS ensures that “VMs and their applications are always getting the compute resources that they need to run efficiently.”¹³

41. VMware describes vRA as a tool that “empowers IT to accelerate the provisioning and delivery of IT services, across infrastructure, containers, applications and custom services.”¹⁴ vRA is said to further “provides a secure portal where authorized administrators, developers, or business users can request new IT services. In addition, they can manage specific cloud and IT

⁹ *Id.* (vRealize Operations Manager 7.0 Help Guide, Ex. 5).

¹⁰ *Id.* (vRealize Operations Manager 7.0 Help Guide, Ex. 5).

¹¹ *Id.* (vRealize Operations Manager 7.0 Help Guide, Ex. 5).

¹² Understanding vSphere DRS Performance, VMware vSphere 6 at 4, available at <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vsphere6-drs-perf.pdf> (Ex. 6).

¹³ *Id.* (Understanding vSphere DRS Performance, Ex. 6).

¹⁴ VMware vRealize Automation Datasheet at 1, available at <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vrealize/vmware-whats-new-vrealize-automation.pdf> (Ex. 37).

resources that enable IT organizations to deliver services that can be configured to their lines of business in a self-service catalog.”¹⁵

42. According to VMware, DPM “provides additional power savings beyond this initial benefit by dynamically consolidating workloads even further during periods of low resource utilization. Virtual machines are migrated onto fewer hosts and the un-needed ESX hosts are powered off. When workload demands increase, ESX hosts are powered back on and virtual machines are redistributed to them. VMware DPM is an optional feature of VMware® Distributed Resource Scheduler (VMware DRS).”¹⁶

43. VMware describes sDRS as a tool that “allows you to manage the aggregated resources of a datastore cluster. When Storage DRS is enabled, it provides recommendations for virtual machine disk placement and migration to balance space and I/O resources across the datastores in the datastore cluster.”¹⁷

44. HA is described as a tool that “provides high availability for virtual machines by pooling the virtual machines and the hosts they reside on into a cluster. Hosts in the cluster are monitored and in the event of a failure, the virtual machines on a failed host are restarted on alternate hosts.”¹⁸

¹⁵ Foundations and Concepts at 5 (May 7, 2019), available at <https://docs.vmware.com/en/vRealize-Automation/7.5/vrealize-automation-75-foundations-and-concepts.pdf> (Ex. 38).

¹⁶ VMware Distributed Power Management Concepts and Use at 3 (2010), available at <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/Distributed-Power-Management-vSphere.pdf> (Ex. 39).

¹⁷ VMware vSphere Resource Management at 105 (Jan. 11, 2019), available at <https://docs.vmware.com/en/VMware-vSphere/6.7/vsphere-esxi-vcenter-server-671-resource-management-guide.pdf> (Ex. 40).

¹⁸ vSphere Availability at 11 (Apr. 11, 2019), available at <https://docs.vmware.com/en/VMware-vSphere/6.7/vsphere-esxi-vcenter-server-672-availability-guide.pdf> (Ex. 41).

C. Densify and VMware's Relationship

45. Densify's business historically has focused on technology that can optimize the placement of virtual machines on hosts while under the multiple operational constraints. It works with virtualization platforms offered by, among others, VMware. Densify's Optimization Engine historically has complemented VMware's products, but as VMware learned Densify's technology, it slowly over the years has become Densify's competitor in virtual infrastructure optimization.

46. IT organizations of major companies typically utilize virtual environments, and most of these companies use VMware's virtualization platform (i.e., hypervisor and related software).

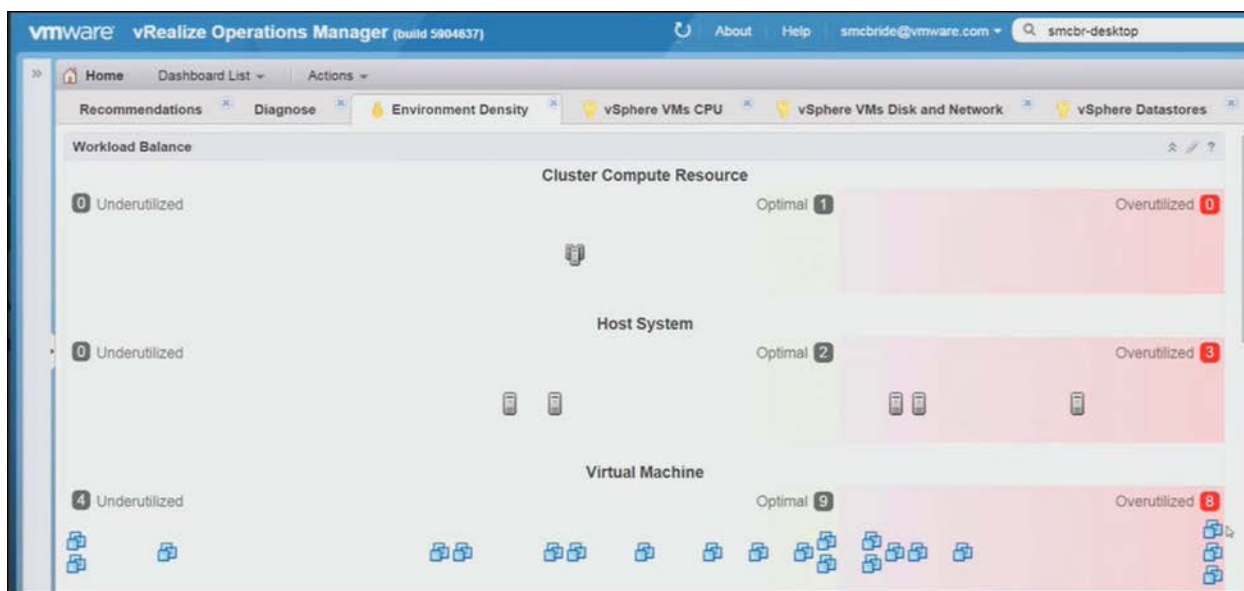
47. Densify integrates with VMware's virtualization platform, including working with VMware management products like DRS and vROps.

48. Because VMware offers the baseline virtualization platform with approximately 99% of *Fortune 1000* companies as customers, it is naturally the dominant "incumbent" with direct access to customers. VMware uses its incumbent position and market dominance to consolidate its power, particularly when it comes to competitors. For example, on information and belief, VMware controls access to the industry's leading tradeshow (called "VMworld"), including speaker and analyst invitations. By doing so, VMware is capable of controlling the agenda for what is known to be a must-attend virtualization tradeshow.

D. VMware Copied Densify's Technology

49. VMware has long been familiar with Densify's product, technology and inventions. VMware embarked on a strategy to introduce products with Densify's technology and intellectual property.

50. As an example, VMware previewed its release of vROps 6.1 at the VMworld 2015 trade show. vROps 6.1 was shown with the following dashboard:¹⁹



51. vROps' dashboard looked the same in relevant respects as Densify's product that was in the market since 2012, shown below.²⁰ In fact, an audience member spoke up during the VMworld conference presentation to note how similar vROps dashboard was.

¹⁹ Screenshot from VMware Videos, How to Troubleshoot Using vRealize Operations Manager (Deep Live Demo) (MGT4928-1) and How to Troubleshoot Using vRealize Operations Manager (Deep Live Demo) (MGT4928-2), available at <https://videos.vmworld.com/global/2015?q=MGT4928> (Ex. 7). Full versions of the videos from which the screenshot (Ex. 7) was taken have been provided as Exs. 32 – 33.

²⁰ Screenshot of Densify's Product (Ex. 8).



Screenshot of Densify's Product

52. As seen above, Densify's product and VMware's vROps 6.1's dashboard share substantially the same key features (e.g., "too little infrastructure" is analogous to "overutilized"). The similarity of VMware's dashboard to Densify's was so striking that Densify received a call from a customer who mistakenly assumed VMware must have acquired Densify.

53. Upon information and belief, the dashboard in paragraph 47 remains a vROps feature.

54. VMware also released "predictive DRS" that was in material respects the same as a core Densify feature that VMware knew about and copied.

55. These releases were troublesome and infringing, but still VMware was not able to offer a product, like Densify's, that satisfied the needs of customers with complex business and operating constraints.

56. Then, in late September 2018, VMware released vROps version 7.0. vROps 7.0 added the "Automated Host Based Placement" feature, which allows vROps to "teach DRS your business intent and control not only balancing across clusters, but also which host within a cluster

the workload will land on.”²¹ The particular use case for this feature is software license control²²— a key Densify product feature known to VMware.

57. Upon information and belief, prior to the vROps 7.0 release, vROps did not have a host-based software license control feature. With the “Automated Host Based Placement” feature, however, VMware, for the first time, provides a commercially viable host-based software license control feature that allows a customer to “automatically drive” its “business intent” from a single and automated user interface enabling the type of software license control feature Densify described as important to its customers. VMware has not stopped there. In a recent announcement, VMware has emphasized that it is further enhancing the vROps host-based placement software license control features in a brand new vROps 7.5 release.

58. Not only did VMware copy Densify’s technology, it even began using its name and trademark DENSIFY (or formatives thereof) in reference to VMware’s products. For example, it began identifying its key feature as “Workload Optimization – Densifying to Repurpose Hosts.” An example is below:²³

²¹ Gandhi, T., What’s New in vRealize Operations 7.0, VMware Blogs, available at <https://blogs.vmware.com/management/2018/08/whats-new-in-vrealize-operations-7-0.html> (Ex. 9).

²² Dias, J., Using Host Rules with Business Intent in vRealize Operations 7.0, VMware Blogs, available at <https://blogs.vmware.com/management/2018/11/using-host-rules-with-business-intent-in-vrealize-operations-7-0.html> (Ex. 10).

²³ Screenshot from Workload Optimization - Densifying to Repurpose Hosts, VMware Cloud Management, available at <https://www.youtube.com/watch?v=IymKQdg3oNE> (Ex. 11). A full version of the video from which the screenshot (Ex. 11) was taken has been provided as Ex. 34.

Workload Optimization
Business Intent: **Densify to Repurpose Hosts**

Use Case: You are looking to free up some hosts for another business need

- Move workloads until the cluster is “Full”
- Honor 20% Headroom setting
- Free up some hosts

Play (k)

vmware 0:13 / 2:52

E. VMware’s Infringement of Densify’s Intellectual Property

59. VMware implemented features in VMware products that infringed Densify’s patents.

1. The ’687 Patent

60. On June 26, 2012, the U.S. Patent and Trademark Office duly and legally issued the ’687 patent, entitled “Method and System for Evaluating Virtualized Environments,” with Tom Silangan Yuyitung and Andrew Derek Hillier as inventors. The earliest application related to the ’687 patent was filed on August 31, 2007. A true and correct copy of the ’687 patent is attached as Exhibit 12.

61. The ’687 patent is directed to technological solutions that addresses problems specifically grounded in enterprise IT environments. For example, the ’687 patent explains that “IT infrastructures used by many organizations have moved away from reliance on centralized computing power and towards more robust and efficient distributed systems. While the benefits

of a distributed approach are numerous and well understood, there has arisen significant practical challenges in managing such systems for optimizing efficiency and to avoid redundancies and/or under-utilized hardware. In particular, one challenge occurs due to the sprawl that can occur over time as applications and servers proliferate. Decentralized control and decision making around capacity, the provisioning of new applications and hardware, and the perception that the cost of adding server hardware is generally inexpensive, have created environments with far more processing capacity than is required by the organization.”²⁴

62. The '687 patent also states that “[w]hen cost is considered on a server-by-server basis, the additional cost of having underutilized servers is often not deemed to be troubling. However, when multiple servers in a large computing environment are underutilized, having too many servers can become a burden. Moreover, the additional hardware requires separate maintenance considerations; separate upgrades and requires the incidental attention that should instead be optimized to be more cost effective for the organization. Heat production and power consumption can also be a concern. Even considering only the cost of having redundant licenses, removing even a modest number of servers from a large computing environment can save a significant amount of cost on a yearly basis.”²⁵

63. The emergence of virtual infrastructure provided a technology foundation to achieve consolidation, and according to the '687 patent, “organizations have become increasingly concerned with such redundancies and how they can best achieve consolidation of capacity to reduce operating costs.”²⁶ The problems IT organization faced at the time were that “[c]omplex systems configurations, diverse business requirements, dynamic workloads and the heterogeneous

²⁴ '687 patent, 1:26-41.

²⁵ '687 patent, 1:42-54.

²⁶ '687 patent, 1:55-57.

nature of distributed systems can cause incompatibilities between systems. These incompatibilities limit the combinations of systems that can be consolidated successfully. In enterprise computing environments, the virtually infinite number of possible consolidation permutations which include suboptimal and incompatibility system combinations make choosing appropriate consolidation solutions difficult, error-prone, and time consuming.”²⁷ This virtually infinite number of possible consolidation permutations make it impossible to choose the appropriate consolidation with mental processes (i.e., performing the calculations using pen and paper). Attempting to manually perform the claimed invention directed to a problem of computer technology, would not achieve the results described in the ’687 patent. The ’687 patent claimed a way to solve technological problems that existed within the field of virtual environments. It provides a technological solution to a problem rooted in computer technology, improving the way server networks function. It also provides a non-conventional technique that generates new data for analyzing the virtualized computer environment.

64. The ’687 patent addresses the technological problems not by a mere nominal application of a generic computer to practice the invention, but by “recogniz[ing] that virtualization often involves more than considering sizing, for example, it is beneficial to understand all the constraints that govern and impact a target environment and ensure that these constraints are taken into account when planning and managing a virtual environment. This has been found to be particularly true of virtualization infrastructures such as VMware Infrastructure®, where sophisticated features such as vMotion, distributed resource scheduling (DRS) and HA require careful planning and diligent administration of virtual environments. It has been found that to fully realize the capabilities of the virtualization infrastructure, the virtualization scheme being used should be combined with accurate

²⁷ ’687 patent, 2:12-21.

intelligence and focused analytics in order to safely and effectively transform existing systems into a new virtual paradigm. In order to provide such intelligence and focused analytics, an analysis program for determining compatibilities in a computing environment . . . can be utilized along with specific virtualization rule sets and user interfaces (UIs) to address the considerations of a virtualization infrastructure.”²⁸ The ’687 patent addresses the shortcomings in prior systems, discussed above, with its improvement in the way virtualized servers function.

65. The ’687 patent recites “evaluating the placement of said virtual machines in said virtualized environment using said data sets by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to said technical, business and workload constraints to determine guest-host placement.”²⁹ By determining host placements based on these constraints, the ’687 patent improves the function of servers as compared to prior systems by improving efficiency, reducing redundancies and under-utilized hardware, reducing or eliminating unnecessary processing capacity and incompatibilities, reducing errors, decreasing costs, requiring less maintenance, decreasing heat and power consumption, ease of maintaining compliance or risk de-concentration, and generally being easier to manage.

66. Rule sets described in detail in the ’687 patent, as well as the application incorporated by reference, U.S. patent application Ser. No. 11/535,308, entitled “Method for Evaluating Computer Systems.” For example, the ’687 patent describes that “[r]ules comprised by a rule set **28** may reference common parameters but perform different tests to identify different forms of incompatibilities that may have different levels of importance. For example, a version four operating system versus a version three operating system may be considered less costly to

²⁸ ’687 patent, 5:52-6:4.

²⁹ ’687 patent, 38:64-39:2.

remedy and thus less detrimental than a version five operating system compared to a version one operating system. As can be seen, even though the operating systems are different in both cases, the nature of the difference can also be considered and different weights and/or remedies applied accordingly. Rules can also test for similarities that indicate contentions which can result in incompatibilities between systems. For example, rules can check for name conflicts with respect to system names, database instance names, user names, etc.”³⁰

67. The '687 patent also discusses generating a new virtual environment design for virtualizing computer system and conducting compatibility analyses, and refer to consolidating virtual environments by moving from one system to another.³¹ The compatibility analyses discussed in the '687 patent include N-to-1, N-by-N, and 1-to-1 analyses, which are multi-dimensional compatibility analyses that make possible the VM “optimal placements” discussed in the '687 patent. They make possible the tangible benefits that are not well-known or conventional, such as greater IT efficiencies, reduced operating costs, faster workload deployment, increased application performance, higher server availability, ease of maintaining compliance or risk de-concentration, and reduced complexity. As such, the “intelligence and focused analytics” the '687 patent is directed to provide a specific improvement over prior systems that is not well-known or conventional, resulting in an improved automated evaluation of IT infrastructure for purposes of optimization. For example, VMware has acknowledged that its prior versions “did not support affinity between VMs and hosts,” but as it described in an article about “lessons learned,” it added that feature, which is covered by the '687 patent, because “[w]hile VM-to-VM affinity was sufficient for most technical use-cases, there were other requirements such as software licensing

³⁰ '687 patent, 12:19-33.

³¹ '687 patent, 37:28-67.

that made administrators want to isolate VMs onto a set of hosts.”³² The ’687 patent is tailored to these improvements in the field of virtual environments, one specific area of computer technology.

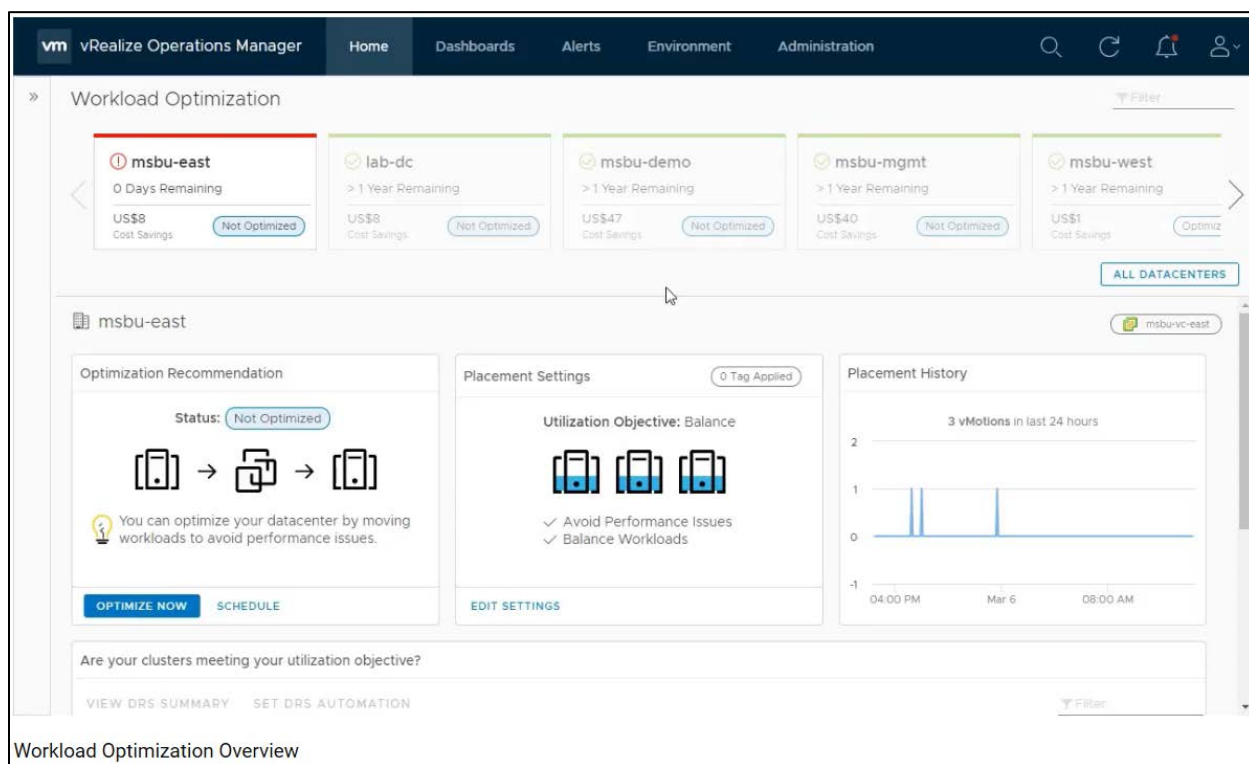
68. Claim 7 is directed to a “method for validating an existing virtualized environment comprising a plurality of virtual machines placed on one or more virtual hosts.” It recites “obtaining a data set for each of said plurality of virtual machines, each data set comprising information pertaining to technical, business, and workload constraints.” Claim 7’s method also involves “evaluating each virtual guest against each virtual host and other virtual guest using one or more rule sets pertaining to said technical, business, and workload constraints to determine guest-host placements.” Claim 7 further recites “identifying the existence of virtual machines with suboptimal placements to enable alternative placements for said virtual machines.”

69. The VMware Accused Products directly infringe at least claim 7. For example, vROps infringes at least claim 7 when it conducts “Capacity Optimization,” “Workload Optimization,” including the newly added host-based placement feature, “What-If Analyses,” and the installation of vROps. When these features are executed, vROps validates an existing virtualized environment by evaluating each virtual guest with each virtual host and other virtual guests by using a constraint-based analysis (e.g., technical, workload, and business constraints) and identifies virtual machines in the virtualized environment that are not optimally placed.

70. vROps utilizes a method for validating an existing virtual environment comprising a plurality of virtual machines placed on one or more hosts. For example, vROps validates virtual environments comprising of virtual machines and hosts as it conducts “continuous performance

³² Gulati, et al., VMware Distributed Resource Management: Design, Implementation, and Lessons Learned, available at <https://labs.vmware.com/vmtj/vmware-distributed-resource-management-design-implementation-and-lessons-learned> (Ex. 13).

optimization . . . driven by business and operational intent”³³ vROps validates and manages virtualized environments by automating and simplifying “IT operations management” and by providing “unified visibility from applications to infrastructure across physical, virtual, and cloud environments.”³⁴ Generally, vROps assesses the placements of virtual machines on hosts using multiple criteria, and therefore, vROps is capable of validating whether an environment in is an optimized state. It validates an existing virtualized environment that includes VMs and hosts, as shown in the video below:³⁵



³³ Self-Driving Operations by VMware vRealize Operations, Datasheet at 1, available at <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vCenter/vmware-vrealize-operations-datasheet.pdf> (Ex. 14 (highlighting added)).

³⁴ *Id.* (Self-Driving Operations by VMware vRealize Operations Datasheet at 1, Ex. 14 (highlighting added)).

³⁵ Screenshot from VMware Video, VMware vRealize Operations, Workload Optimization – Overview, available at <https://www.youtube.com/watch?v=w7Y4uEayltM&feature=youtu.be> (Ex. 15). A full version of the video from which the screenshot (Ex. 15) was taken has been provided as Ex. 35.

71. vROps also obtains a data set for each of said plurality of virtual machines, each data set comprising information pertaining to technical, business, and workload constraints. For example, vROps obtains “performance data from monitored software and hardware resources in your enterprise and provides predictive analysis and real-time information about problems.”³⁶ It “monitors your ESXi hosts and the virtual machines located on them,” and “monitors virtual machines running in a vCenter Server, analyzes longer-term historical data, and provides forecast data about predictable patterns of resource usage to Predictive DRS. Based on these predictable patterns, Predictive DRS moves to balance resource usage among virtual machines.”³⁷ Predictive DRS is a feature included in vROps. It also computes “analytics [that] provide precise tracking, measuring, and forecasting of data center capacity, usage, and trends to help manage and optimize resource use, system tuning, and cost recovery.”³⁸

72. Examples of technical (e.g., configuration) and workload (e.g., utilization/performance) data vROps uses can be found illustrated in the various components of the interface shown below.³⁹ Further, for technical constraints: “The VM dashboard focuses on highlighting the key configurations of the virtual machines in your environment. You can use this dashboard to find inconsistencies in configuration within your virtual machines and take quick remedial measures.”⁴⁰ And vROps can “[use] the Workload Utilization widget to identify which workload objects are underutilized and overutilized.”⁴¹

³⁶ *Supra* n. 8 (vRealize Operations Manager 7.0 Help Guide at 486, Ex. 5).

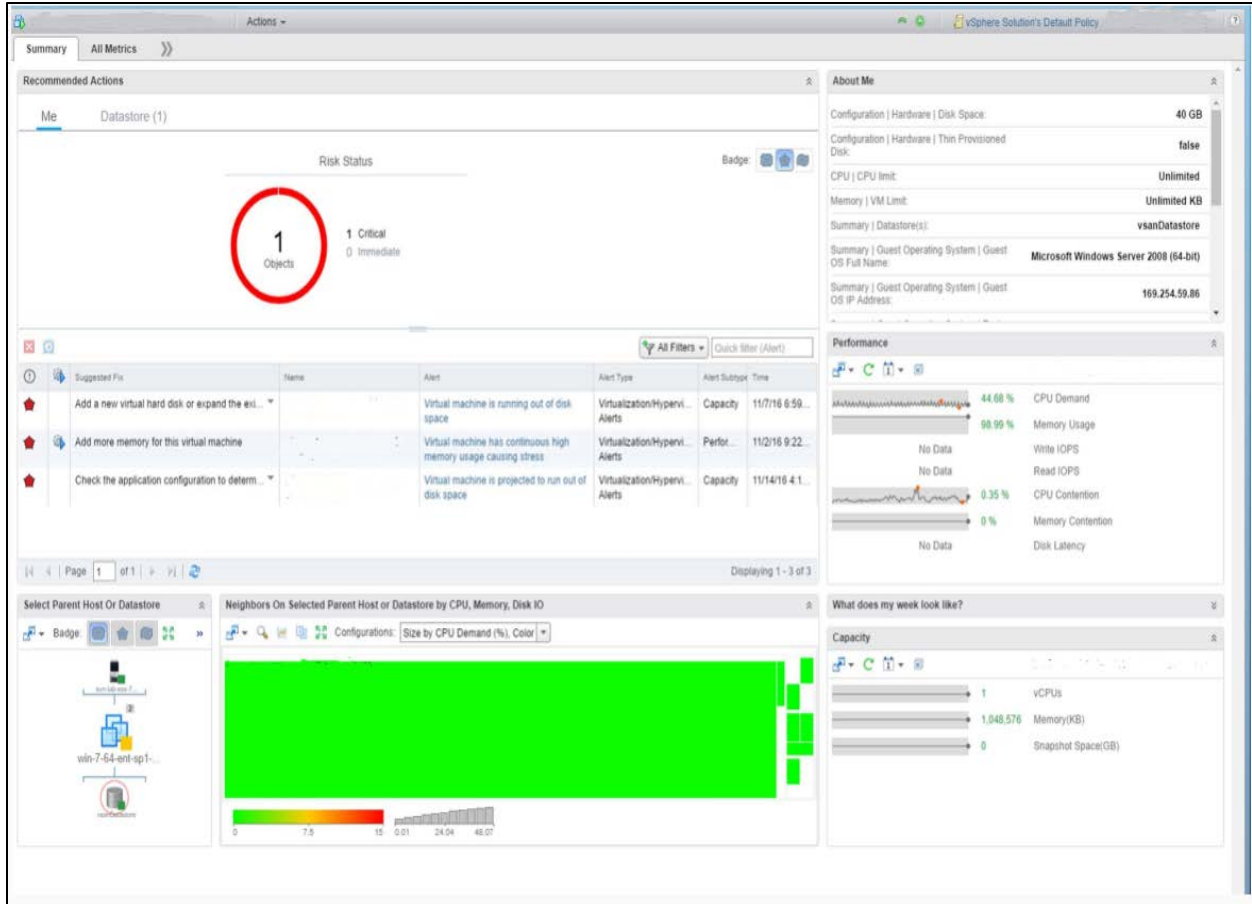
³⁷ *Id.* (vRealize Operations Manager 7.0 Help Guide at 120, Ex. 5).

³⁸ *Id.* (vRealize Operations Manager 7.0 Help Guide at 770, Ex. 5).

³⁹ *Id.* (vRealize Operations Manager 7.0 Help Guide at 684, Ex. 5).

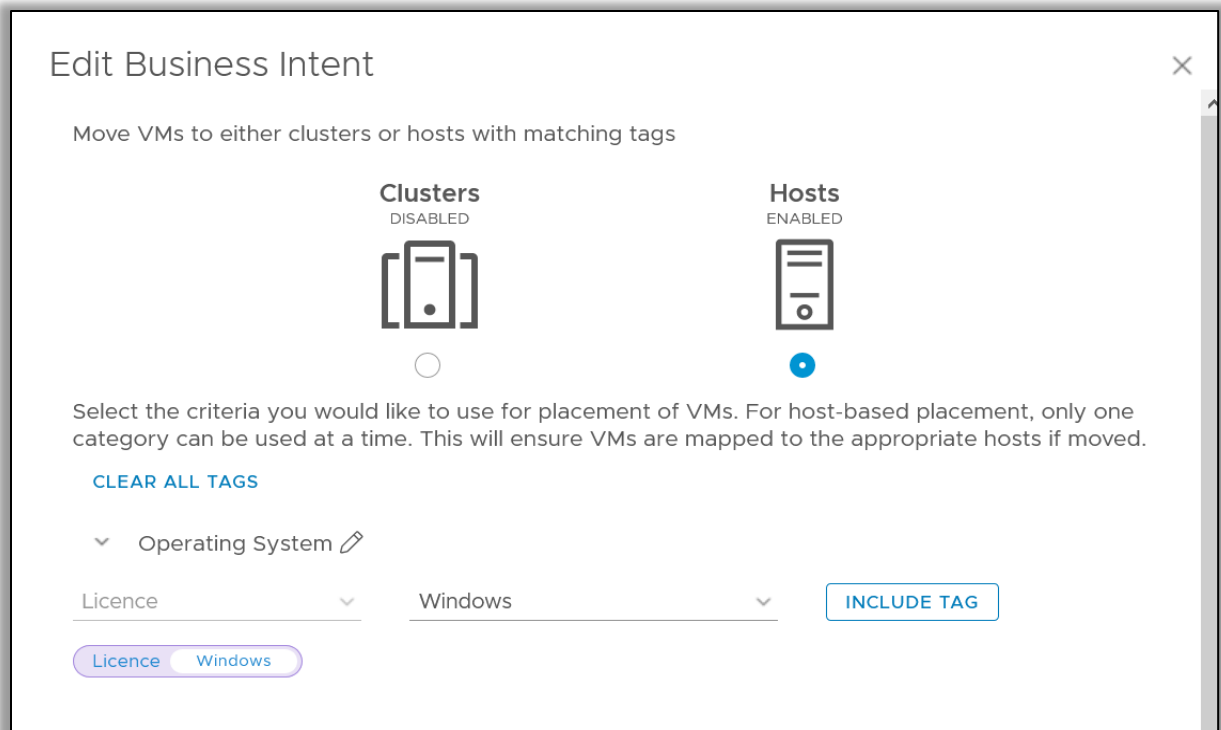
⁴⁰ *Id.* (vRealize Operations Manager 7.0 Help Guide at 496, Ex. 5).

⁴¹ *Id.* (vRealize Operations Manager 7.0 Help Guide at 485, Ex. 5).



73. Further, vROps uses business constraints corresponding to VMs as shown below:⁴²

⁴² *Supra* n. 22 (Dias, Ex. 10).



74. Further, for business constraints: “You can use vCenter Server tagging to tag VMs, hosts, and/or clusters with specific tags. vRealize Operations Manager can be configured to leverage tags to define business-related placement constraints: VMs can only be placed on hosts/clusters with matching tags.”⁴³

75. vROps also evaluates the placement of said virtual machines in said virtualized environment using said data sets by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to said technical, business and workload constraints to determine guest-host placements. For example, as depicted above, vROps evaluates placements of virtual machines when conducting “Workload Optimization.” vROps “monitors virtual objects and collects and analyzes related data that is presented to users in graphical form at the Workload Optimization screen. Depending on what appears on the screen, you might use

⁴³ *Supra n. 8* (vRealize Operations Manager 7.0 Help Guide at 631, Ex. 5).

optimization functions to distribute a workload differently in a data center or customer data center. Or you may decide to perform more research, including checking the Alerts page to determine if any alerts have been generated for objects or interest.”⁴⁴ It “provides for moving virtual compute resources and their file systems dynamically across datastore clusters within a data center or custom data center. Using Workload Optimization, you can rebalance virtual machines and storage across clusters, relieving demand on an overloaded individual cluster and maintaining or improving cluster performance. You can also set your automated rebalancing policies to emphasize VM consolidation, which potentially frees up hosts and reduces resource demand.”⁴⁵

76. In addition, vROps’ “Workload Optimization offers you the potential to automate fully a significant portion of your cluster workload rebalancing tasks” by, among other things, tagging virtual machines for “Host-Based Virtual Machine Placement” and “Tag-Based VM Placement in Clusters.”⁴⁶ “[W]hen the system runs an optimization, it uses VM-to-host tag matching to ensure that VMs are moved to - or stay with – the appropriate host.”⁴⁷ vROps uses “host-based VM placement to tie your VMs more closely to your infrastructure. By using vCenter Server to tag hosts and VMs with specific tags, you make certain that when the system runs an optimization, it uses VM-to-host tag matching to ensure that VMs are moved to – or stay with – the appropriate host.”⁴⁸ Accordingly, vROps evaluates each virtual guest against each virtual host and other virtual guests when optimizing the workload across clusters and hosts.

⁴⁴ *Supra* n. 8 (vRealize Operations Manager 7.0 Help Guide at 633, Ex. 5).

⁴⁵ *Id.* (vRealize Operations Manager 7.0 Help Guide at 625, Ex. 5).

⁴⁶ *Id.* (vRealize Operations Manager 7.0 Help Guide at 626, Ex. 5).

⁴⁷ *Id.* (vRealize Operations Manager 7.0 Help Guide at 630, Ex. 5).

⁴⁸ *Id.* (vRealize Operations Manager 7.0 Help Guide at 630, Ex. 5).

77. Lastly, vROps identifies the existence of virtual machines with suboptimal placements to enable alternative placements for virtual machines. vROps meets this limitation when it conducts workload optimization, which identifies sub-optimal VM placements.

78. Based on the above, the VMware Accused Products directly infringe at least, but not limited to, claim 7 of the '687 patent.

2. The '367 Patent

79. On May 16, 2017, the U.S. Patent and Trademark Office duly and legally issued the '367 patent, entitled "System and Method for Determining and Visualizing Efficiencies and Risks in Computing Environments," with Andrew Derek Hillier as inventor. The earliest application related to the '367 patent was filed on August 16, 2011. A true and correct copy of the '367 patent is attached as Exhibit 16.

80. The '367 patent is directed to technological solutions with practical applications that address problems specifically grounded in enterprise IT environments. For example, the '367 patent explains that "[m]odern data centers typically comprise hundreds if not thousands of servers. Each server supplies a finite amount of resource capacity, typically in the form of, but not limited to: central processing unit (CPU) capacity, memory or storage capacity, disk input/output (I/O) throughput, and network I/O bandwidth. Workloads running on these servers consume varying amounts of these resources. With the advent of virtualization and cloud technologies, individual servers are able to host multiple workloads."⁴⁹

81. The '367 patent further explains that "[p]ercent CPU utilization, which corresponds to the ratio of CPU usage relative to CPU capacity, is a common measure of how effectively servers are being utilized. Various other metrics may be used to determine resource utilization for

⁴⁹ '367 patent, 1:23-31.

computing systems. Organizations may wish to measure and evaluate efficiencies and risks in computing environments but often do not have convenient ways to perform such measurements and evaluations.”⁵⁰

82. The ’367 patent addresses these technological problems by not only a mere nominal application of a generic computer to practice the invention, but by requiring the computations of “at least one score quantifying efficiencies and/or risks associated . . . based on the resource utilization or performance data, the capacity data, and the at least one operational policy.”⁵¹ The ’367 patent also requires the “displaying an indicator for at least one of the plurality of computing entities in a graphical representation based on the corresponding score.”⁵² Such scores and indicators provide for a convenient and efficient way for IT organization to determine resource optimization for computing systems.

83. This notion was explained during the prosecution of the underlying patent application: “the technical problem is solved by the present invention by computing efficiency scores for computing entities based on resource utilization data, resource capacity data, and operational policies relating to resource allocation. These scores thus reflect conditions prevailing in an apparatus or system – in particular conditions as to the utilization and availability of resources (such as CPU/memory capacity) – and these conditions are visually indicated by generating a graphical display with an indicator for at least one computing entity.”⁵³ As such, the ’367 patent is directed to a novel way to measure, evaluate, and visualize efficiencies and risks in an IT

⁵⁰ ’367 patent, 1:32-39.

⁵¹ ’367 patent, 1:49-52.

⁵² ’367 patent, 11:25-27.

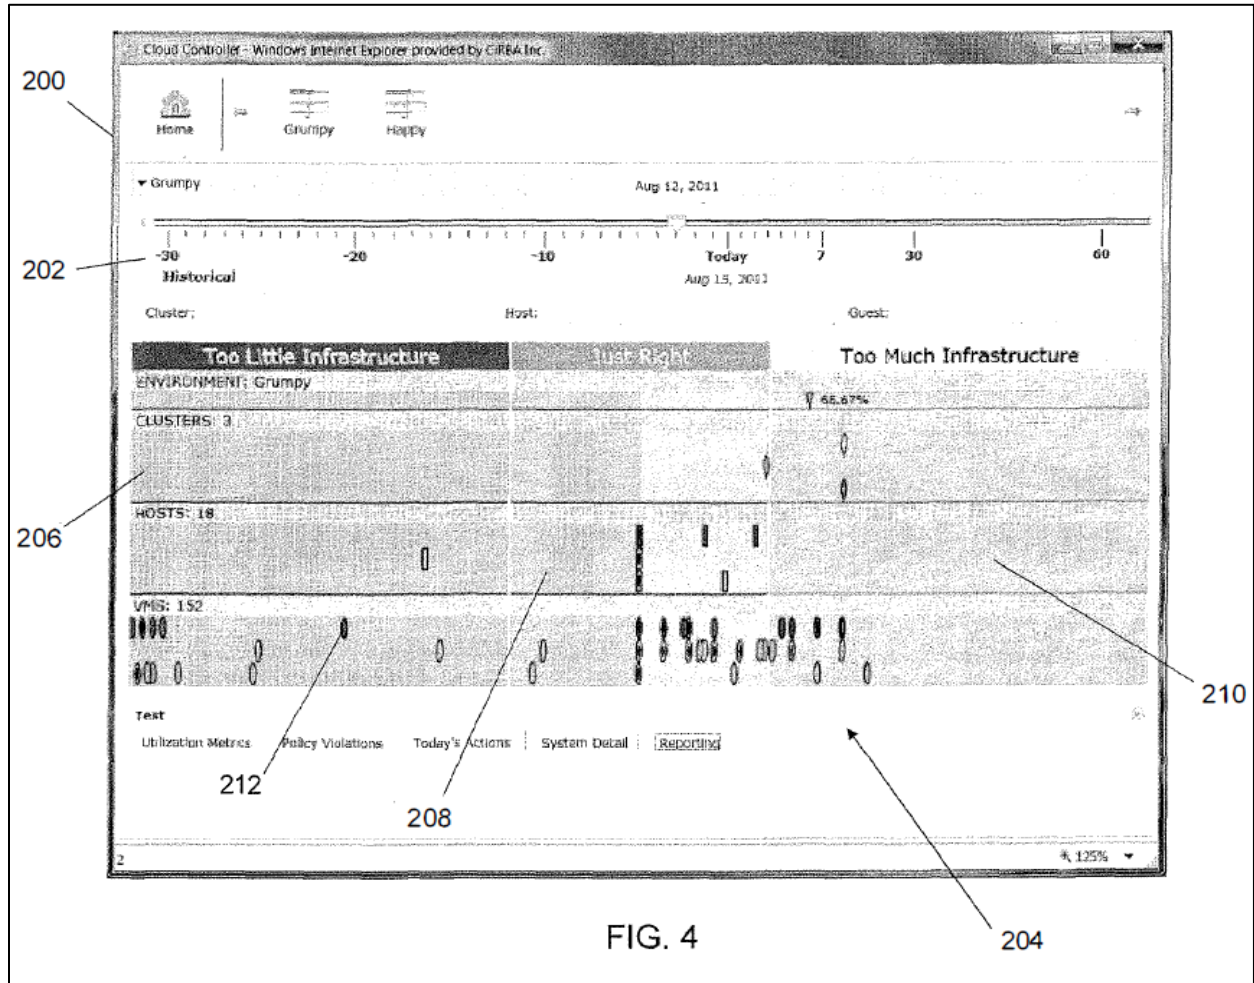
⁵³ U.S. Patent Application No. 14/180,438, Apr. 22, 2016 Response to Office Action (Ex. 17).

infrastructure, resulting in a specific improvement over prior systems that is not well-known or conventional.

84. Claim 1 of the '367 patent is directed to a “method performed by a processor in a computing system.” vROps is performed by a processor in a computing system. Claim 1 further recites “obtaining resource utilization or performance data pertaining to a plurality of computing entities in a computing environment, and capacity data specifying resource capacities for the plurality of computing entities in the computing environment” and “obtaining at least one operational policy defining criteria to determine whether the utilization or performance of an entity is in an acceptable range relative to its capacity or performance limits.” Claim 1’s method also recites “computing at least one score quantifying efficiencies and/or risks associated with corresponding ones of the entities in the computing environment, based on the resource utilization or performance data, the capacity data, and the at least one operational policy.”

85. Claim 1 further involves “displaying an indicator for at least one of the plurality of computing entities in a graphical representation based on the corresponding score,” “wherein each indicator is positioned in the graphical representation according to the corresponding score such that the positioned indicator shows in a spatial manner, relative efficiencies and/or risks for the corresponding entity by positioning the indicator in one of a first portion indicative of risk associated with having infrastructure in the computing environment that cannot service workload demands and meet criteria specified in the at least one operational policy, a second portion indicative of an amount of infrastructure in the computing environment that can service workload demands based on the at least one operational policy, or a third portion indicative of inefficiencies associated with having more than the required amount of infrastructure in the computing environment to service workload demands based on the at least one operational policy.”

86. Figure 4 of the '367 patent is instructive because it is representative of an embodiment of the claimed invention and because vROps has an identical feature as depicted below:



'367 patent, Figure 4



87. As illustrated above⁵⁴ and explained in detail below, the VMware Accused Products directly infringe at least claim 1 of the '367 patent.

88. vROps obtains resource utilization or performance data pertaining to a plurality of computing entities in a computing environment, and capacity data specifying resource capacities for the plurality of computing entities in the computing environment. For example, vROps' "analytics provide precise tracking, measuring and forecasting of data center capacity, usage, and trends to help manage and optimize resource use, system tuning, and cost recovery. The system monitors stress thresholds and alerts you before potential issues can affect performance. Multiple pre-set reports are available. You can plan capacity based on historical usage and run what-if scenarios as your requirements expand."⁵⁵ vROps also "use[s] the Capacity Optimization and

⁵⁴ Overbeek, D., vRealize Operations 6.2: Intelligent Workload Placement with DRS, VMware Blogs, available at <https://blogs.vmware.com/management/2016/02/vrealize-operations-6-2-intelligent-workload-placement-with-drs-2.html> (Ex. 18).

⁵⁵ *Supra* n. 8 (vRealize Operations Manager 7.0 Help Guide at 770, Ex. 5).

Reclaim features to assess workload status and resource contention in data centers across your environment. You can determine time remaining until cpu, memory, or storage resources run out and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.”⁵⁶

89. vROps likewise obtains performance data from monitored software and hardware resources in your enterprise and provides predictive analysis and real-time information about problems.⁵⁷ It “monitors your ESXi hosts and the virtual machines located on them,” and “monitors virtual machines running in a vCenter Server, analyzes longer-term historical data, and provides forecast data about predictable patterns of resource usage to Predictive DRS. Based on these predictable patterns, Predictive DRS moves to balance resource usage among virtual machines.”⁵⁸

90. Next, vROps obtains at least one operational policy defining criteria to determine whether the utilization or performance of an entity is in an acceptable range relative to its capacity or performance limits. For example, in vROps “[a]dministrators assign policies to object groups and applications to support Service Level Agreements (SLAs) and business priorities. When you use policies with object groups, you ensure that the rules defined in the policies are quickly put into effect for the objects in your environment.”⁵⁹ As a further example, vROps “[u]se host-based VM placement to tie your VMs more closely to your infrastructure. By using vCenter Server to tag hosts and VM with specific tags, you make certain that when the system runs an optimization,

⁵⁶ *Supra* n. 8 (vRealize Operations Manager 7.0 Help Guide at 770, Ex. 5).

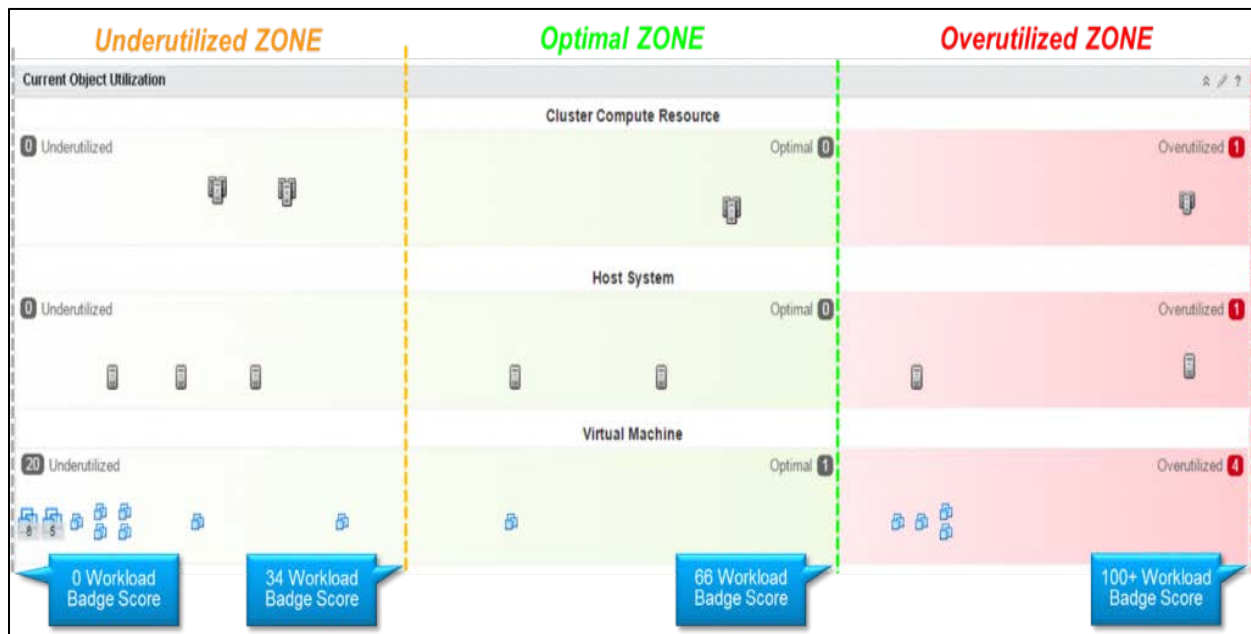
⁵⁷ *Id.* (vRealize Operations Manager 7.0 Help Guide at 486, Ex. 5).

⁵⁸ *Id.* (vRealize Operations Manager 7.0 Help Guide at 120, Ex. 5).

⁵⁹ *Id.* (vRealize Operations Manager 7.0 Help Guide at 304, Ex. 5).

it used VM-to-host tag matching to ensure that VMs are moved to – or -stay with – the appropriate host.”⁶⁰

91. Further, vROps computes at least one score quantifying efficiencies and/or risks associated with corresponding ones of the entities in the computing environment, based on the resource utilization or performance data, the capacity data, and the at least one operational policy. For example, vROps provides “the efficiency widget that is the status of the efficiency-related alerts for the objects it is configured to monitor. Efficiency alerts in vRealize Operations Manager usually indicate that you can reclaim resources. You can create one or more efficiency widgets for objects that you add to your custom dashboards.”⁶¹ In addition, scores “quantifying efficiencies and/or risks” can be observed in the vROps dashboard below.⁶²



⁶⁰ *Id.* (vRealize Operations Manager 7.0 Help Guide at 630, Ex. 5).

⁶¹ *Id.* (vRealize Operations Manager 7.0 Help Guide at 408, Ex. 5).

⁶² *Supra* n. 55 (Overbeek, vRealize Operations 6.2: Intelligent Workload Placement with DRS, Ex. 18).

92. Lastly, vROps displays an indicator for at least one of the plurality of computing entities in a graphical representation based on the corresponding score, wherein each indicator is positioned in the graphical representation according to the corresponding score such that the positioned indicator shows in a spatial manner, relative efficiencies and/or risks for the corresponding entity by positioning the indicator in one of a first portion indicative of risk associated with having infrastructure in the computing environment that cannot service workload demands and meet criteria specified in the at least one operational policy, a second portion indicative of an amount of infrastructure in the computing environment that can service workload demands based on the at least one operational policy, or a third portion indicative of inefficiencies associated with having more than the required amount of infrastructure in the computing environment to service workload demands based on the at least one operational policy.”

93. For example, the “first portion” is depicted in the illustration of vROps above by “overutilized ZONE,” the “second portion” is depicted by “optimal ZONE,” and the “third portion” is depicted the “underutilized ZONE.”

94. Based on the above, the VMware Accused Products directly infringe at least, but not limited to, claim 1 of the '367 patent.

F. VMware's Infringement Has Injured Densify

95. VMware's use of infringing technology has increased over the years, culminating in its recent release of vROps 7.0 and now the new release of vROps 7.5. And VMware shows no signs of stopping. VMware has even more recently indicated its intent to move into the native and hybrid cloud environments. In its marketing materials, VMware is touting enhancements related to cloud migration as one of the biggest feature enhancements in the recent vROps 7.5 release. And, in these materials, VMware emphasizes that vROps can manage workloads in VMware

Cloud on AWS. With its launch of vROps 7.0 and vROps 7.5, VMware's products are now workable and attractive to these customers because they incorporate Densify's patented technology.

96. Upon information and belief, VMware is offering steep discounts (e.g., single, one-time, and bundled discounts) to subscribe to vROps 7.0. By doing so, VMware is foreclosing the infrastructure optimization market. Although Densify has been competitive with VMware, customers have used Densify alongside VMware's management tools like DRS and vROps. On information and belief, VMware is telling customers that there is no need for them to have Densify now that VMware has its new functionality – functionality that infringes Densify's patents. As such, Densify cannot compete effectively when its patents are infringed. Once customers are lost and market share diminished, the barrier to entry would be insurmountable for Densify to compete again in the same market space.

97. Accordingly, Densify has brought this action to stop VMware from using copied technology to take customers from Densify.

CLAIM I

VMWARE'S INFRINGEMENT OF THE '687 PATENT

98. Densify incorporates by reference the foregoing paragraphs of this Complaint as if fully set forth herein.

99. VMware has infringed and continues to infringe the '687 Patent. VMware directly infringes the '687 patent under 35 U.S.C. §271(a) by making, using, selling, offering for sale, and/or importing in this District and into the United States products and/or methods covered by one or more claims of the '687 patent, including, but not limited to, the VMware Accused Products. As an example, the VMware Accused Products infringe at least claim 7 of the '687 patent.

100. VMware also indirectly infringes the '687 patent by inducing others to infringe and/or contributing to the infringement of others, including third party users of the VMware Accused Products in this District and elsewhere in the United States.

101. Specifically, on information and belief, VMware has had knowledge of its infringement of the '687 patent since at least the time of service of this Complaint and before the time of service of the Complaint, including for the reasons that, for example: Densify's products were well-known and publicized; Densify has received press coverage (including at VMworld) of its pioneering technological developments in virtual infrastructure optimization; Densify has marketed its products with notices regarding the '687 patent; former Densify customer representatives and industry analysts that were intimately familiar with Densify and its products are now at VMware; and Densify and VMware are competitors in the same industry, giving VMware reason to investigate Densify's patents on the technologies. To the extent that VMware lacked knowledge of its infringement of the '687 patent before the time of service of this Complaint, it remained willfully blind by affirmatively avoiding investigating Densify's patents or inspecting Densify's website.

102. Upon information and belief, Densify alleges that VMware has actively induced the infringement of the '687 patent under 35 U.S.C. § 271(b) by actively inducing the infringing use of the VMware Accused Products by third parties in the United States. Densify is informed and believes, and thereon alleges, that these third parties infringe the '687 patent in violation of 35 U.S.C. § 271(a) by using the VMware Accused Products.

103. VMware actively induces others, including customers, administrators, and operators, to infringe by, among other activities, providing instructions, training, and support for

the use of the VMware Accused Products through websites, technical documents and manuals, tutorials, and support services.

104. For example, VMware provides websites and blogs that instructs customers or other third parties how to use the VMware Accused Products (e.g., <https://blogs.vmware.com/management/2018/09/self-driving-all-the-way-to-the-host-oh-yeah-host-based-placement.html>, Ex. 19).

105. As another example, VMware provides customers and third parties with technical documents, like the vRealize Operations Manager Help Guide (e.g., <https://docs.vmware.com/en/vRealize-Operations-Manager/7.0/vrealize-operations-manager-70-help.pdf>, Ex. 5), that teach customers, administrators, and other third parties how to use the VMware Accused Products in an infringing manner.

106. Further, VMware hosts industry events (e.g., VMworld), local seminars, as well as live and on-demand webcasts and webinars to teach customers and third parties how to use the VMware Accused Products.⁶³ Specifically, such events provide for sessions that teach and show step-by-step how the customers can use, configure, manage, operate, etc. their VMware Accused Products. Accordingly, VMware actively induces third parties to infringe the '687 patent.

107. Upon information and belief, VMware contributorily infringes the '687 patent under 35 U.S.C. § 271(c) by importing, selling and/or offering to sell within the United States the VMware Accused Products (or components thereof) that constitute a material part of the claimed invention and are not staple articles of commerce suitable for substantial non-infringing use. For

⁶³ E.g., <https://blogs.vmware.com/management/2018/08/vmworld-las-vegas-recommended-self-driving-operations-sessions.html> (Ex. 20); Screenshot from HOL-1971-01-CMP - What's New in vRealize Operations Manager 7.0, VMware Hands-on Labs - Cloud Management Platform, available at <https://labs.hol.vmware.com/HOL/catalogs/catalog/873> (Ex. 21).

example, when VMware provides vROps to its customers or other third parties, as part of a suite of software products, vROps constitutes a material component of the suite of software products that infringe or vROps is a material component used in practicing the '687 patent. vROps is especially adapted for use in an infringing manner without substantial non-infringing uses. And VMware has provided components of vROps, which constitute a material part of the claimed invention and are not staple articles of commerce suitable for substantial non-infringing use.

108. In addition, VMware offers to sell and sells the VMware Accused Products to resellers who then incorporate the VMware Accused Products into infringing products which are used, sold, offered for sale, and or/or imported in the United States in an infringing manner. Accordingly, VMware contributorily infringes the '687 patent.

109. Densify has no adequate remedy at law against VMware's acts of infringement, and unless VMware is enjoined from its infringement of the '687 patent, Densify will suffer irreparable harm.

110. VMware, by way of its infringing activities, has caused and continues to cause Densify to suffer damages, the exact amount to be determined at trial. The damages include lost profits and/or reasonable royalty damages based on VMware's infringement. Densify's products have been marked with Densify's patents since at least or around March 2015.

111. Densify is further informed, and on this basis alleges, that VMware's infringement of the '687 patent has been and continues to be deliberate and willful, and, therefore, this is an exceptional case warranting an award of enhanced damages for up to three times the actual damages awarded and attorney's fees to Densify pursuant to 35 U.S.C. §§ 284-285. As noted above, VMware has had knowledge of the '687 patent and its infringement thereof, and yet has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless

disregard for Densify's patent rights. Thus, VMware's infringing activities have been and continue to be consciously wrongful.

CLAIM II

VMWARE'S INFRINGEMENT OF THE '367 PATENT

112. Densify incorporates by reference the foregoing paragraphs of this Complaint as if fully set forth herein.

113. VMware has infringed and continues to infringe the '367 Patent. VMware directly infringes the '367 patent under 35 U.S.C. §271(a) by making, using, selling, offering for sale, and/or importing in this District and into the United States products and/or methods covered by one or more claims of the '367 patent, including, but not limited to, the VMware Accused Products. As an example, the VMware Accused Products infringe at least claim 1 of the '367 patent.

114. VMware also indirectly infringes the '367 patent by inducing others to infringe and/or contributing to the infringement of others, including third party users of the VMware Accused Products in this District and elsewhere in the United States.

115. Specifically, on information and belief, VMware has had knowledge of its infringement of the '367 patent since at least the time of service of this Complaint and before the time of service of the Complaint, including for the reasons that, for example: Densify's products were well-known and publicized; Densify has received press coverage (including at VMworld) of its pioneering technological developments in virtual infrastructure optimization; Densify has marketed its products with notices regarding the '367 patent; former Densify customer representatives and industry analysts that were intimately familiar with Densify and its products are now at VMware; and Densify and VMware are competitors in the same industry, giving VMware reason to investigate Densify's patents on the technologies. To the extent that VMware lacked knowledge of its infringement of the '367 patent before the time of service of this

Complaint, it remained willfully blind by affirmatively avoiding investigating Densify's patents or inspecting Densify's website.

116. Upon information and belief, Densify alleges that VMware has actively induced the infringement of the '367 patent under 35 U.S.C. § 271(b) by actively inducing the infringing use of the VMware Accused Products by third party users in the United States. VMware actively induces others, including customers, administrators, and operators, to infringe by, among other activities, providing instructions, training, and support for the use of the VMware Accused Products through websites, technical documents and manuals, tutorials, and support services. Acts of inducement are set forth in Paragraphs 98-102.

117. Upon information and belief, VMware contributorily infringes the '367 patent under 35 U.S.C. § 271(c) by importing, selling and/or offering to sell within the United States the VMware Accused Products (or components thereof) that constitute a material part of the claimed invention and are not staple articles of commerce suitable for substantial non-infringing use. For example, VMware has provided components of vROps, which constitute a material part of the claimed invention and are not staple articles of commerce suitable for substantial non-infringing use. Accordingly, VMware contributorily infringes the '367 patent.

118. In addition, VMware offers to sell and sells the Accused Products to resellers who then incorporate the VMware Accused Products into infringing products which are used, sold, offered for sale, and or/or imported in the United States in an infringing manner. Accordingly, VMware contributorily infringes the '367 patent.

119. Densify has no adequate remedy at law against VMware's acts of infringement, and unless, VMware is enjoined from its infringement of the '367 patent, Densify will suffer irreparable harm.

120. VMware, by way of its infringing activities, has caused and continues to cause Densify to suffer damages, the exact amount to be determined at trial. The damages include lost profits and/or reasonable royalty damages based on VMware's infringement. Densify's products have been marked with Densify's patents since at least or around March 2015.

121. Densify is further informed, and on this basis alleges, that VMware's infringement of the '367 patent has been and continues to be deliberate and willful, and, therefore, this is an exceptional case warranting an award of enhanced damages for up to three times the actual damages awarded and attorney's fees to Densify pursuant to 35 U.S.C. §§ 284-285. As noted above, VMware has had knowledge of the '367 patent and its infringement thereof, and yet has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for Densify's patent rights. Thus, VMware's infringing activities have been and continue to be consciously wrongful.

CLAIM III

UNFAIR COMPETITION IN VIOLATION OF THE LANHAM ACT, 15 U.S.C. § 1125(a)

122. Densify incorporates by reference the foregoing paragraphs of this Complaint as if fully set forth herein.

123. Densify owns all right, title, and interest in and to the trademark and service marks including "DENSIFY," "DENSIFICATION," and "DENSIFYING" (collectively, the "Densify Marks"). Densify has used the marks in connection in interstate commerce in connection with its goods and services since at least as early as 2015.

124. Densify has invested and continues to invest substantial resources into promoting the Densify Marks. As a result, the Densify Marks have come to be associated exclusively with a single source, specifically, Densify's products and services. The Densify Marks have achieved

secondary meaning. They embody the substantial goodwill Densify has built in its company and its well-regarded products. The Densify Marks are distinctive.

125. The below picture shows Densify's use of its mark on the exterior of its offices in Markham, Ontario.⁶⁴



126. The below picture shows Densify's use of its mark when it sponsored a NASCAR at the Daytona 500.⁶⁵ Other examples of Densify's use of the Densify Marks are included as Exhibits 24 – 25.⁶⁶

⁶⁴ Photograph of Densify's exterior offices in Markham, Ontario (Ex. 22).

⁶⁵ Twitter, @Densify, available at <https://twitter.com/Densify/status/958020795648892934> (Ex. 23).

⁶⁶ Densify, Densify Brand Resources, available at <https://www.densify.com/company/brand> (Ex. 24); Densify, Dedicated Cloud & Infrastructure Experts at Your Service, available at <https://www.densify.com/service/expert-insight> (Ex. 25).



127. Upon information and belief, VMware was aware of Densify's use of the Densify Marks.

128. On May 12, 2016, Densify launched its website, <http://www.densify.com>. The launch of the website was accompanied by a press release.

129. On June 28, 2017, Densify launched a branding project under the name "Densify." This rebranding included a press release as well as a substantial advertising campaign. These materials were sent to customers and potential customers in Delaware.

130. On August 14, 2018, Densify filed an application for trademark protection with the United States Patent and Trademark Office, with the serial number 88078131, for the mark "DENSIFY." That application is currently pending.

131. Upon information and belief, VMware has regularly and continuously infringed upon Densify's marks in connection with the sale, offering for sale, distribution, or advertising of its products and services, words and symbols that infringe upon Densify's trademarks and service marks.

132. Upon information and belief, since at least 2018, VMware has infringed upon Densify's marks by using them in its sales materials, videos, presentations, and documentation. Examples of VMware's infringing use are included as Exhibits 14 and 26 – 30.⁶⁷

133. Upon information and belief, VMware uses the marks to increase the commercial value of its products and services by creating an association with Densify's award-winning optimization products and services. For example, VMware promotes its "Top Reasons to Upgrade vROps 7.0 and How to Get the Most Out of It" by using "densification" as shown in the red box in the figure below.⁶⁸

⁶⁷ Jones, M., Upgrade to Self-Driving Operations at up to 65% off!, VMware Blogs, available at <https://blogs.vmware.com/management/2018/10/upgrade-to-self-driving-operations-at-up-to-65-off.html> (Ex. 26 at 2 (highlighting added));

Supra n. 34 (Self-Driving Operations by VMware vRealize Operations, Datasheet at 1 (Ex. 14 (highlighting added)));

Overbeek, D., Cloud Management Platform (CMP) – Intelligent Provisioning and Optimization, VMware Blogs, available at <https://blogs.vmware.com/management/2018/11/cloud-management-platform-cmp-intelligent-provisioning-and-optimization.html> (Ex. 27 at 1 (highlighting added));

Overbeek, D., Start Running a Self-Driving Datacenter – vRealize Operations 7.0 Workload Optimization!, VMware Blogs, available at <https://blogs.vmware.com/management/2018/09/start-running-a-self-driving-datacenter-vrealize-operations-7-0-workload-optimization.html> (Ex. 28 at 2 (highlighting added));

Screenshot from VMware EMEA, Top Reasons to Upgrade to vROps 7.0 and How to Get the Most Out of It, available at <https://www.youtube.com/watch?v=1KRRuntV4aM> (Ex. 29 (highlighting added)). A full version of the video from which the screenshot (Ex. 29) was taken has been provided as Ex. 36;

Screenshot from Workload Optimization - Densifying to Repurpose Hosts, VMware Cloud Management, available at <https://www.youtube.com/watch?v=IymKQdg3oNE> (Ex. 30). A full version of the video from which the screenshot (Ex. 30) was taken has been provided as Ex. 34.

⁶⁸ *Supra* n. 68 (Screenshot from VMware EMEA, Ex. 29 (highlighting added)).

The Three Tenets of Self-Driving Operations
 Run Production Operations Hands-off and Hassle-free

- Continuous Performance Optimization** – Assure application performance based on business and operational intent
- Efficient Capacity Management** – Run infrastructure like AWS – optimal densification, proactive planning and procurement
- Intelligent Remediation** – Predict, prevent and troubleshoot across SDDC and multiple clouds

Optimized Hybrid Cloud
 Cut 40% trouble tickets to drive in the field
 30% Capacity Savings
 10% reduction in hardware costs
 75% reduction in unplanned downtime

vmware Confidential | ©2018 VMware, Inc. POWERED BY INXPO

0:02:07 0:54:53

134. Upon information and belief, VMware’s use of Densify’s marks is likely to create confusion. Both VMware and Densify sell services within the same market and most of Densify’s customers are also VMware users.

135. Upon information and belief, VMware has customers within the state of Delaware, and regularly solicits further business within the state of Delaware. Upon information and belief, it uses materials that infringe upon Densify’s trade and service marks within the state of Delaware and in interstate commerce.

136. Upon information and belief, Densify has been harmed by, and will continue to be harmed by, VMware’s infringing conduct. VMware’s actions have caused monetary damages to Densify, by creating confusion in the minds of consumers as to what products and services they were being sold, and by depriving Densify of the full value of the Densify Marks. VMware’s use of the Densify Marks has damaged the goodwill associated with those marks.

137. Upon information and belief, VMware's infringing conduct causes Densify irreparable harm, including, but not limited to, depriving Densify of its rights in the Densify Marks, creating a likelihood of confusion, mistake, or deception within its industry, causing a false association in the minds of its consumers between Densify and VMware, and incalculable loss of goodwill.

138. VMware's use of the Densify Marks violates the Lanham Act, 15 U.S.C. § 1125(a).

139. VMware's violation of § 1125(a) was deliberate, willful, and intended to cause confusion, mistake, or deception in violation of 15 U.S.C. § 1125.

140. As a result of VMware's acts, Densify is entitled to recover VMware's profits, its own damages, the costs of this action, and enhanced damages.

CLAIM IV

DECEPTIVE TRADE PRACTICES UNDER DELAWARE LAW

141. Densify incorporates by reference the foregoing paragraphs of this Complaint as if fully set forth herein.

142. Densify has common law rights in the state of Delaware and elsewhere to the Densify Marks.

143. VMware has used the Densify Marks or facsimiles thereof to create a likelihood of customer confusion that VMware's products use Densify's products and services, are associated with Densify or its products or services, or are substantially the same as Densify's products and services.

144. VMware has passed off its goods and services as those of Densify or otherwise caused a likelihood of confusion or misunderstanding as to the "the source, sponsorship, approval .

. . . [or] affiliation, connection, association with, or certification by” between its products and services and Densify’s products and services in violation of Delaware Code Title 6 § 2532(1)-(3).

145. VMware further “[r]epresents that [its] goods or services have sponsorship, approval, characteristics . . . that they do not have,” namely the sponsorship, approval, or characteristics of Densify and its products or services, in violation of Delaware Code Title 6 § 2532(5).

146. These infringements by VMware add to the commercial value of its products and services.

147. Densify has suffered monetary harm as a result of VMware’s conduct.

148. VMware’s conduct, in both its violations of Delaware’s trademark protections and the federal Lanham Act, constitutes a “deceptive trade practice” within the meaning of Delaware Code Title 6 § 2532.

149. VMware’s acts constitute an exceptional case and have been committed willfully. Accordingly, Densify is entitled to reasonable attorneys’ fees under Delaware Code Title 6 § 2533(b).

150. Densify is entitled to damages and injunctive relief under Delaware law. Both Densify and the public will suffer irreparable harm if VMware is permitted to continue its infringement. Therefore, Densify is entitled to injunctive relief that requires VMware to cease use of any and all Densify trade and service marks. VMware’s willful actions constitute an exceptional case, and Densify is therefore entitled to its reasonable attorneys’ fees and such other relief as is provided by law.

CLAIM V

COMMON LAW TRADEMARK INFRINGEMENT

151. Densify incorporates by reference the foregoing paragraphs of this Complaint as if fully set forth herein.

152. Densify has valid and protectable common law rights in the Densify Marks.

153. Densify is the senior user of the Densify Marks.

154. VMware's conduct constituted infringement of Densify's common law rights in the Densify Marks.

155. VMware's use of the Densify Marks on unauthorized products and services is likely to cause confusion as to source, sponsorship, affiliation, or origin of VMware's products and services.

156. VMware's use of the Densify Marks is likely to deceive others as to the relationship between VMware and Densify.

157. VMware's wrongful acts of infringement have permitted, continue to permit, and will permit VMware to earn substantial profits on the basis of Densify's reputation and goodwill embodied in the marks.

158. As a direct and proximate result of VMware's wrongful acts of infringement, Densify has been, continues to be, and will be harmed.

159. Densify is entitled to damages and injunctive relief under Delaware law. Both Densify and the public will suffer irreparable harm if VMware is permitted to continue its infringement. Therefore, Densify is entitled to injunctive relief that requires VMware to cease use of the Densify Marks.

PRAYER FOR RELIEF

160. WHEREFORE, Densify prays for judgment as follows:

a) That VMware has infringed, and unless enjoined will continue to infringe, each of the Asserted Patents;

b) That VMware has willfully infringed each of the Asserted Patents;

c) That VMware, its officers, agents, servants, and employees, and those persons in active concert or participation with any of them, be preliminarily, and thereafter permanently, enjoined from commercially manufacturing, using, offering for sale, selling, or importing the VMware Accused Products or any other product that infringes, or induces or contributes to the infringement of the Asserted Patents, prior to the expiration date of the last to expire of those patents;

d) That Densify be awarded monetary relief sufficient to compensate Densify for damages resulting from VMware's infringement of the Asserted Patents, including lost profits and/or a reasonable royalty under 35 U.S.C. § 284, and that such monetary relief be awarded to Densify with prejudgment and post-judgment interest;

e) That Densify be awarded enhanced damages, up to and including trebling of the damages awarded to Densify;

f) That Densify be awarded the attorneys' fees, costs, and expenses that it incurs prosecuting this action under 35 U.S.C. § 285;

g) Damages, VMware's profits, and the costs of this action under 15 U.S.C. § 1117;

h) Injunctive relief that requires VMware to cease use of the Densify Marks; and

i) That Densify be awarded such other and further relief as this Court deems just and proper.

DEMAND FOR JURY TRIAL

161. Densify demands a trial by jury of any and all issues triable of right before a jury.

Dated: June 18, 2019

Respectfully submitted,

/s/ Kenneth L. Dorsney

Sarah O. Jorgensen
(*pro hac vice*)
sjorgensen@reichmanjorgensen.com
Reichman Jorgensen LLP
1201 West Peachtree, Suite 2300
Atlanta, GA 30309
Telephone: (650) 623-1403
Telecopier: (650) 623-1449

Christine E. Lehman
(*pro hac vice*)
clehman@reichmanjorgensen.com
Reichman Jorgensen LLP
1615 M Street, NW, Suite 300
Washington, DC 20036
Telephone: (202) 894-7311
Telecopier: (650) 623-1449

Jaime Cardenas-Navia
(*pro hac vice*)
jcardenas-navia@reichmanjorgensen.com
Wesley L. White
(*pro hac vice*)
wwhite@reichmanjorgensen.com
Reichman Jorgensen LLP
100 Park Avenue, Suite 1600
New York, NY 10017
Telephone: (646) 921-1474
Telecopier: (650) 623-1449

Kenneth L. Dorsney (#3726)
kdorsney@morrisjames.com
Morris James LLP
500 Delaware Avenue, Suite 1500
Wilmington, DE 19801
Telephone: (302) 888-6800

Courtland L. Reichman (*pro hac vice*)
creichman@reichmanjorgensen.com
Shawna Ballard (*pro hac vice*)
sballard@reichmanjorgensen.com
Jennifer P. Estremera (*pro hac vice*)
jestremera@reichmanjorgensen.com
Phillip Lee (*pro hac vice*)
plee@reichmanjorgensen.com
Michael G. Flanigan (*pro hac vice*)
mflanigan@reichmanjorgensen.com
Joachim B. Steinberg (*pro hac vice*)
jsteinberg@reichmanjorgensen.com
Kate Falkenstien (*pro hac vice*)
kfalkenstien@reichmanjorgensen.com
Reichman Jorgensen LLP
303 Twin Dolphin Drive, Suite 600
Redwood Shores, CA 94065

*Attorneys for Plaintiffs
Cirba Inc. (d/b/a Densify) and Cirba IP, Inc.*

Exhibit 1

densify
Case 1:19-cv-00742-LPS Document 68 Filed 06/18/19 Page 56 of 288 PageID #: 4230

ANALYSIS FINDINGS

- Executive Summary
- Public Cloud
- Virtual Infrastructure
- My Reports

SUPPLY & DEMAND PLANNING

- Transform to Cloud
- Route and Reserve Demand
- Plan and Manage Supply
- Forecast

AUTOMATION

- Real-Time Automation
- Proactive Automation

ADVANCED CONTROLS

- Control Policies
- Control Console
- Configure My Reports

DENSIFICATION ADVISOR



7 Insight(s)

Clusters with Insufficient Capacity

4 of 16

3 are **OPEN** Clusters
1 is a **CLOSED** Cluster

Top Clusters with Insufficient Capacity

Hosts at Risk

17 of 95

Impact of Rebalancing
17 to 1 (94.1% Reduction)

Undersized Guests

120 of 1,015

Impact of Right Sizing
120 to 4 (96.7% Reduction)

Clusters with Undersized Guests

Cluster, host and guest-level risks

Environment	Cluster	Available Space	Environment	Cluster	Hosts at Risk	%	Environment	Cluster	Undersized Guests	%
Houston	HV-PRD-App1	23	New York	General Apps	4	5.1	London	Internal Apps	13	28.6
New York	General Apps	16	London	BC-Apps	1	10	San Francisco	Zone 3	3	29.8
New York	Production Apps2	6	London	BA-Apps	1	50	Seattle	ProdCloud-Apps1	5	17.3
New York	Production Apps1	6	London	Internal Apps	1	32.9	Houston	HV-PRD-App2	5	11.8
			Seattle	ProdCloud-Apps1	2	75	New York	General Apps	6	16.2


Exactly where the risks are and how to fix them





Exhibit 2

Operational Risk

Stranded Capacity



**Before
Densify**



**After
Densify**

Exhibit 3



ANALYSIS FINDINGS

- Executive Summary
- Public Cloud
- Virtual Infrastructure
- My Reports

SUPPLY & DEMAND PLANNING

- Transform to Cloud
- Route and Reserve Demand
- Plan and Manage Supply
- Forecast

AUTOMATION

- Real-Time Automation
- Proactive Automation

ADVANCED CONTROLS

- Control Policies
- Control Console
- Configure My Reports

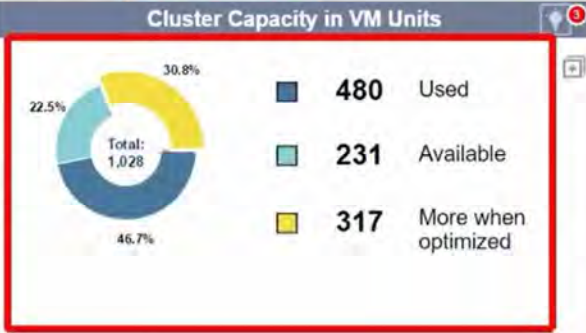
DENSIFICATION ADVISOR



7 Insight(s)



Virtual Infrastructure for **Total Estate** [Set Filter](#)

Risk Analysis **Densification Opportunity**




How to free up stranded capacity

San Francisco Zone 2 **24**



How to maximize VM density

San Francisco Zone 3 **179**



How to minimize software costs

San Francisco Zone 3 **179**




Exhibit 4

What the history of VMware reveals about its future projects

VMware dominated server virtualization, but was hesitant to embrace the cloud.

Examine VMware's history and its approach to new technology to see where it's headed next.

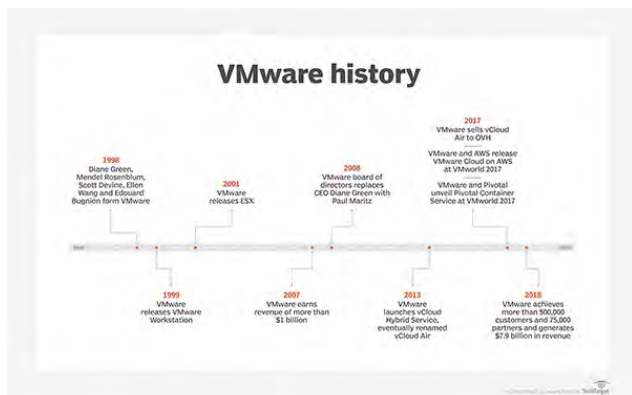
The history of VMware presents a company that disrupted the IT industry with server virtualization, but that was hesitant to embrace new technologies, such as the public cloud. VMware's future success will depend on how it finds a role in the current cloud market and how it approaches containers.

This year marks VMware's 20th year anniversary, a significant accomplishment in the volatile high-tech marketplace. The firm's virtualization software revolutionized the data center and elevated VMware to a lofty position in the market. In a March 2018 press release, VMware publicized that it had attained more than 500,000 customers and 75,000 partners and generated \$7.9 billion in revenue so far in 2018.

In the spring of 1998, Diane Greene, Mendel Rosenblum, Scott Devine, Ellen Wang and Edouard Bugnion formed VMware. The impending [Y2K](#) crisis loomed and the dot-com boom was building hype, so few paid attention to the fledgling business.

The company's first product -- [VMware Workstation](#), which offered VM software for Intel x86 computers, is still relevant today, and its release marks a prescient point in the history of VMware.

"VMware was an innovator," said Gary Chen, research manager of software-defined compute at International Data Corp.



VMware was one of the first to market with server virtualization and offered ESX, a well-designed product with numerous advantages over traditional systems.

"The product saved IT lots of money and made data center technicians' lives easier, so it is still the foundation for many enterprise data centers today," Chen said.



Gary Chen

The concept of virtualization has moved from the server to other areas in the [data center infrastructure](#). The history of VMware shows a company trying to keep up with its original innovation.

VMware parlayed the resulting initiatives into becoming a global IT leader, [earning revenues](#) of more than a billion dollars in 2007. This success ensured that the management team largely stayed in place until the summer of 2008.

After a [disappointing financial performance](#) due to mismanagement of the company's growing scale, the board of directors replaced VMware president and CEO Diane Greene with Paul Maritz, a former Microsoft veteran who headed EMC's cloud computing business unit. This leadership reorganization marked a point in the history of VMware that foreshadowed a shift in the company's direction.

VMware history fails to prepare it for cloud

VMware's foundation is in the on-premises data center, so it was slow to embrace public cloud. Established vendors often feared that offering cloud services posed risks to their existing businesses.

"VMware was part of that thinking, along with many other vendors," Chen said.

In May 2013, VMware launched an infrastructure-as-a-service product. VMware initially dubbed it vCloud Hybrid Service, but it eventually became known as vCloud Air. Interest was tepid, and in the second quarter of 2017, VMware [sold the service](#) to OVH, a French cloud supplier.

The firm eventually bet on an alliance with Amazon Web Services (AWS). The two vendors developed [VMware Cloud on AWS](#) as a service to help businesses migrate applications from on-premises VMware clusters to AWS. Compatibility is a big draw with the service.

VMware maintains its position with data center technicians, but it's not growing with business unit developers who increasingly control significant portions of IT spending.

"If a firm has VMware now in its data center, they see an easy way to move that workload to the public cloud," Chen said.

VMware Cloud on AWS isn't as attractive for organizations developing new applications. VMware maintains its position with data center technicians, but it's not growing with business unit developers who increasingly control significant portions of IT spending.

VMware develops container strategy

Containers are changing the way companies design and support enterprise applications. In March 2013, VMware spun off Pivotal Software, which had been at the forefront of the container movement.

Recently the two corporations [came back together](#). At VMworld 2017, VMware and Pivotal unveiled Pivotal Container Service (PKS). PKS enables organizations to use Kubernetes on VMware vSphere and Google Cloud Platform to manage container applications.

Like the AWS agreement, the service has advantages and disadvantages.

"The purist may want a direct connection from a host operating system and its libraries direct to containers. With VMware/Pivotal, a hypervisor sits in between the two," said Marco Alcalá, CEO at Alcalá Consulting.

VMware's virtualization software reshaped the computer industry. Cloud and containers are now having a similar effect. VMware's [response to these trends](#) will determine its impact in the coming years.

[Prev](#)

1. [1](#)

[Next](#)



[Paul Korzeniowski](#) asks:

How do you think VMware will fare in the container market?

[Join the Discussion](#)

[Prev](#)

Exhibit 5

vRealize Operations Manager 7.0 Help

15 NOV 2018

vRealize Operations Manager 7.0



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2018 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

About VMware vRealize Operations Manager



With vRealize Operations Manager enterprise software, you can proactively identify and solve emerging issues with predictive analysis and smart alerts, ensuring optimal performance and availability of system resources - across physical, virtual, and cloud infrastructures.

vRealize Operations Manager gives you complete monitoring capability in one place, across applications, storage, and network devices, with an open and extensible platform supported by third-party management packs. In addition, vRealize Operations Manager increases efficiency by streamlining key processes with preinstalled and customizable policies while retaining full control.

Using data collected from system resources (objects), vRealize Operations Manager identifies issues in any monitored system component, often before the customer notices a problem.

vRealize Operations Manager also frequently suggests corrective actions you can take to fix the problem right away. For more challenging problems, vRealize Operations Manager offers rich analytical tools that allow you to review and manipulate object data to reveal hidden issues, investigate complex technical problems, identify trends or drill down to gauge the health of a single object.

Table 4-1. Manage Solution Page Options (Continued)

Option	Description
Collectors/Groups	Determines which vRealize Operations Manager collector is used to manage the adapter processes. If you have only one adapter instance, select Default collector group . If you have multiple collectors in your environment, and you want to distribute the workload to optimize performance, select the collector to manage the adapter processes for this instance.
Auto Discovery	<p>Determines whether new objects added to the monitored system are discovered and added to vRealize Operations Manager after the initial configuration of the adapter.</p> <ul style="list-style-type: none"> ■ If the value is true, vRealize Operations Manager collects information about any new objects that are added to the monitored system after the initial configuration. For example, if you add more hosts and virtual machines, these objects are added during the next collections cycle. This is the default value. ■ If the value is false, vRealize Operations Manager monitors only the objects that are present on the target system when you configure the adapter instance.
Process Change Events	<p>Determines whether the adapter uses an event collector to collect and process the events generated in the vCenter Server instance.</p> <ul style="list-style-type: none"> ■ If the value is true, the event collector collects and publishes events from vCenter Server. This is the default value. ■ If the value is false, the event collector does not collect and publish events.
Enable Collecting vSphere Distributed Switch Enable Collecting Virtual Machine Folder Enable Collecting vSphere Distributed Port Group	When set to false, reduces the collected data set by omitting collection of the associated category.
Exclude Virtual Machines from Capacity Calculations	When set to true, reduces the collected data set by omitting collection of the associated category.
Maximum Number Of Virtual Machines Collected	<p>Reduces the collected data set by limiting the number of virtual machine collections.</p> <p>To omit data on virtual machines and have vRealize Operations Manager collect only host data, set the value to zero.</p>
Provide data to vSphere Predictive DRS	<p>vSphere Predictive DRS proactively load balances a vCenter Server cluster to accommodate predictable patterns in the cluster workload.</p> <p>vRealize Operations Manager monitors virtual machines running in a vCenter Server, analyzes longer-term historical data, and provides forecast data about predictable patterns of resource usage to Predictive DRS. Based on these predictable patterns, Predictive DRS moves to balance resource usage among virtual machines.</p> <p>Predictive DRS must also be enabled for the Compute Clusters managed by the vCenter Server instances monitored by vRealize Operations Manager. Refer to the <i>vSphere Resource Management Guide</i> for details on enabling Predictive DRS on a per Compute Cluster basis.</p> <p>When set to true, designates vRealize Operations Manager as a predictive data provider, and sends predicative data to the vCenter Server. You can only register a single active Predictive DRS data provider with a vCenter Server at a time.</p>
Enable Actions	Enabling this option helps in triggering the actions that are related to vCenter.
Cloud Type	Provides an ability to identify the type of vCenter is used in vRealize Operations Manager. By default, the cloud type is set to Private Cloud.

- 6 Click **Advanced Settings**, and in the **Collectors/Groups** text box, select the vRealize Operations Manager collector used to manage the adapter process.

If you have one adapter instance, select **Default collector group**. If you have multiple collectors in your environment, to distribute the workload and optimize performance, select the collector to manage the adapter processes for this instance.

- 7 Click **Save Settings** to finish configuration of the adapter, and click **Close**.

vRealize Business for Cloud adapter is available and is used only as a pre-configuration for SDDC health MP.

Cost Settings for Financial Accounting Model

You can configure Server Hardware cost driver and resource utilization parameters to calculate the accurate cost and improve the efficiency of your environment.

Cost Drivers analyzes the resources and the performance of your virtual environment. Based on the values you define, Cost Drivers can identify reclamation opportunities and can provide recommendations to reduce wastage of resources and cost.

Configuring Depreciation Preferences

To compute the amortized cost of the Server Hardware cost driver, you can configure the depreciation method and the depreciation period. Cost Drivers supports two yearly depreciation methods and you can set the depreciation period from two to seven years.

Note Cost Drivers calculates the yearly depreciation values and then divides the value by 12 to arrive at the monthly depreciation.

Method	Calculation
Straight line	Yearly straight line depreciation = [(original cost - accumulated depreciation) / number of remaining depreciation years]
Max of Double or Straight	Yearly max of Double or Straight = Maximum (yearly depreciation of double declining balance method, yearly depreciation of straight line method) Yearly depreciation of double declining method= [(original cost - accumulated depreciation) * depreciation rate]. Depreciation rate = 2 / number of depreciation years. Note Double declining depreciation for the last year = original cost - accumulated depreciation

Example: Example for Straight Line Depreciation Method

Year	Original Cost	Accumulated Depreciation	Straight Line Depreciation Cost
Year 1	10000	0	$[(10000-0)/5] = 2000$
Year 2	10000	2000	$[(10000-2000)/4] = 2000$

Table 4-79. Policy Library Tab Options (Continued)

Option	Description
Policy Library Tab > Details Tab	<p>The Details tab displays the name and description of the policy from which the settings are inherited, the policy priority, who last modified the policy, and the number of object groups associated with the policy. From the Details tab, you can view the settings that are locally defined in your policy, and the complete group of settings that include both customized settings and the settings inherited from the base policies selected when the policy was created.</p> <ul style="list-style-type: none"> ■ Locally Defined Settings. Displays the locally changed policy element settings for each object type in the policy. ■ Complete Settings Including Inherited. Displays all of the policy element settings for each object type in the policy, including locally changed settings and settings that are inherited. A summary of the enabled and disabled alert definitions, symptom definitions, and attributes appear indicate the number of changes in the policy. The policy element settings include symptom thresholds, and indicate changes made to the Workload, Capacity Remaining, and Time Remaining settings.
Related Objects Tab	<p>Summarizes the related groups and objects, and details about the selected object group and objects.</p> <ul style="list-style-type: none"> ■ Groups. Displays the groups of objects associated with the selected active policy, and provides options to add and release an association. <ul style="list-style-type: none"> ■ Add Association. Opens the Apply the policy to groups dialog box where you select object groups to associate with the selected policy. ■ Release Association. Opens a confirmation dialog box to confirm the release of the object group that is associated with the selected policy. ■ Data grid. Displays the groups assigned to this policy, the object types associated with the group, and the number of objects in the group. ■ Details for the selected object group. Displays the object group name, type, and number of members associated with the selected policy, and the type of association with the policy. An object group can have a direct association with a policy, and inherited policy associations based on the base policies that you selected when you created a local policy. For example, if the Base Settings policy appears in the list, with an inherited association, the Base Settings policy was included in the base policies selected when this policy was created. ■ Affected Objects. Displays the names of the objects in your environment, their object types, and associated adapters. When a parent group exists for an object, it appears in this data grid.

Operational Policies

Determine how to have vRealize Operations Manager monitor your objects, and how to notify you about problems that occur with those objects.

vRealize Operations Manager Administrators assign policies to object groups and applications to support Service Level Agreements (SLAs) and business priorities. When you use policies with object groups, you ensure that the rules defined in the policies are quickly put into effect for the objects in your environment.

With policies, you can:

- Enable and disable alerts.
- Control data collections by persisting or not persisting metrics on the objects in your environment.

The data grid provides information on which you can sort and search.

Option	Description
Name	Displays the names of the clusters in the selected parent vCenter Server instance.
Datacenter	Displays the data centers that belong to each cluster.
vCenter	Displays the parent vCenter Server instance where the cluster resides.
DRS Settings	Displays the level of DRS automation for the cluster. To change the level of DRS automation for the cluster, select Cluster Actions > Set DRS Automation from the toolbar. You can change the automation level by selecting an option from the drop-down menu in the Automation Level column.
Migration Threshold	Recommendations for the migration level of virtual machines. Migration thresholds are based on DRS priority levels, and are computed based on the workload imbalance metric for the cluster.
CPU Workload %	Displays the percentage of CPU in GHz available on the cluster.
Memory Workload %	Displays the percentage of memory in GB available on the cluster.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget .

Efficiency Widget

The efficiency widget is the status of the efficiency-related alerts for the objects it is configured to monitor. Efficiency alerts in vRealize Operations Manager usually indicate that you can reclaim resources. You can create one or more efficiency widgets for objects that you add to your custom dashboards.

How the Efficiency Widget Works

You can add the efficiency widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget .
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Objects List	List of objects in your environment that you can search or sort by column so that you can locate the object on which you are basing the data that appears in the widget. If you select an object in the list, the object becomes the selected object for the widget.

Workload Utilization Widget

The Workload Utilization widget displays a visual summary of the workload resources used by the objects in your environment.

How the Workload Utilization Widget and Configuration Options Work

Use the Workload Utilization widget to identify which workload objects are underutilized and overutilized.

You can add the Workload Utilization widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

Where You Find the Workload Utilization Widget and Configuration Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, click **Content** in the left pane, and click **Dashboards**. On the Dashboards toolbar, click the plus sign to add a dashboard or the pencil to edit the selected dashboard. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the pencil to access the configuration options.

Workload Utilization Widget and Configuration Options

The Workload Utilization widget includes toolbar and configuration options.

Option	Description
Action	Displays the available actions for a specific object. For example, if you select the host object icon, the Action icon is enabled and displays all the available actions you can carry out. Some of the options are: Power Off VM , Power On VM , and so on . The actions displayed change based on the type of object you select. The button is dimmed when actions are not available for an object you select.
Constrained by	Sorts the objects in the chart based on a metric you select. For example, if you select CPU Demand, all the objects constrained by CPU demand are displayed in the chart. You can sort the chart based on options like: CPU , CPU Demand , Memory , Memory Consumed , and vSphere Configuration Limit .
Reset to initial object	Displays the original view of the chart.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget .
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Select Object	Your inventory explorer where you can locate the object on which you are basing the data that appears in the widget.
Object Type	Select specific object types to see in the charts. Press Ctrl+click to select multiple object types. If you leave the object type deselected, you see all base object children in the charts.

Dashboards

Dashboards present a visual overview of the performance and state of objects in your virtual infrastructure. You use dashboards to determine the nature and timeframe of existing and potential issues with your environment. You create dashboards by adding widgets to a dashboard and configuring them.

vRealize Operations Manager collects performance data from monitored software and hardware resources in your enterprise and provides predictive analysis and real-time information about problems. The data and analysis are presented through alerts, in configurable dashboards, on predefined pages, and in several predefined dashboards.

- You can start with several predefined dashboards in vRealize Operations Manager.

You can use the dashboard widgets in several ways.

- **Select a Distributed Switch:** Use this widget to select the switch for which you want to view details. You can use the filter to narrow your list based on several parameters. After you identify the switch that you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Distributed Port Groups on the Switch:** Use this widget to view the port groups on the switch, how many ports each switch has, and the usage details.
- **ESXi Hosts/VMs Using the Selected Switch:** Use these widgets to find out which ESXi hosts and VMs use the selected switch. You can also view configuration details about the ESXi hosts and VMs that use the selected switch.

Host Configuration Dashboard

The Host Configuration dashboard provides an overview of your ESXi host configurations, and displays inconsistencies so that you can take corrective action.

The dashboard also measures the ESXi hosts against the vSphere best practices and indicates deviations that can impact the performance or availability of your virtual infrastructure. Although you can view this type of data in other dashboards, in this dashboard you can export the ESXi configuration view and share it with other administrators.

VM Configuration Dashboard

The VM dashboard focuses on highlighting the key configurations of the virtual machines in your environment. You can use this dashboard to find inconsistencies in configuration within your virtual machines and take quick remedial measures. You can safeguard the applications which are hosted on these virtual machines by avoiding potential issues due to misconfigurations.

Some of the basic problems the dashboard focuses on includes identifying VMs running on older VMware tools versions, VMware tools not running, or virtual machines running on large disk snapshots. VMs with such symptoms can lead to potential performance issues and hence it is important that you ensure that they do not deviate from the defined standards. This dashboard includes a predefined Virtual Machine Inventory Summary report which you can use to report the configurations highlighted in this dashboard for quick remediation.

You can use the dashboard widgets in several ways.

- Use the Large VMs widgets to view graphical representations of VMs that have a large CPU, RAM, and disk space.
- **Guest OS Distribution:** Use this widget to view a break up of the different flavors of operating systems you are running.
- **Guest Tools Version** and **Guest Tools Status:** Use these widgets to identify if you have inconsistent or older version of VMware tools which might lead to performance issues.
- View the VMs with limits, large snapshots, orphaned VMs, VMs with more than one NIC, and VMs with a nonstandard operating system. These VMs have a performance impact on the rest of the VMs in your environment even though they do not fully use their allocated resources.

Table 4-236. Support Bundle Data Grid Options

Option	Description
Bundle	System-generated identifier for the support bundle
Bundle Type	<ul style="list-style-type: none">■ Light. Include 24 hours of logs■ Full. Include all available logs and configuration files
Date and Time Created	Time when support bundle creation began
Status	Progress of support bundle creation

Update the Reference Database for vRealize Operations Manager

You can update the reference database to have the most updated version of the reference library. The reference database supplies default values for cost calculations.

Procedure

- 1 In the menu, click **Administration** and in the left pane click **Support > Cost Reference Database**.
The existing version of the reference database along with the date is displayed.
- 2 Click **Download Here**.
The latest version of the reference database is downloaded to the default location.
- 3 Click **Upload Reference Database** and select the reference database from the default download location.

Note that the updated reference library values are reflected in the cost drivers only after the cost calculation process runs as per the schedule.

Configuring and Using Workload Optimization

Workload Optimization provides for moving virtual compute resources and their file systems dynamically across datastore clusters within a data center or custom data center.

Using Workload Optimization, you can rebalance virtual machines and storage across clusters, relieving demand on an overloaded individual cluster and maintaining or improving cluster performance. You can also set your automated rebalancing policies to emphasize VM consolidation, which potentially frees up hosts and reduces resource demand.

Workload Optimization further enables you potentially to automate a significant portion of your data center compute and storage optimization efforts. With properly defined policies determining the threshold at which resource contention automatically runs an action, a data center performs at optimum.

vRealize Automation Integration

When you add an instance to a vRealize Automation adapter or solution pack as well as to a vCenter Server adapter instance that is connected to the vRealize Automation server, using vRealize Automation-managed resources, vRealize Operations Manager automatically adds a custom data center for the vCenter Server, using vRealize Automation-managed resources.

On the vRealize Operations Manager side, to get the day2 chain configured, you must make the following initial configurations:

- 1 In vCenter Server, **Administration -> Solutions** and then add the VMware vSphere adapter instance for the vCenter Server that is configured as an endpoint in vRealize Automation Server.
- 2 In vCenter Server, **Administration -> Solutions** and then add the VMware vRealize Automation adapter instance for the server that will appear in the vRealize Operations Manager and vRealize Automation integration day2 chain.

vRealize Operations Manager can manage workload placement and optimization for the custom data centers that reside in vRealize Automation-managed clusters.

However, vRealize Operations Manager is not permitted to set tag policies for the custom data center. (At the Workload Optimization screen, the Business Intent window is not operational for vRealize Automation custom data centers.) When rebalancing a vRealize Automation custom data center, vRealize Operations Manager uses all applicable policies and placement principles from both systems: vRealize Automation and vRealize Operations Manager. For more information on configuring vRealize Automation to work with vRealize Operations Manager, see [vRealize Automation Solution](#) For complete information on creating and managing vRealize Automation custom data centers that are managed by vRealize Operations Manager, see the vRealize Automation documentation.

Configuring Workload Optimization

Workload Optimization offers you the potential to automate fully a significant portion of your cluster workload rebalancing tasks. The tasks to accomplish workload automation are as follows:

- 1 Configure the Workload Automation Details. See [Workload Automation Details](#).
- 2 Tag VMs for cluster placement. See [Business Intent - Host-Based Virtual Machine Placement](#) and [Business Intent: Tag-Based VM Placement in Clusters](#).
- 3 If you do not use the AUTOMATE function in the Optimization Recommendation pane at the Workload Automation screen, configure the two Workload Optimization alerts to be triggered when cluster CPU/memory limits are breached, and configure them as automated. When the alerts are automated, the actions calculated by Workload Optimization are run automatically. See [Configuring Workload Optimization Alerts](#)

Prerequisites

Workload Optimization acts on objects associated with the VMware vSphere Solution that connects vRealize Operations Manager to one or more vCenter Server instances. The virtual objects in this environment include a vCenter Server, data centers and custom data centers, cluster compute and storage resources, host systems, and virtual machines. Specific requirements:

- A vCenter Adapter configured with the actions enabled for each vCenter Server instance.
- A vCenter Server instance with at least two datastore clusters with sDRS enabled and fully automated.
- Any non-datastore clusters must have DRS enabled and fully automated

Cluster	Environment	Service Tier	Licensing
11	Dev	Gold	
12, 13	Dev	Silver	
14, 15	Dev		Oracle

Opening the vRealize Operations Manager policies to Tag-Based VM Placement in the Business Intent window, the administrator prioritizes the Environment: Production and Service Tier: Gold category-tag combinations. Because the Optimization policies emphasize balance, clusters with those tags are balanced first.

Business Intent - Host-Based Virtual Machine Placement

Use host-based VM placement to tie your VMs more closely to your infrastructure. By using vCenter Server to tag hosts and VMs with specific tags, you make certain that when the system runs an optimization, it uses VM-to-host tag matching to ensure that VMs are moved to - or stay with - the appropriate host.

Using Tags to Enhance Structure

When configuring data centers or custom data centers without tags, you configure clusters and their hosts as relatively homogenous. All cluster resources must support, for example, the same OS or the same security requirements so that optimization actions do not place VMs in the wrong cluster.

The tagging approach enables you to define zones of infrastructure within cluster boundaries. VM-to-cluster tagging, for example, allows you to tag VMs and clusters to assure that Windows VMs are moved only to Windows-licensed clusters and Oracle VMs are moved only to Oracle-licensed clusters.

With host-based VM placement (VM-to-host tagging), you bind your VMs to individual hosts rather than clusters.

vCenter Server tags are implemented as *key:value* labels that enable operators to add meta-data to vCenter Server objects. In vCenter Server terminology, the *key* is the tag category and the *value* is the tag name. You can define many keys and values in vCenter Server, but choose a subset to be considered in the Business Intent pane of the Workload Optimization screen (**Home -> Optimize Performance -> Workload Optimization**).

Note If you choose host-based placement in the Business Intent pane, the system - after getting confirmation from you - disables conflicting user-created affinity rules. Then, as you define host-VM tagging relationships in the Business Intent pane, vRealize Operations Manager automatically creates the required affinity rules, saving you the manual effort. So, for example, suppose you configure a tag in the Business Intent pane that requires VM1 to remain with Host1. If there exists a user-configured affinity rule keeping VM1 with Host2, the system disables the rule. However, if another user-configured affinity rule dictates that VM2 remains with Host2, the system does not change that rule.

Additional Considerations

- You are not permitted to employ both VM-to-cluster tagging and VM-to-host tagging in the same data center or custom data center - only one tagging method or the other. If you select host-based VM placement, any cluster tags are ignored.
- With host-based VM placement, only one category and one tag per VM is allowed per VM.
- A tagless VM can be sent to any host, even a tagged host.
- A host with multiple tags is treated as tagless.
- Even if all workloads are balanced, if there is also a tag violation, the system is by definition not optimized.
- The system does not consider any tags of storage - that is, datastores or datastore clusters.

Business Intent Workspace

You can use vCenter Server tagging to tag VMs, hosts, and/or clusters with specific tags. vRealize Operations Manager can be configured to leverage tags to define business-related placement constraints: VMs can only be placed on hosts/clusters with matching tags.

Where You Find Business Intent

From the Home page, click the chevron next to Optimize Performance on the left. Click Workload Optimization, select a data center or custom data center from the top row, and click **Edit** in the Business Intent window.

To edit Business Intent values, you must have privileges for Administration -> Configuration -> Workload Placement Settings -> Edit.

Establishing Business Intent

Tags are implemented in vCenter Server as *key:value* labels that enable operators to add meta-data to vCenter Server objects. In vCenter Server terminology, the *key* is the tag category and the *value* is the tag name. Using this construct, the tag OS: Linux can indicate a cluster or VM that is assigned to the category OS with a tag name of Linux. For complete information on vCenter Server tagging capabilities, refer to the vCenter Server and Host Management guide.

To specify tags considered for placement, first select the radio button for the type of object you want to associate with VMs in this business intent session: Clusters or Hosts.

The system provides several suggested categories. These categories are only suggestions. You must specify the actual categories in vCenter Server after you expand the section for a suggested category . For example, in section "Tier", you can specify the actual vCenter Server tag category that represents tier semantics, for instance, "service level".

- Operating System
- Environment
- Tier

Procedure

- 1 Select **Administration** from the menu, then **Policies** from the left pane.
- 2 Click **Policy Library** and select the policy that includes settings for the relevant data centers and custom data centers, for example, **vSphere Solution's Default Policy**.
- 3 Click **Edit**.
- 4 Click #6 on the lower left, Alert/Symptom Definitions.
- 5 Search on "can potentially be optimized" to locate the two alerts you want.
- 6 The alerts are turned ON by default/inheritance (see the State column).
- 7 The alerts are not automated by default/inheritance (see the Automate column). To automate the alerts, click the menu symbol to the right of the inherited value and select the green check mark.

Workload Optimization is fully automated for your environment.

What to do next

To confirm that actions are taken automatically, monitor rebalance activity at the Workload Optimization screen.

Using Workload Optimization

Use the Workload Optimization UI pages to monitor optimizing moves in a fully automated system. If your system is not fully automated, you can use the UI to conduct research and run actions directly.

vRealize Operations Manager monitors virtual objects and collects and analyzes related data that is presented to you in graphical form at the Workload Optimization screen. Depending on what appears on the screen, you might use optimization functions to distribute a workload differently in a data center or custom data center. Or you may decide to perform more research, including checking the Alerts page to determine if any alerts have been generated for objects of interest.

For comprehensive general instructions on responding to alerts and analyzing problems related to objects in your environment, see [Chapter 5 Monitoring Objects in Your Managed Environment by Using vRealize Operations Manager](#).

The following examples demonstrate the primary ways you can use Workload Optimization to keep your data centers balanced and performing their best.

Example: Run Workload Optimization

As a virtual infrastructure administrator or other IT professional, you use Workload Optimization functions to identify points of resource contention or imbalance. In this example, you manually run an optimization action to consolidate demand.

When you log into vRealize Operations Manager, you see the Quick Start page. In the left-most column, Optimize Performance, is the alert 3 DATA CENTERS REQUIRING OPTIMIZATION.

Understanding the Host Summary Tab

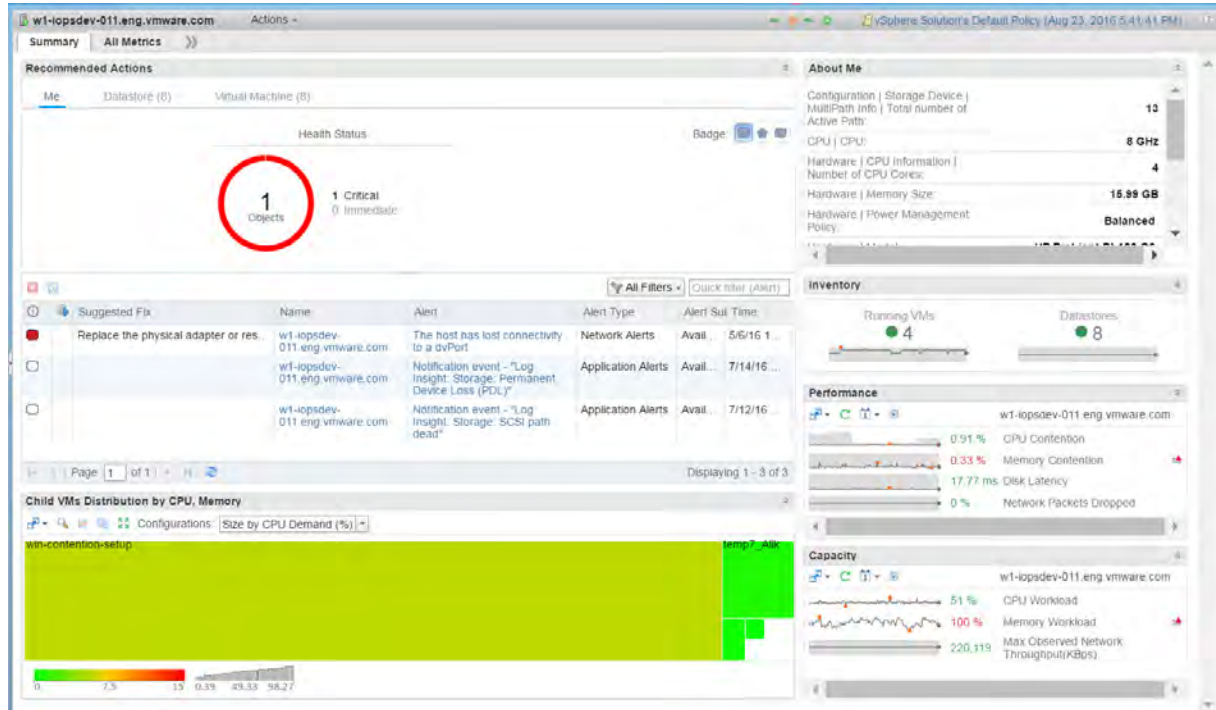


Table 5-3. Host Summary Tab Options

Option	Description
Recommended Actions	<p>This widget displays the health status for the selected object and its descendants. It also displays recommendations to solve problems in an instance.</p> <p>The badges provide a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> Health alerts that usually require immediate attention. Risk alerts indicating that you must look into any problems soon. Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the badge .</p>
About Me	<p>This widget displays the key metrics and properties of the selected object.</p>
Inventory	<p>This widget displays the number of running VMs and Datastores associated with the selected host.</p>
Capacity	<p>This widget displays a visual summary of the capacity and workload resources used by the objects in your environment. It displays the latest value and a trend line of the various key indicators in a color that indicates its health based on the symptom associated with the metrics. Double-click each metric to see the detailed chart.</p>

Capacity Optimization for Your Managed Environment

6

Capacity Optimization in vRealize Operations Manager is achieved using powerful integrated functions - capacity overview, workload balancing and optimization, repurposing of underutilized resources, and what-if predictive scenarios - to reach optimal system performance.

Capacity planners must assess whether physical capacity is sufficient to meet current or forecasted demand. With robust capacity planning and optimization, you can manage your production capacity effectively as your organization addresses changing requirements. The objective of strategic capacity optimization is to reach an optimal level where production capabilities meet ongoing demand.

vRealize Operations Manager analytics provide precise tracking, measuring and forecasting of data center capacity, usage, and trends to help manage and optimize resource use, system tuning, and cost recovery. The system monitors stress thresholds and alerts you before potential issues can affect performance. Multiple pre-set reports are available. You can plan capacity based on historical usage, and run what-if scenarios as your requirements expand.

How Capacity Optimization Works

The Capacity Optimization provides four integrated functions - Overview, Reclaim, Workload Optimization, and What-If Scenarios - that give an overview of the status of all data center activity and trending. You can conduct on-the-spot analysis, including drilling down into further detail on any object to identify possible performance problems or anomalies. You can rebalance and optimize compute resources. The system further identifies underutilized workloads (virtual machines) and calculates the potential cost savings that can accrue when these resources are reclaimed to be deployed more effectively. You can interact with and manipulate data and outcomes based on your requirements.

Use the Capacity Optimization and Reclaim features to assess workload status and resource contention in data centers across your environment. You can determine time remaining until cpu, memory, or storage resources run out and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.

Workload Optimization provides for moving virtual workloads and their file systems dynamically across datastore clusters within a data center or custom data center. You can potentially automate a significant portion of your data center compute and storage optimization efforts. With properly defined policies determining the threshold at which resource contention triggers an alert and automatically runs an action, a data center performs at optimum.

Exhibit 6

UNDERSTANDING vSPHERE DRS PERFORMANCE

VMware vSphere 6

Introduction

VMware vSphere® Distributed Resource Scheduler™ (DRS) is the resource scheduling and load balancing solution for vSphere. DRS works on a cluster of ESXi hosts and provides resource management capabilities like load balancing and virtual machine (VM) placement. DRS also enforces user-defined resource allocation policies at the cluster level, while working with system-level constraints.

Although DRS is widely deployed and generally understood, questions about "how" DRS does what it does are not uncommon. Not knowing exactly how DRS works often leads to confusion and improper expectations about DRS behavior and its performance.

In this paper, we demystify DRS. First, we give an overview of how DRS works. Next, we explain some of the factors that influence DRS behavior and how they can be monitored. Finally, we cover some common performance scenarios that customers raised as issues and show what we can learn about DRS from these scenarios.

How DRS Works

The main goal of DRS is to ensure that VMs and their applications are always getting the compute resources that they need to run efficiently. In other words, *DRS strives to keep your VMs happy*¹. It does this by ensuring that newly powered-on VMs get all the required resources soon after they are powered on, and the resource utilization is always balanced across the cluster.

From time to time, VMs' workloads may change, and with many VMs with changing workloads, there can be imbalance in the cluster. Each of these can degrade application performance. DRS solves these problems by regularly monitoring the cluster balance state once every five minutes, by default, and then takes the necessary actions to fix any imbalance. DRS automatically determines which virtual machines would benefit from a move to another host and live migrates the VM onto the new host using vMotion. In this way, DRS ensures each virtual machine in the cluster gets the host resources—like memory and CPU—that it needs.

Let's take a closer look at how DRS achieves its goal of ensuring VMs are happy, with effective placement and efficient load balancing.

Effective VM Placement

When a VM is being powered up in a DRS cluster, DRS runs its algorithm to determine the right ESXi host for it to be powered up on. This decision, also known as VM placement (or initial placement) is made based on the expected change in resource distribution (after ensuring that there are no constraint violations if the VM was placed on the host).


One of the first steps in ensuring good VM performance is to make sure that the VM gets all the resources it needs as soon as it is powered on. DRS considers the demand of a VM, so it will never be short of resources whenever it is started. A VM's **demand** includes the amount of resources it needs to run, and the way DRS calculates this is described in [Calculating VM Resource Demand](#).

¹ For VMs and their applications to perform well, they rely on DRS to provide the necessary resources. In reality, application performance depends on more than just the availability of resources. DRS can only ensure that lack of resource availability is not the reason for any application performance issues in your cluster.


Exhibit 7

Workload Balance

Cluster Compute Resource



Host System



Virtual Machine




Exhibit 8



Exhibit 9

What's New in vRealize Operations 7.0

Today, VMware announced the upcoming release of vRealize Operations 7.0, which will help customers derive even more value from a “Self-Driving” approach to operations management.

The [Self-Driving vision is to automate and simplify operations management](#) by delivering per the “Three Tenets” of Self-Driving Operations.

1. Intent-Driven Continuous Performance Optimization
2. Efficient Capacity Management
3. Intelligent Remediation

As you will recall, we introduced self-driving operations in March this year, with the release of vRealize Operations 6.7. Self-driving summarizes the strategy and vision to provide continuous and automated closed loop performance and capacity optimization at minimal cost that is based on business and operational intent. Think about it, as a customer all you need to do is define operational and business intent, and let the platform take care of the rest to assure performance, densify clusters or enforce software license separation. Once intent is defined, the platform continuously verifies workload performance against defined intent, applying predictive analytics to project future requirements, and automatically takes actions to balance workloads and right-size VMs to optimize performance and capacity.

The upcoming vRealize Operations 7.0 will augment the capabilities introduced in the last release, particularly focusing on:

- **Business intent-driven continuous performance optimization** with new automated workload balancing capabilities, enhanced integration with vRealize Automation for both initial placement and on-going workload balancing, new host-based placement feature, and workload right-sizing workflows
- **Efficient capacity management** by enhancing the real-time, predictive capacity & cost analytics engine and adding multiple what-if planning scenarios
- **Multi-cloud** support with migration planning across VMware Cloud on AWS and native AWS, as well as updated management packs for AWS and Kubernetes
- **Quick time to value** with updates to the UI, simplified custom dashboard creation and sharing, enhanced SDDC integrations and built-in vSphere config & compliance

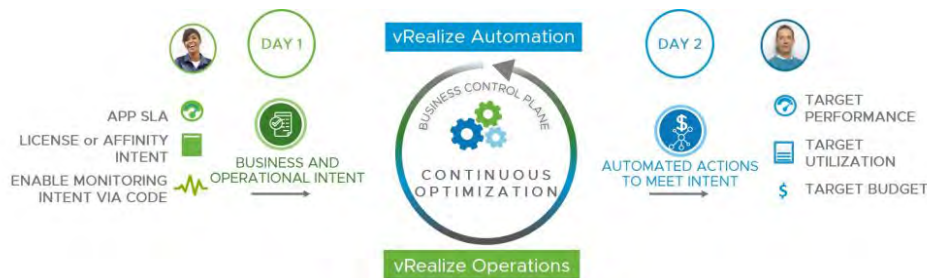
Learn more at VMworld: <https://blogs.vmware.com/management/2018/08/vmworld-las-vegas-recommended-self-driving-operations-sessions.html>

Let's double click into some of these capabilities.

Closed Loop Performance Optimization

With the upcoming vRealize Operations 7.0 release, VMware will enhance the integration between vRealize Operations and vRealize Automation to deliver closed loop performance optimization. Deep integration between vRealize Operations 7.0 and vRealize Automation 7.5 will enable full Cloud Management Platform (CMP) integration for **initial and ongoing placement** of workloads across clusters based on operational and business intent (e.g., utilization, compliance, license cost)

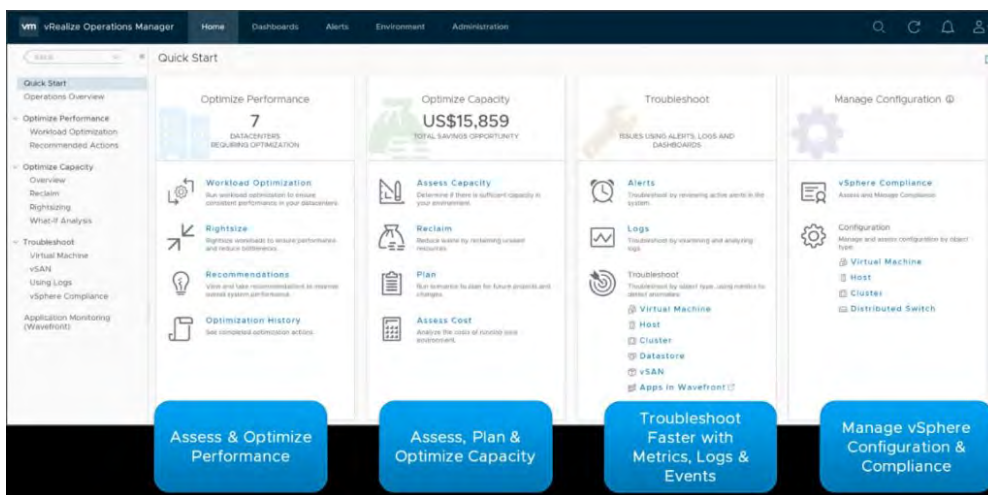
vRealize Automation can already take advantage of the analytics and intent from vRealize Operations for best initial placement of workloads, provisioning new deployments based on available capacity and usage trends. Now let's say some time passes and either vRealize Operations finds an opportunity to optimize performance or a datacenter goes red, indicating resource imbalance. vRealize Operations will honor the reservation policies for placement and balancing and only migrate virtual machines between clusters which are part of the same reservation.



Key Highlights of vRealize Operations 7.0

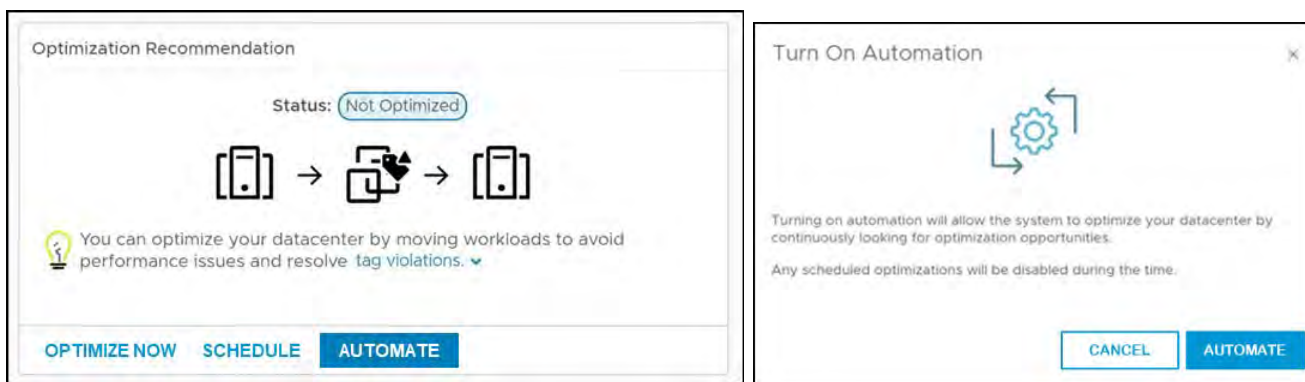
1. Enhanced User Interface

The new release will make vRealize Operations even simpler to use, featuring an updated use case and persona-based 'Quick Start' dashboard.



2. Automating Performance Based on Business and Operational Intent

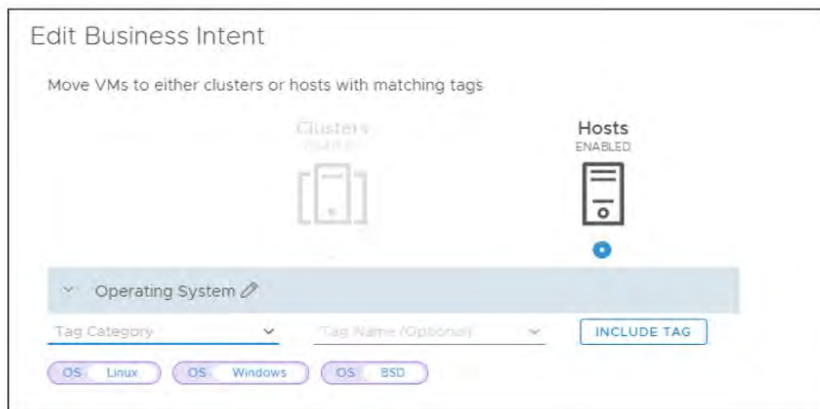
In the previous release of vRealize Operations 6.7, we introduced performance optimization based on intent. In this release of vRealize Operations 7.0, we are enhancing this capability by introducing the **Automate** button. When the **Automate** button is selected, vRealize Operations will automatically search for optimization opportunities and execute accordingly. Customers will still have the options to trigger optimization actions immediately and manually or schedule it to occur in a convenient maintenance window.



3. Automated Host Based Placement, Driven by Business Intent

This new capability will provide customers a way to automate Distributed Resource Scheduler (DRS) based on business intent. Typically, vRealize Operations balances workloads across clusters and DRS resolves contention within the cluster. Now, you will be able to teach DRS your business intent and control not only balancing across clusters, but also which host within a

cluster the workload will land on. This can be used for license separation, compliance, tiering and more. For example, if you have two clusters in a datacenter or custom datacenter, you will be able to specify which hosts are for Windows workloads, Linux, or MySQL and vRealize Operations will balance workloads across these clusters while ensuring that Windows workloads land on Windows hosts, Linux on Linux hosts ... you get the picture. You will also be able use this capability to fix tag violations and even segment a single giant cluster and manage license separation.



4. Capacity Analytics Enhanced with Exponential Decay and Calendar Awareness


The new capacity analytics engine delivered in vRealize Operations 6.7 was groundbreaking. We've continued to build from that momentum in vRealize Operations 7.0 with additional enhancements, particularly with **Exponential Decay and Calendar Awareness**.

Delivered in vROps 6.7

- Real-Time Predictive Capacity Analytics, based on industry-standard ARIMA statistical analysis model
- Improved capacity accuracy; including for Workload Optimization (aka Workload Balancing)
- Self-learning, updated in real-time; calculations available immediately
- Integrated costing with capacity

What's New in vROps 7.0

- Exponential decay to give more relevance to changing patterns and react better to more recent spikes without losing periodicity
- Improved calendar aware periodicity to detect, the n^{th} day of the month, the end of the month, and the first day of the month



5. Plan Capacity across Private Cloud and VMware Cloud on AWS

The **What-If Analysis** feature from vRealize Operations 7.0 can potentially help make room for new applications, procure hardware or migrate workloads to the cloud. The new release will offer three what-if scenarios to:

- **Workload Planning** to figure out the best fit for these new VMs
- **Physical Infrastructure Planning** to model hardware purchases with CapEx visibility
- **Migration Planning** helps plan cloud migration to VMware Cloud on AWS or AWS natively with detailed drill downs into costs and capacity requirements and recommendations



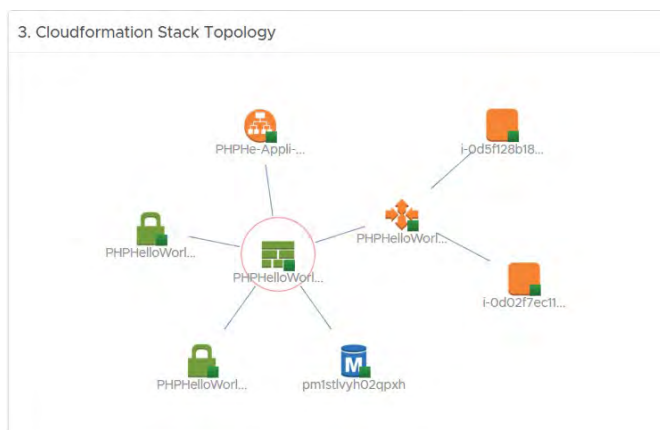
6. Simplified Dashboard Creation and Sharing

In vRealize Operations 7.0, we've simplified the **dashboard creation** process, adding an intuitive canvas and multiple out of the box widgets to improve the user experience. Dashboard sharing and embedding will also become easier using smart links without requiring login, improving the cross-team collaboration and reporting for users.



7. vRealize Operations AWS Management Pack Update

Leveraging the updated AWS management pack, you will be able to manage your AWS inventory across regions and accounts and check the availability of AWS services, with **28 new dashboards**, summary pages and **34 new alerts**. You will also be able to reclaim unused AWS resources and get recommendations on EC2 instance types. You will have support for 24 AWS services, including EC2, RDS, EBS, LB, Lambda, Redshift, etc.



8. Other Miscellaneous Enhancements

- **Workload Right-sizing** to avoid performance bottlenecks and reclaim over-allocated resources
- Built-in **vSphere config & compliance**: PCI, HIPAA, DISA, FISMA, ISO, CIS
- Ability to extend to the **entire data center** and cloud with updated management packs for Storage, vRO, Kubernetes, Federation etc.

- vSAN performance, capacity, and troubleshooting including support for stretched clusters and through vRealize **Operations plug-in in vCenter**
- **Wavefront integration** for application operation
- .. And More!

These are just some of the highlights of the new release. Stay tuned for General Availability announcement.

For more information, take the guided tour: <https://www.vmware.com/go/vrealize-guides>

Visit the vRealize Operations home page: <https://www.vmware.com/products/vrealize-operations.html>

Exhibit 10

Using Host Rules with Business Intent in vRealize Operations 7.0

John Dias (<https://blogs.vmware.com/management/author/jdias>) posted November 27, 2018

2 Comments

Tweet

Like 0

Share

This blog post originally appeared on the blog of Brett Johnson (<https://sdbrett.com/BrettsITBlog/2018/11/host-rules-with-business-intent/>), a Senior Consultant with VMware. It is re-posted here with Brett's permission. You can follow Brett on Twitter @BrettJohnson008 (<https://twitter.com/brettjohnson008>).

vROPS 7.0 launched with a feature called Business Intent which controls workload placement based on vSphere tags Announcement (<https://blogs.vmware.com/management/2018/08/whats-new-in-vrealize-operations-7-0.html>). This feature tackles several use cases one being controlling the host placement of workloads which have physical licencing requirements.


Business Intent settings are configured at the data centre (or custom datacenter) level within vROPS 7.0 and can be configured at the per cluster or per host level.


For the scope of this post, I am going to cover a customer use case to combine clusters, increasing availability for some workloads while maintaining licencing compliance for others.

This workloads within the environment are about 98% RHEL and only a couple are Windows-based. Both RHEL and Windows operating systems are licenced by the physical host that the VM can run on. The more hosts within a cluster, the more OS licences required.

In this scenario, there were 2 clusters,

- Appl: 4 Physical hosts, all licenced for RHEL
- ApplW: 2 Physical hosts, all licenced for both RHEL and Windows

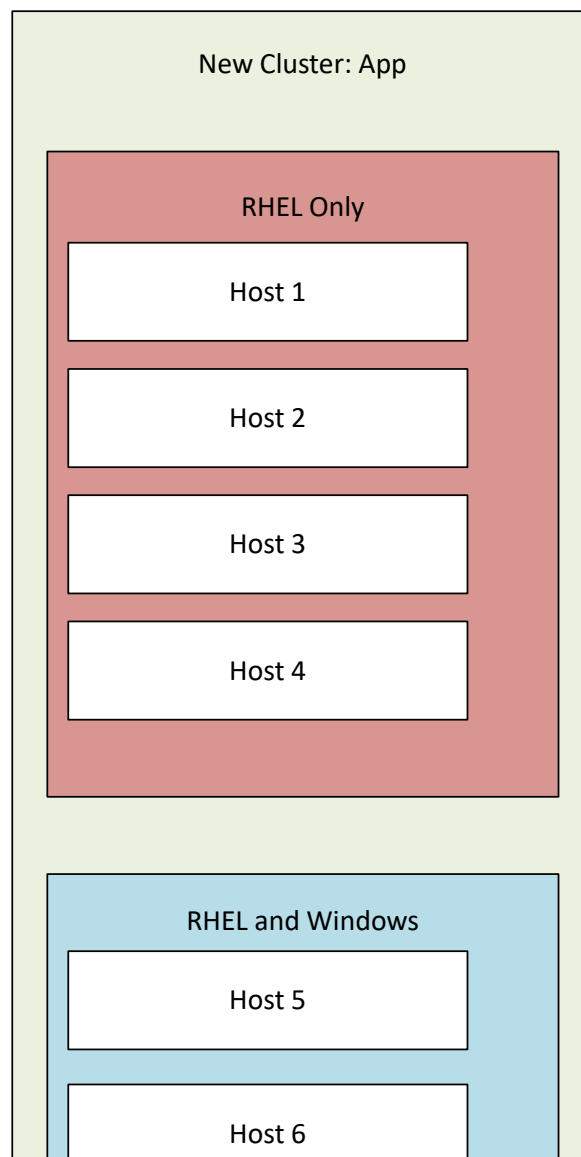




(<https://sdbrett.com/assets/images/vROPS-BI-Existing-Clusters.svg>)

The customer was experiencing sub-optimal resource utilization with the existing cluster configuration. With the launch of vROPS 7.0, we decided to see if improvements could be made with using Business Intent.

The goal of implementing Business Intent rules was to consolidate the two clusters, increasing the availability of the RHEL VMs while maintaining the host alignment restrictions of the Windows VMs.





(<https://sdbrett.com/assets/images/vROPS-BI-New-Cluster.svg>)

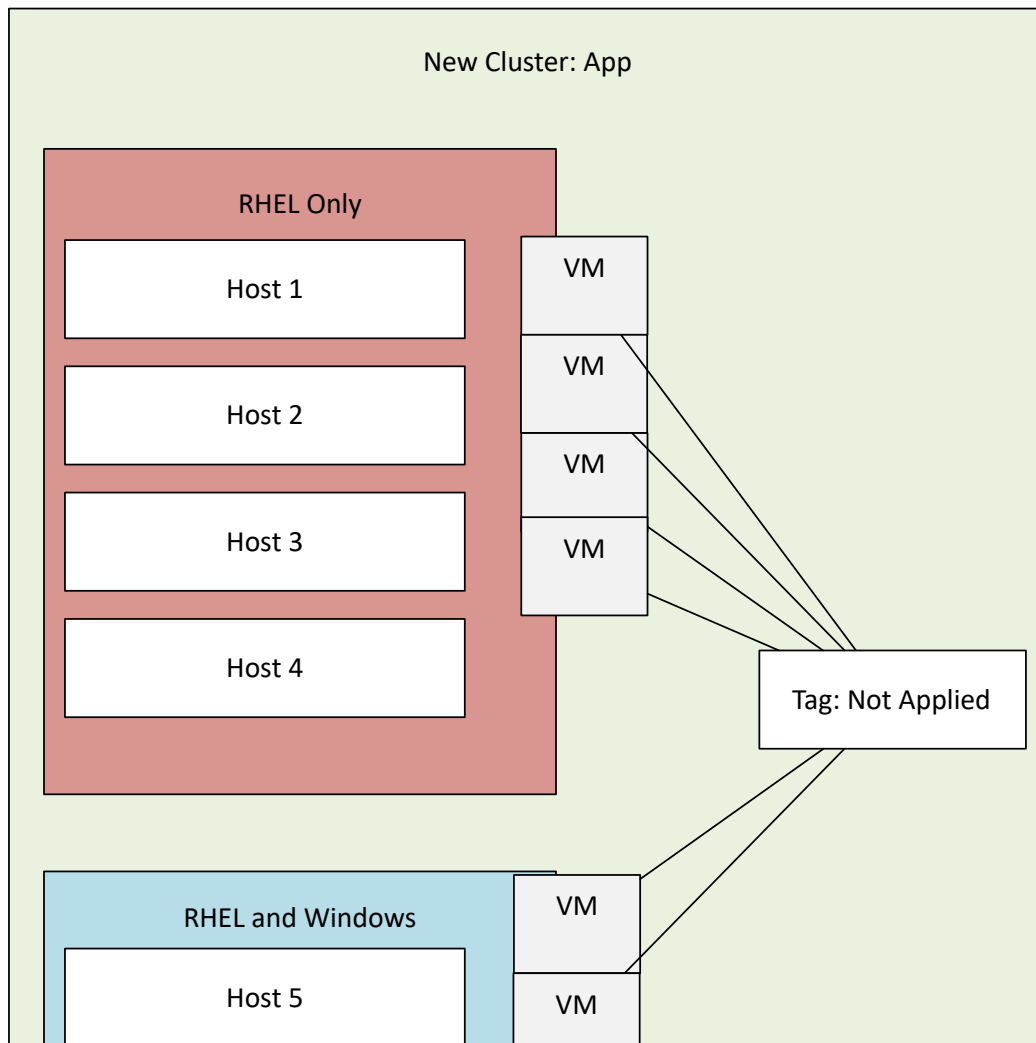
Under the hood, Business Intent uses vSphere tags on hosts and VMs to create DRS affinity groups. Those groups are used to create DRS ‘must run on’ rules to enforce the policy.

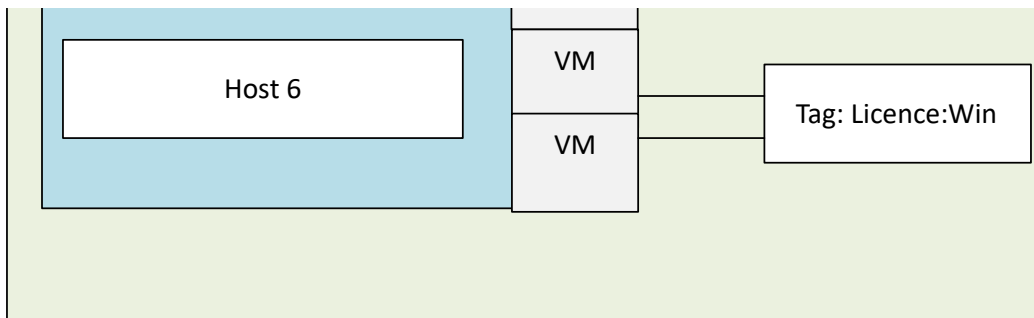
The initial plan was to create a vSphere tag category called ‘Licence’ and two tags within; the first tag was ‘RHEL’ and the second was ‘Windows’. I would assign RHEL only hosts the RHEL tag and on hosts which could run Windows and RHEL I would assign both.

During the setup, I learnt that a single host could not have multiple tags from the same selected category. My initial plan would not work.

After some e-mails, a new way to look at the solution clicked. Lightbulbs and everything.

Business Intent rules only apply DRS ‘must rules’ to VMs and Hosts which have the required tag. It does not impact VMs or Hosts which do not have these tags. I could have the RHEL workloads run on all 6 hosts within the cluster by doing nothing. The Windows workloads were limited to specific hosts by applying the tags.





(https://sdbrett.com/assets/images/vROPS-BI-New-Cluster-Tags.svg)

Example configuration and results:

```
PS C:\Program Files\PowerShell\6.0.2> get-vmhost super-esxi04.virtualiseme.com.au | Get-TagAssignment
```

Tag	Entity
---	-----
Licence/Windows	super-esxi04.virtualiseme.c...

(https://sdbrett.com/assets/images/vrops-bi-catagory.png)

```
PS C:\Program Files\PowerShell\6.0.2> get-vmhost super-esxi04.virtualiseme.com.au | Get-TagAssignment
```

Tag	Entity
---	-----
Licence/Windows	super-esxi04.virtualiseme.c...

(https://sdbrett.com/assets/images/vrops-bi-host-tag.png)

```
PS C:\Program Files\PowerShell\6.0.2> get-vm virtualisem0765 | Get-TagAssignment
```

Tag	Entity
---	-----
vRADeploymentId/80ea37fb-4550-4a6d-83...	Virtualisem0765
Licence/Windows	Virtualisem0765

```
PS C:\Program Files\PowerShell\6.0.2>
```

(https://sdbrett.com/assets/images/vrops-bi-vm-tag.png)

Edit Business Intent



Move VMs to either clusters or hosts with matching tags

Clusters
DISABLED



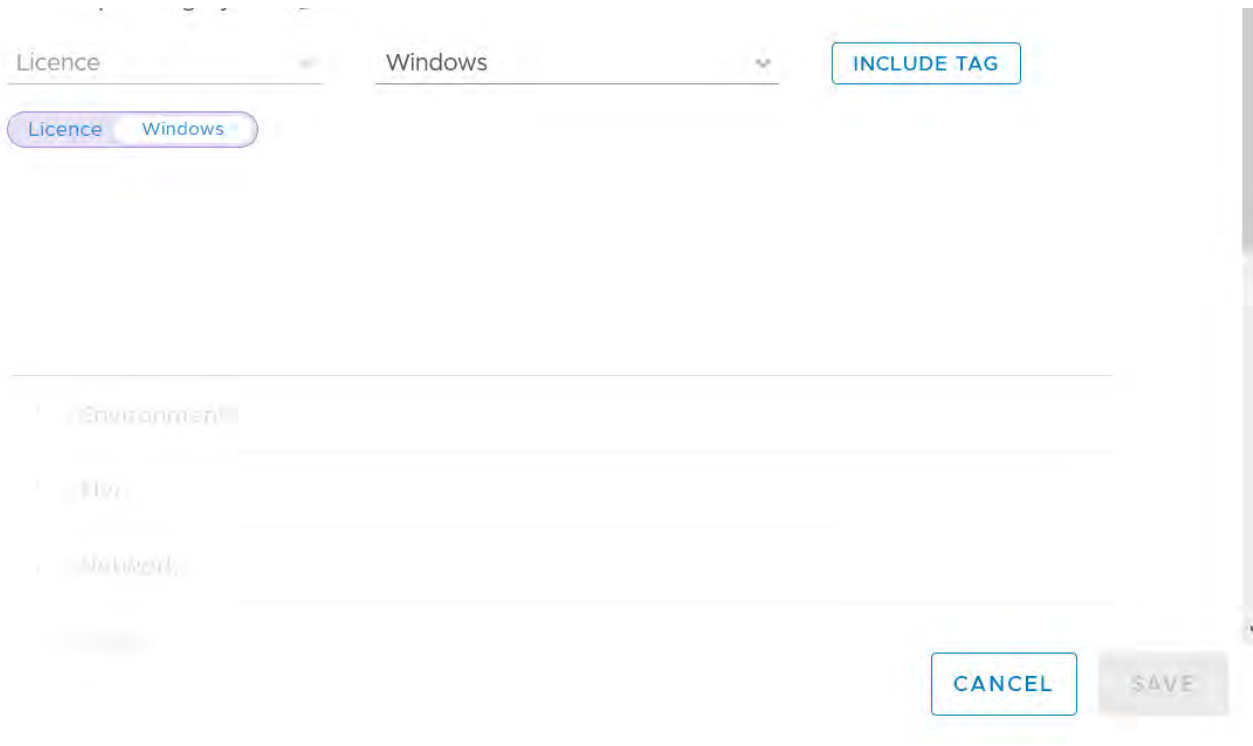
Hosts
ENABLED




Select the criteria you would like to use for placement of VMs. For host-based placement, only one category can be used at a time. This will ensure VMs are mapped to the appropriate hosts if moved.

[CLEAR ALL TAGS](#)

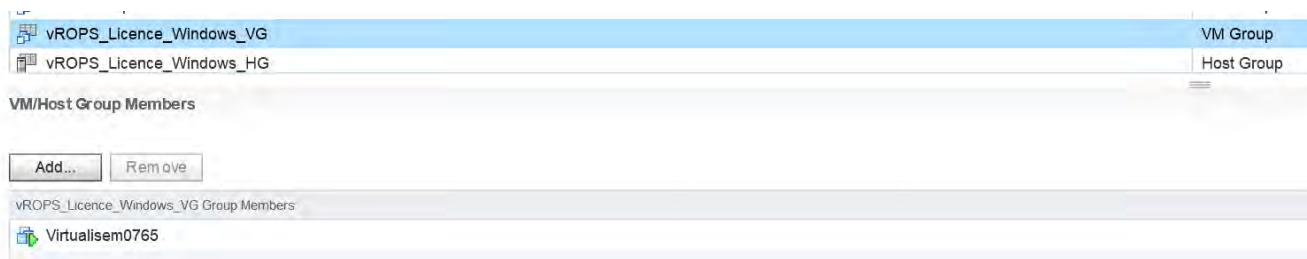
Operating System



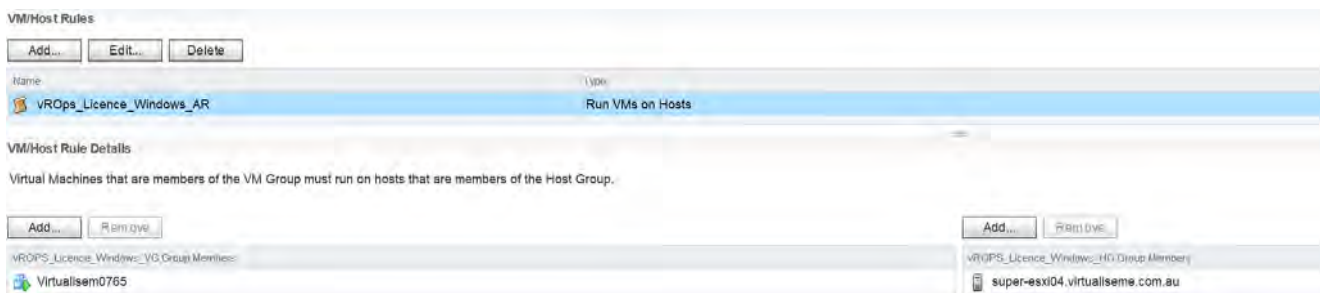
(<https://sdbrett.com/assets/images/vrops-bi-setup.png>)



(<https://sdbrett.com/assets/images/vrops-bi-affinity-hostgroup.png>)



(<https://sdbrett.com/assets/images/vrops-bi-affinity-vm-group.png>)



(<https://sdbrett.com/assets/images/vrops-bi-affinity-rule.png>)

From the above affinity rule images, you can see that the settings are there to specifically restrict the hosts a VM can run on, not control VMs which do not have the required vSphere tag.


When considering host based Business Intent rules, remember that they are a 'this' construct; not

'this OR that'.

Related Posts:



This Baby is Loaded! John Dias posted September 20, 2018
What's New with
vRealize Operations
7.0 Technical
Overview
([https://blogs.vmware](https://blogs.vmware.com/management/2018/09/this-baby-is-loaded/)



Customizing request forms using the new Custom Forms Ivan Ivanov posted May 3, 2018
([https://blogs.vmware](https://blogs.vmware.com/management/2018/05/customizing-request-forms-using-the-new-custom-forms/)




vRealize Automation Daniel Zilberman posted April 26, 2018
with Infrastructure
Blueprint -
Configuring Multi-developer
Environment

PROMO

Exhibit 11

Workload Optimization

Business Intent: Densify to Repurpose Hosts



Use Case: You are looking to free up some hosts for another business need

- Move workloads until the cluster is “Full”
- Honor 20% Headroom setting
- Free up some hosts

Play (k)



0:13 / 2:52



Exhibit 12



(12) **United States Patent**
Yuyitung et al.

(10) **Patent No.:** **US 8,209,687 B2**
(45) **Date of Patent:** **Jun. 26, 2012**

(54) **METHOD AND SYSTEM FOR EVALUATING VIRTUALIZED ENVIRONMENTS**

(75) Inventors: **Tom Silangan Yuyitung**, Toronto (CA);
Andrew Derek Hillier, Toronto (CA)

(73) Assignee: **CiRBA Inc.**, Richmond Hill (CA)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 931 days.

(21) Appl. No.: **12/201,323**

(22) Filed: **Aug. 29, 2008**

(65) **Prior Publication Data**

US 2009/0070771 A1 Mar. 12, 2009

Related U.S. Application Data

(60) Provisional application No. 60/969,344, filed on Aug. 31, 2007.

(51) **Int. Cl.**
G06F 9/44 (2006.01)

(52) **U.S. Cl.** **718/1; 718/104**

(58) **Field of Classification Search** None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,138,078	A	10/2000	Canada et al.	
6,148,335	A	11/2000	Haggard et al.	
6,564,174	B1*	5/2003	Ding et al.	702/186
7,080,221	B1*	7/2006	Todd et al.	711/161
7,577,722	B1*	8/2009	Khandekar et al.	709/220
7,647,516	B2*	1/2010	Ranganathan et al.	713/320
7,886,293	B2*	2/2011	Anderson et al.	718/1
2003/0084157	A1*	5/2003	Graupner et al.	709/226
2004/0107125	A1*	6/2004	Guheen et al.	705/7

2004/0236971	A1*	11/2004	Kopley et al.	713/300
2005/0027466	A1	2/2005	Steinmetz et al.	
2006/0010176	A1*	1/2006	Armington	707/204
2006/0020924	A1	1/2006	Lo et al.	
2006/0167665	A1*	7/2006	Ata	703/2
2006/0230407	A1*	10/2006	Rosu et al.	718/105
2007/0006218	A1*	1/2007	Vinberg et al.	717/174
2007/0079308	A1*	4/2007	Chiaramonte et al.	718/1
2007/0094375	A1	4/2007	Snyder et al.	
2007/0106769	A1	5/2007	Liu	
2007/0250615	A1	10/2007	Hillier	
2007/0250621	A1	10/2007	Hillier	
2007/0250829	A1	10/2007	Hillier et al.	
2008/0011569	A1	1/2008	Hillier et al.	
2008/0256530	A1	10/2008	Armstrong et al.	

FOREIGN PATENT DOCUMENTS

CA	2382017	A1	4/2001
CA	2420076	A1	2/2003
CA	2486103	A1	4/2006
CA	2583582	A1	5/2006

(Continued)

OTHER PUBLICATIONS

Hillier, Andrew; "Large Scale Production Virtualization: Managing the Transformation"; White Paper; Oct. 2006; CiRBA Inc.

(Continued)

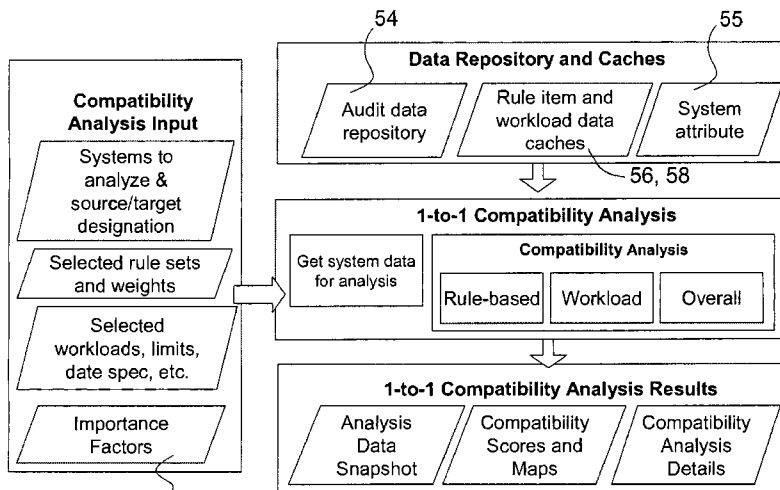
Primary Examiner — Eric Coleman

(74) *Attorney, Agent, or Firm* — Brett J. Slaney; Blake, Cassels & Graydon LLP

(57) **ABSTRACT**

A system and method are provided for incorporating compatibility analytics and virtualization rule sets into a transformational physical to virtual (P2V) analysis for designing a virtual environment from an existing physical environment and for ongoing management of the virtual environment to refine the virtualization design to accommodate changing requirements and a changing environment.

18 Claims, 56 Drawing Sheets



US 8,209,687 B2

Page 2

FOREIGN PATENT DOCUMENTS

WO WO 03/009140 A2 1/2003
WO WO 2004/084083 A1 9/2004

OTHER PUBLICATIONS

Hillier, Andrew; "Transformational Analytics: Virtualizing IT Environments"; White Paper; Apr. 2008; CiRBA Inc.
"Virtualization: Architectural Considerations and Other Evaluation Criteria"; White Paper; 2005; VMware.
"Building the Virtualized Enterprise with VMware Infrastructure"; White Paper; 2006; VMware.
"A Blueprint for Better Management from the Desktop to the Data Center"; White Paper; Feb. 2007; Novell.

Matheson, Leigh; Search Report from corresponding PCT Application No. PCT/CA2008/001522; Dec. 3, 2008.

Tanenbaum, Andrew S. et al; Distributed Systems: Principles and Paradigms; US Ed edition; Jan. 15, 2002; pp. 22-42, 326-336; Prentice Hall.

Hillier, Andrew; "A Quantitative and Analytical Approach to Server Consolidation" dated Jan. 2006, published at least as early as Feb. 3, 2006; CiRBA Inc.; Technical Whitepaper.

Hillier, Andrew; "Data Center Intelligence" dated Mar. 2006, published at least as early as Apr. 1, 2006; CiRBA Inc.; Technical Whitepaper.

Spellman, Amy et al.; "Server Consolidation Using Performance Modeling"; IT Professional; Sep./Oct. 2003; pp. 31-36; vol. 5, No. 5.

* cited by examiner




Figure 1




Figure 2




Figure 3(a)

Figure 3(b)




Figure 4

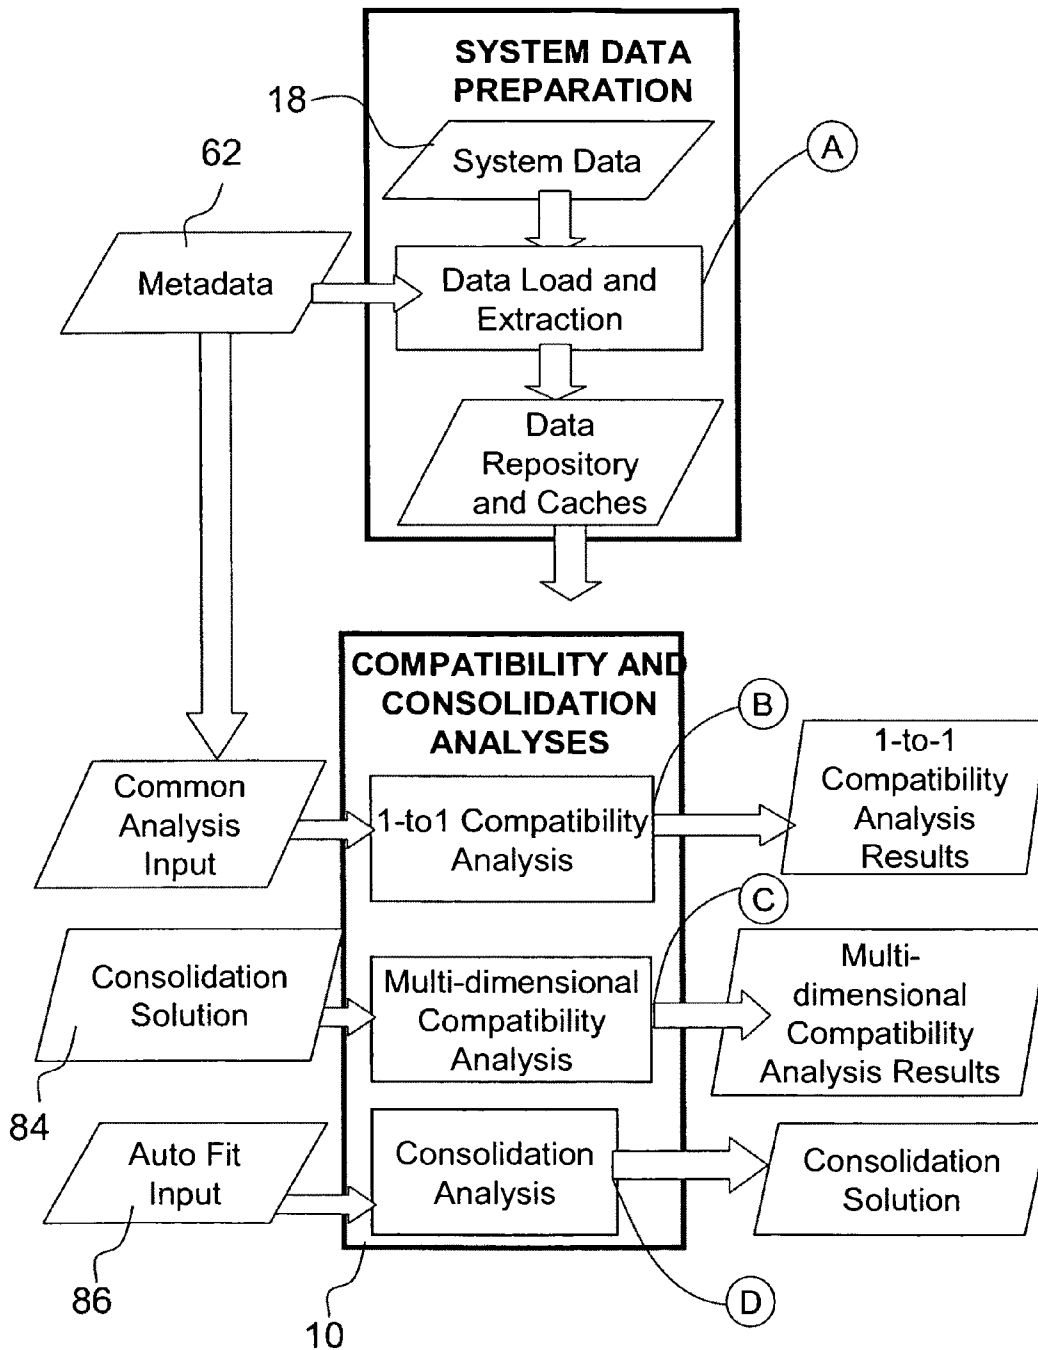


Figure 5

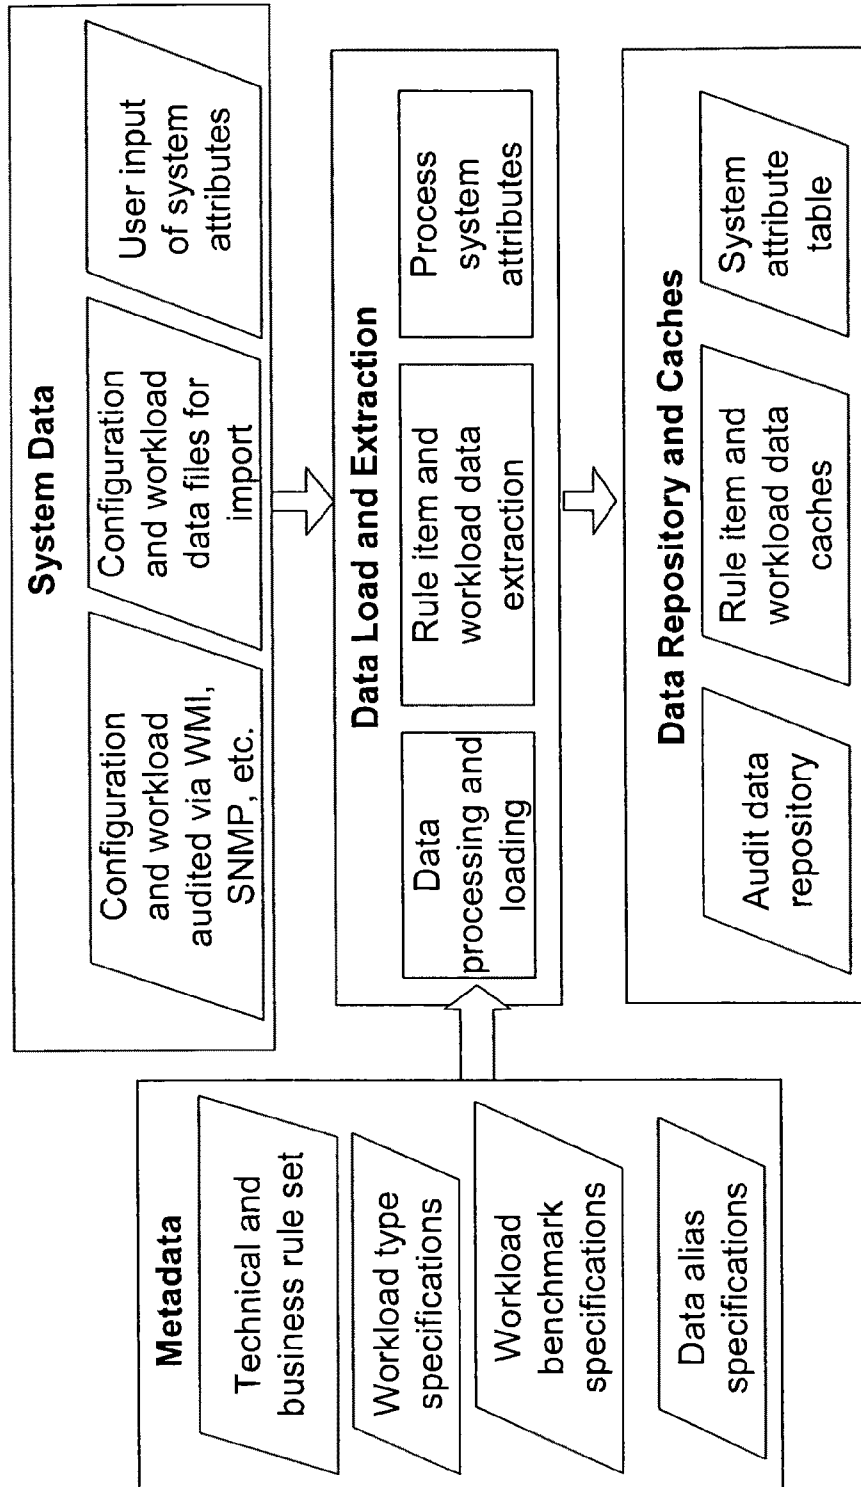


Figure 6




Figure 7

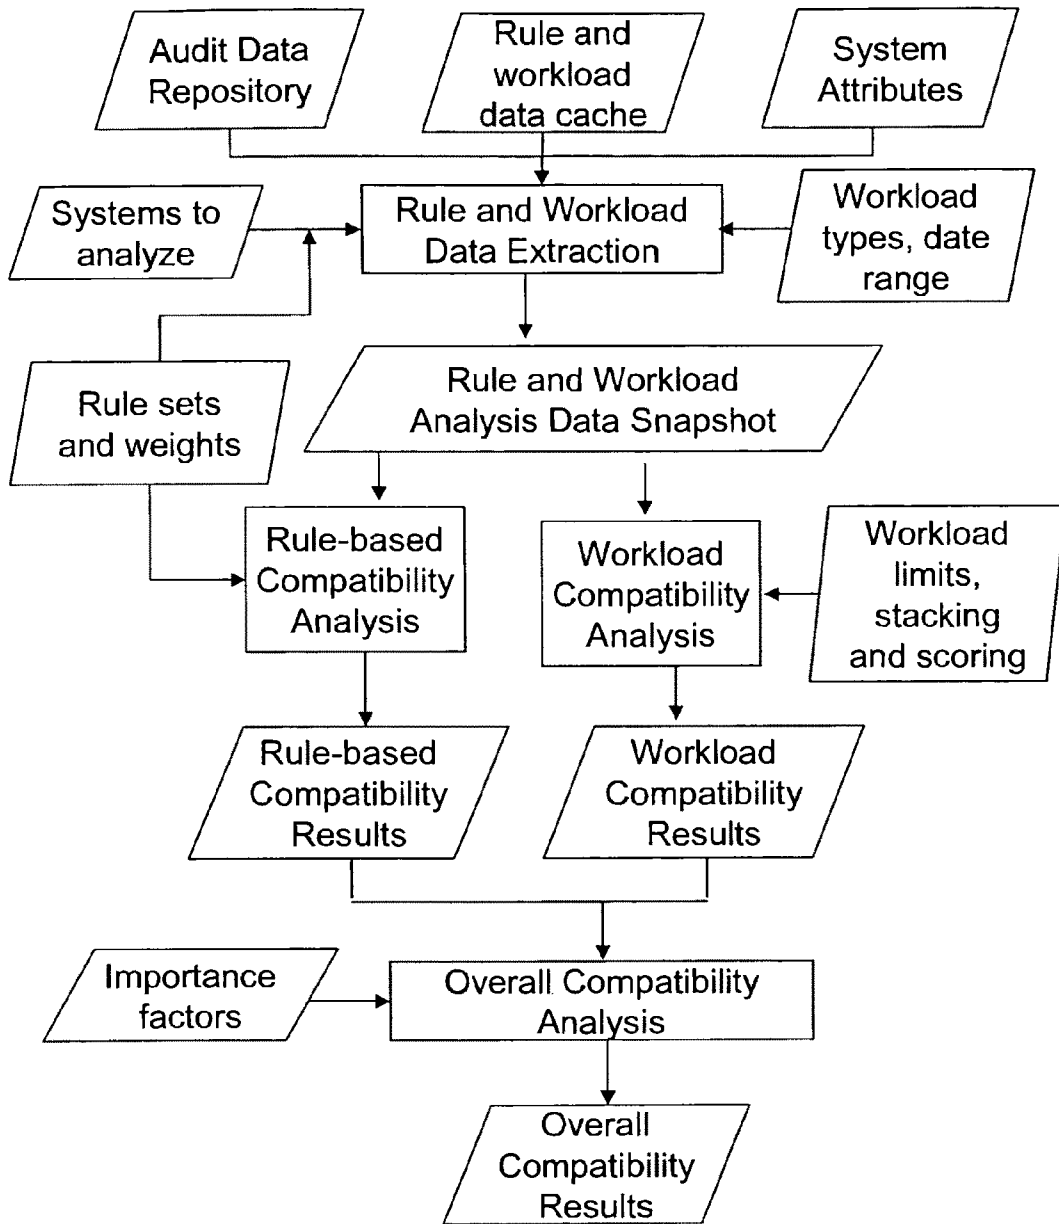


Figure 8

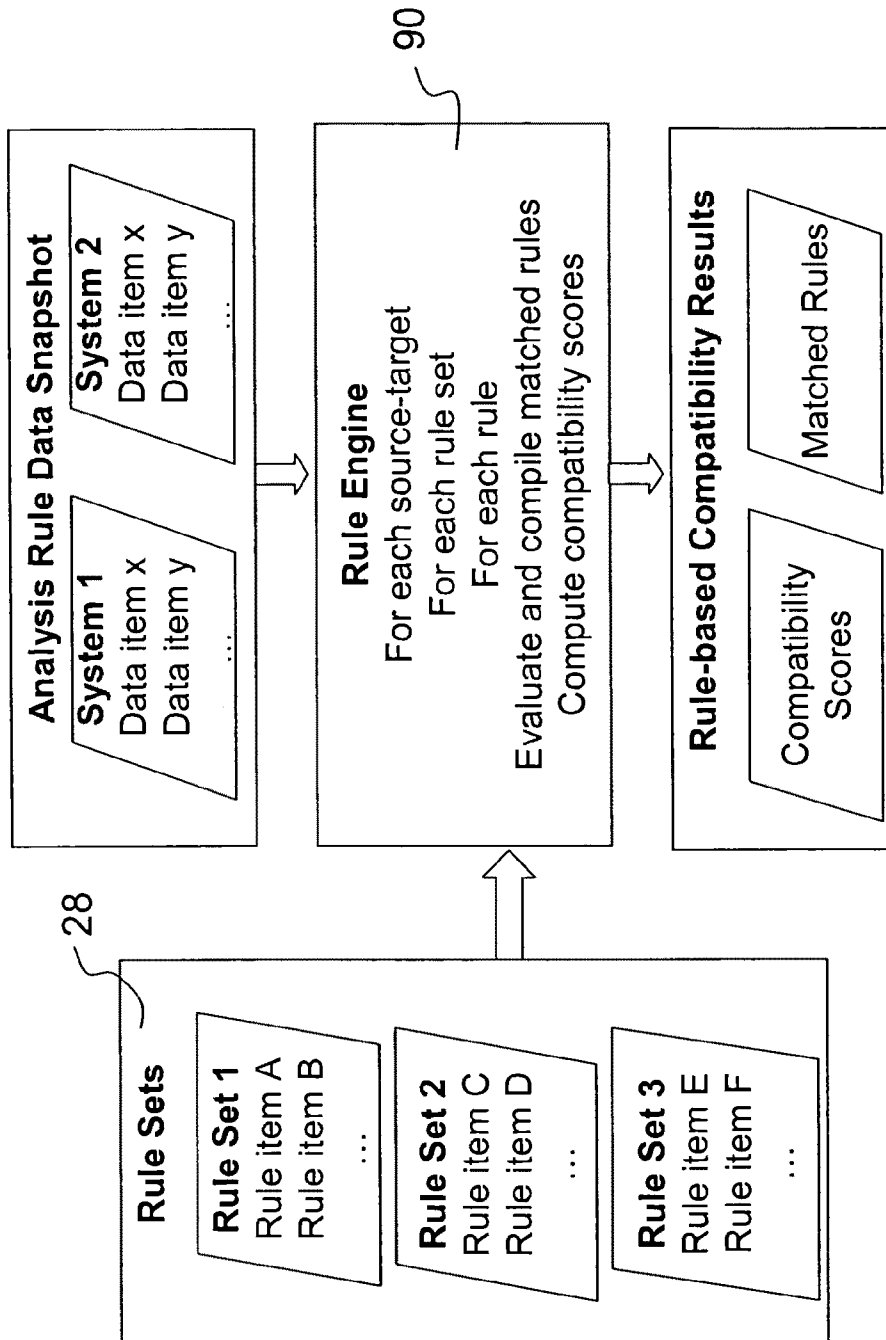
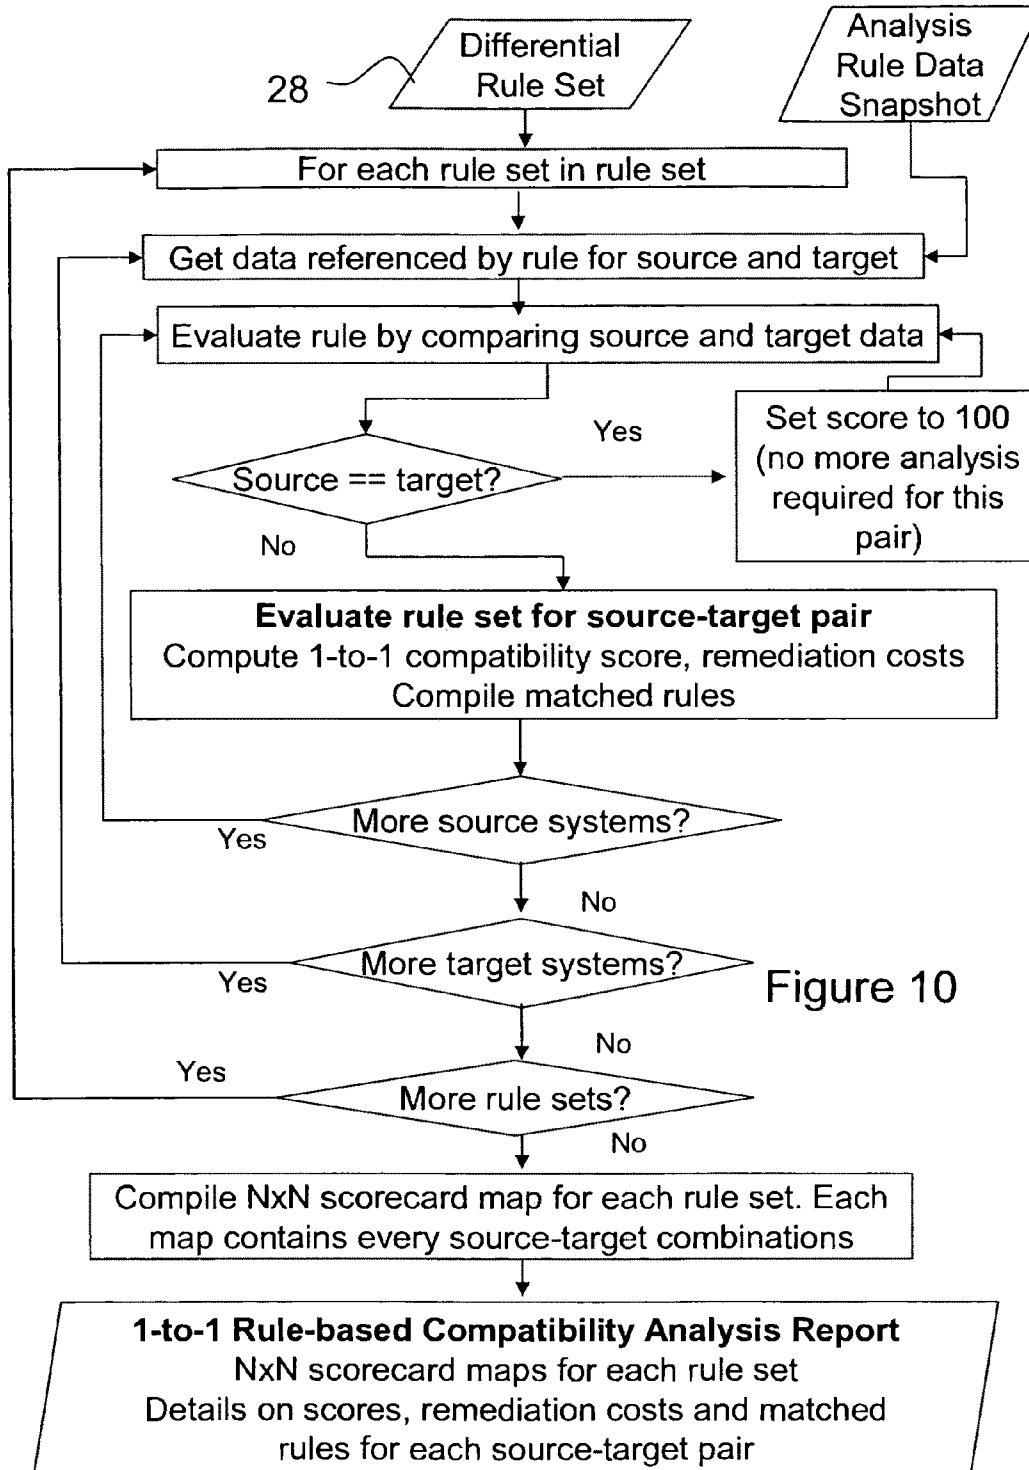


Figure 9



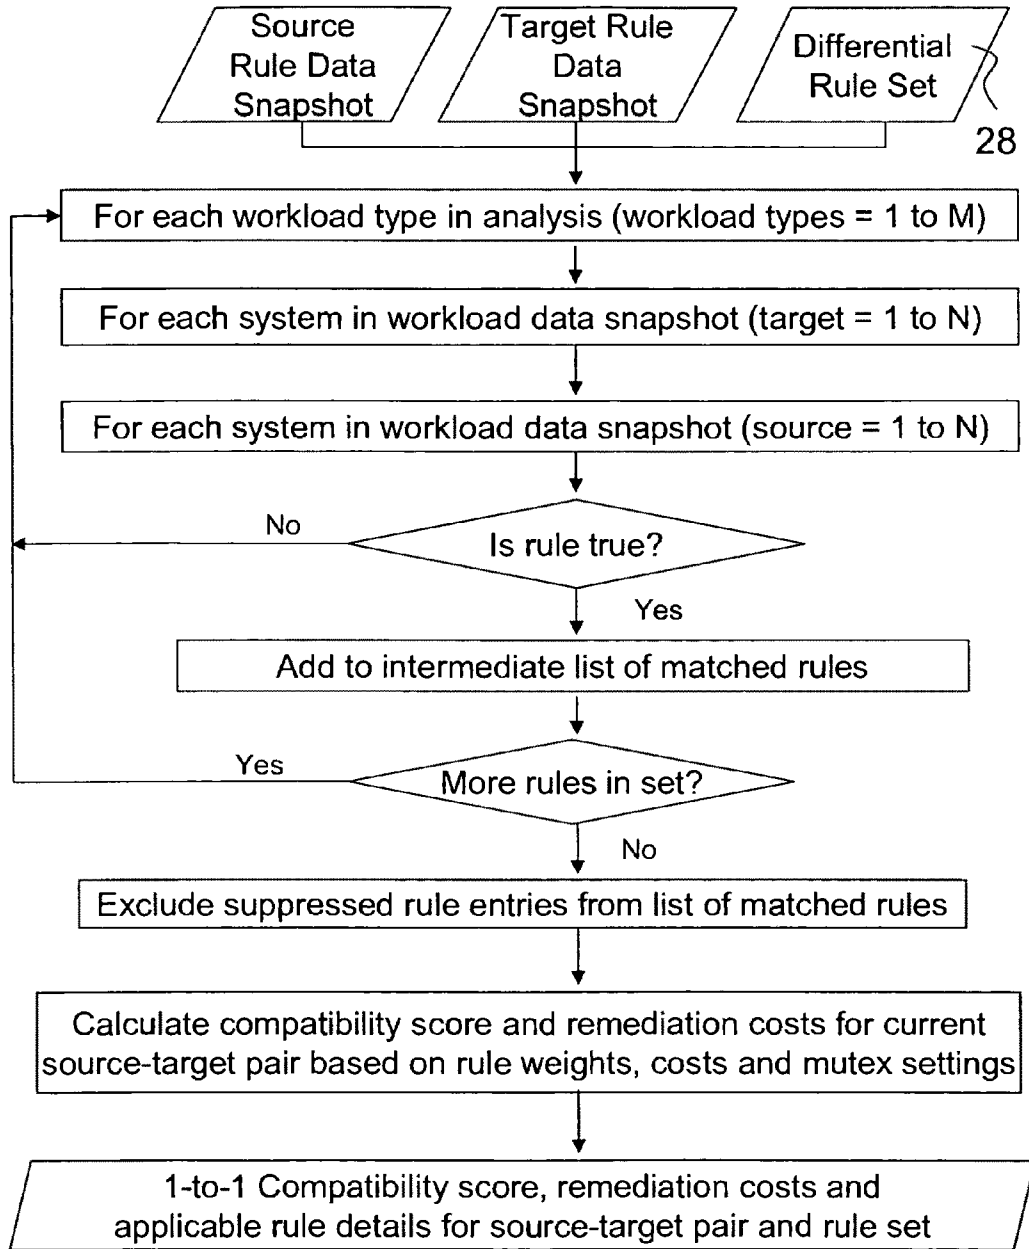


Figure 11

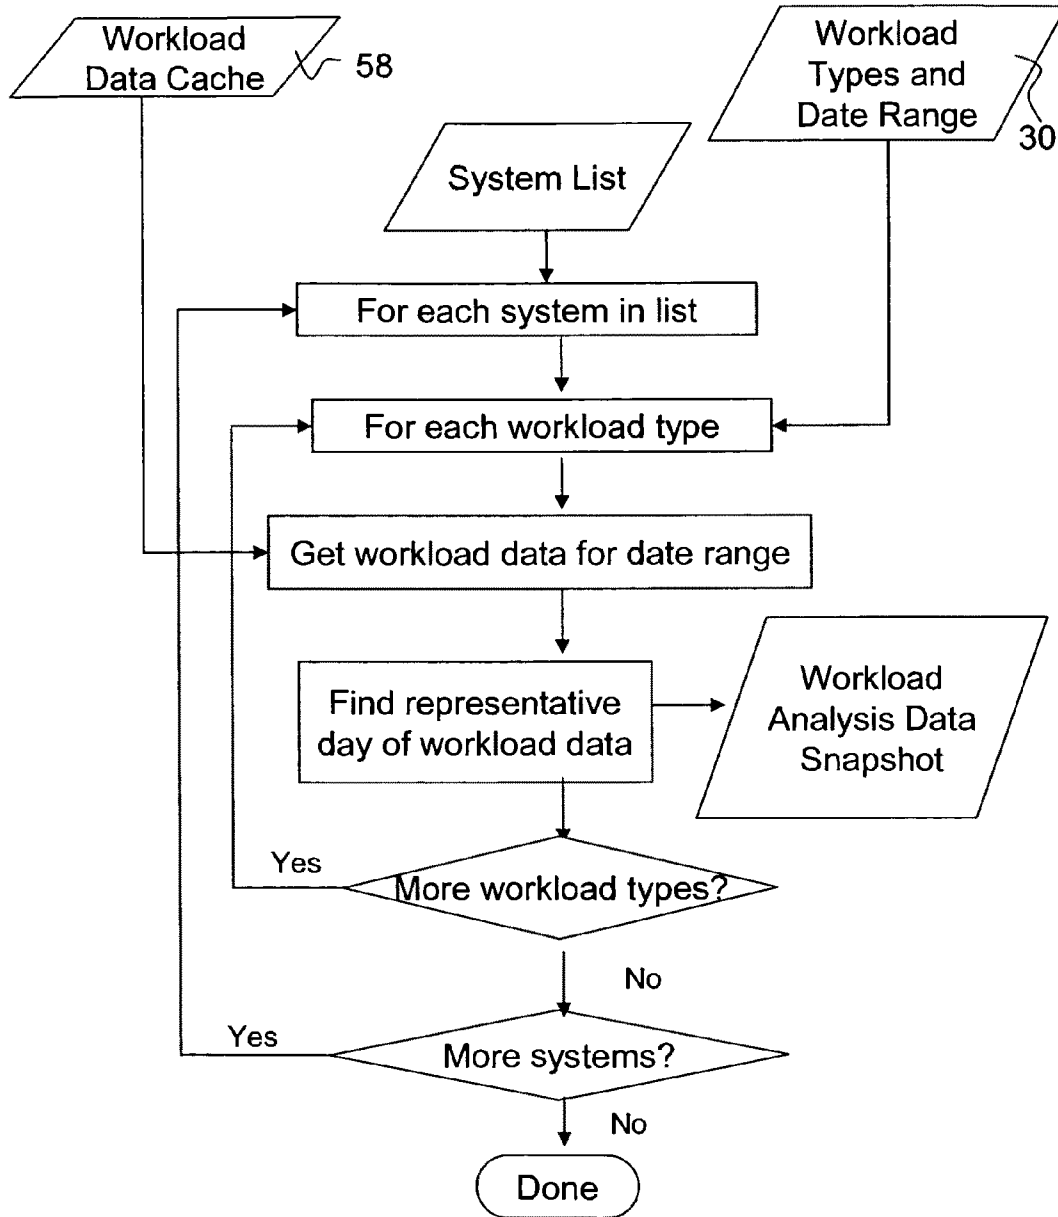


Figure 12

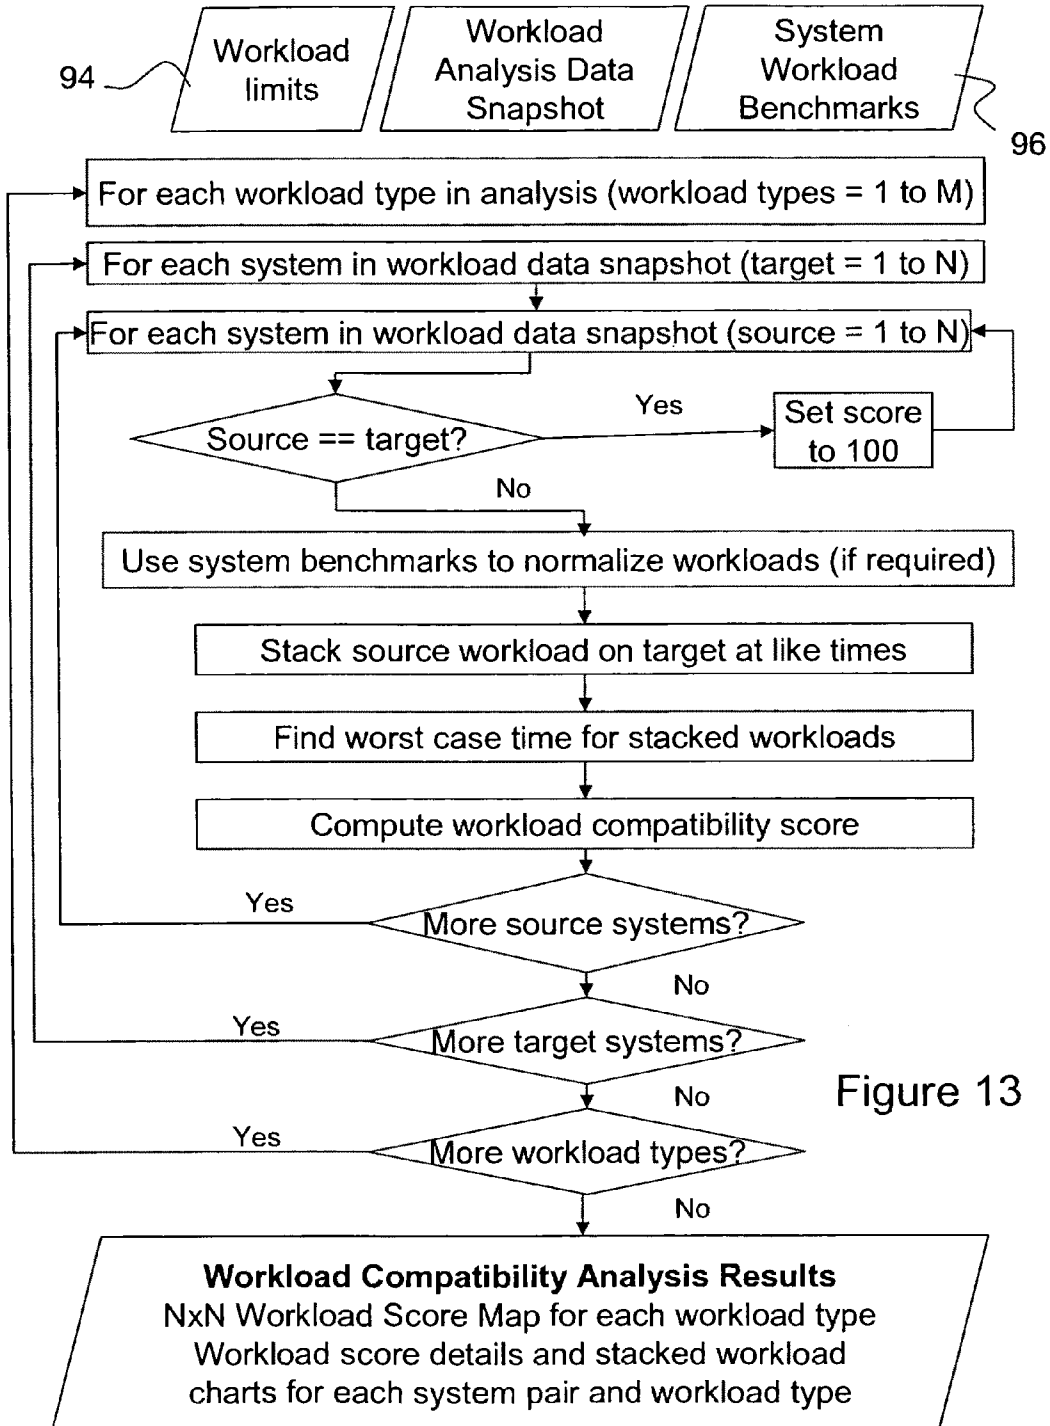


Figure 13

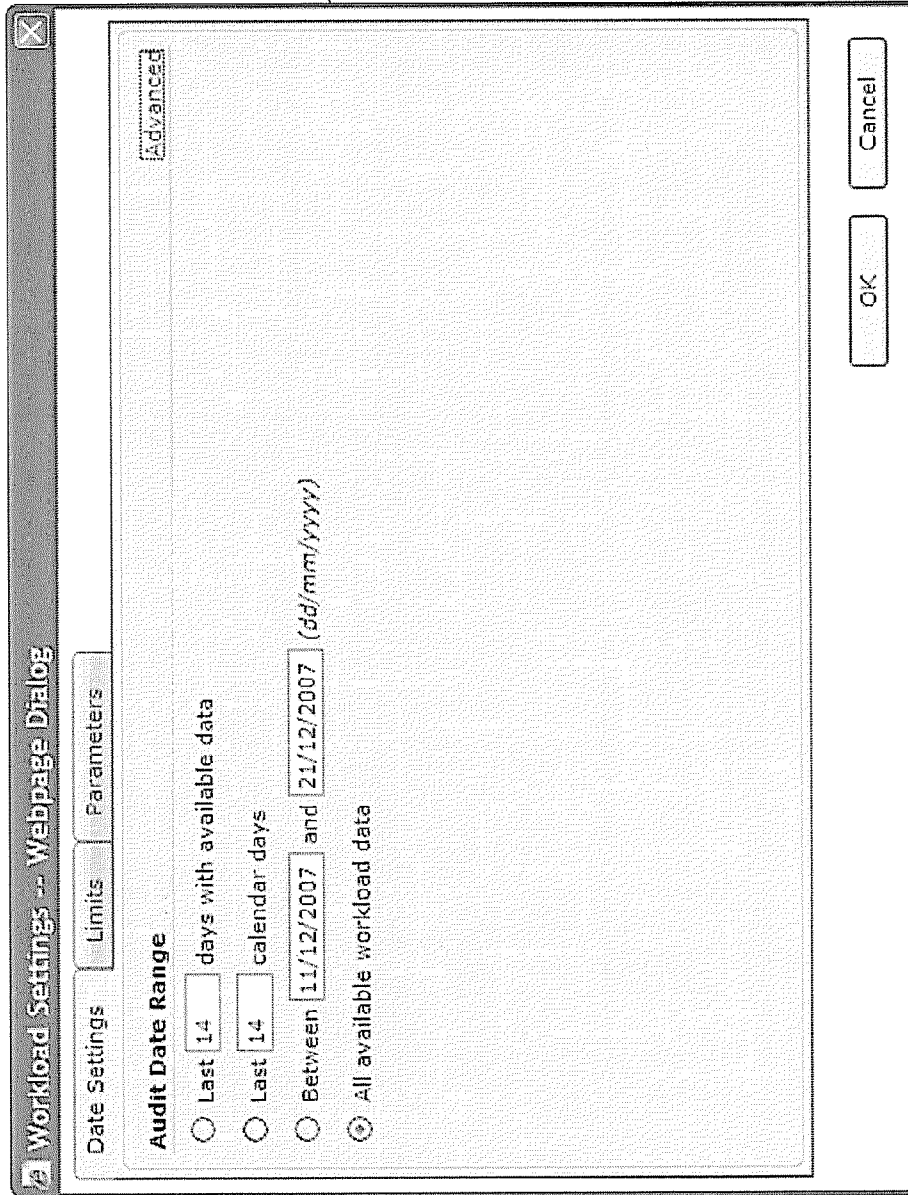


Figure 14

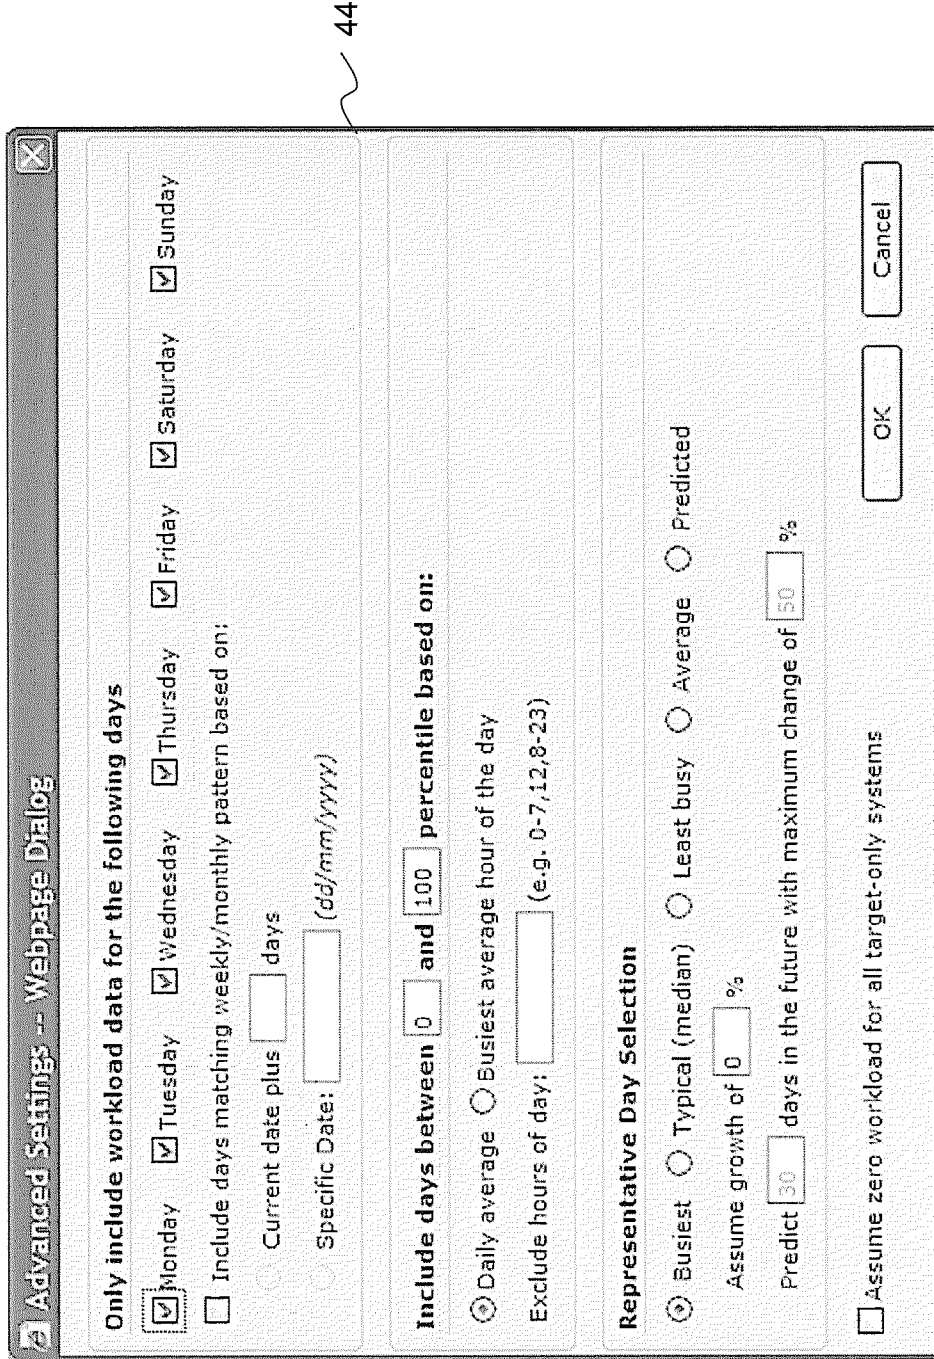


Figure 15




Figure 16

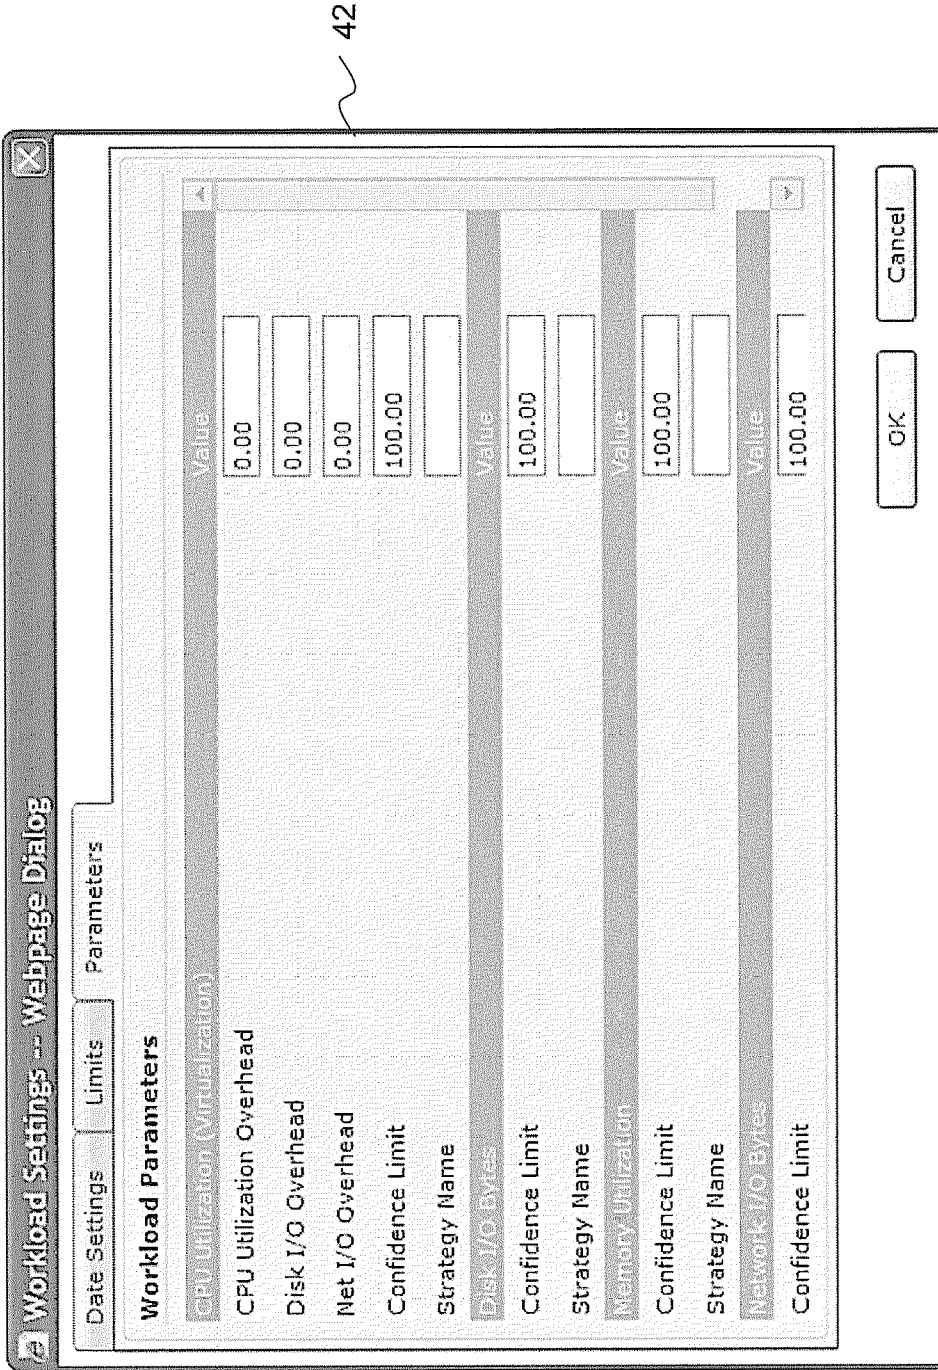


Figure 17

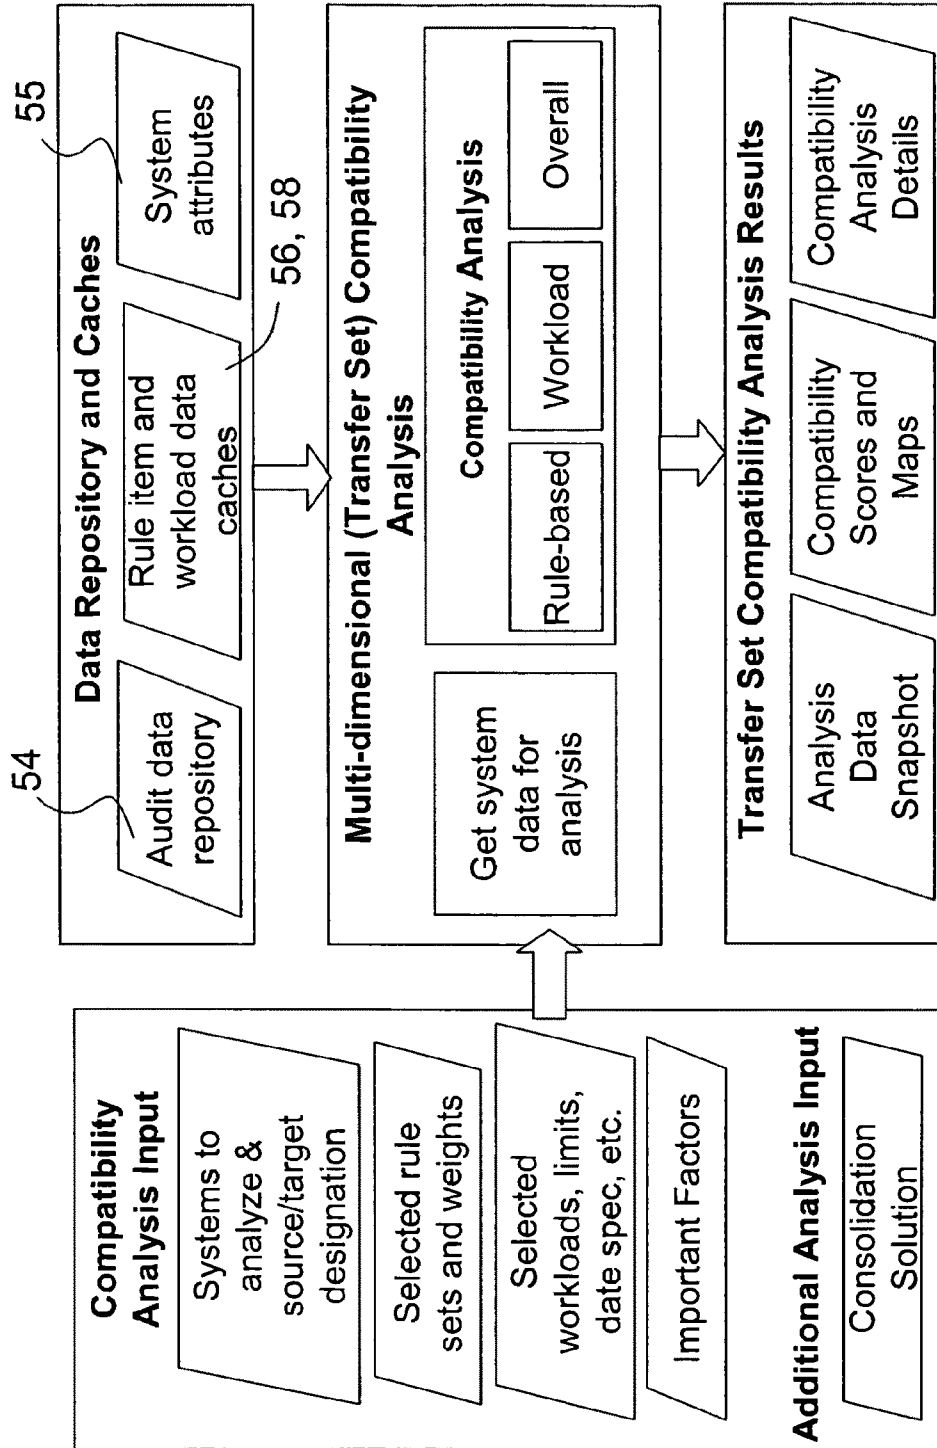


Figure 18

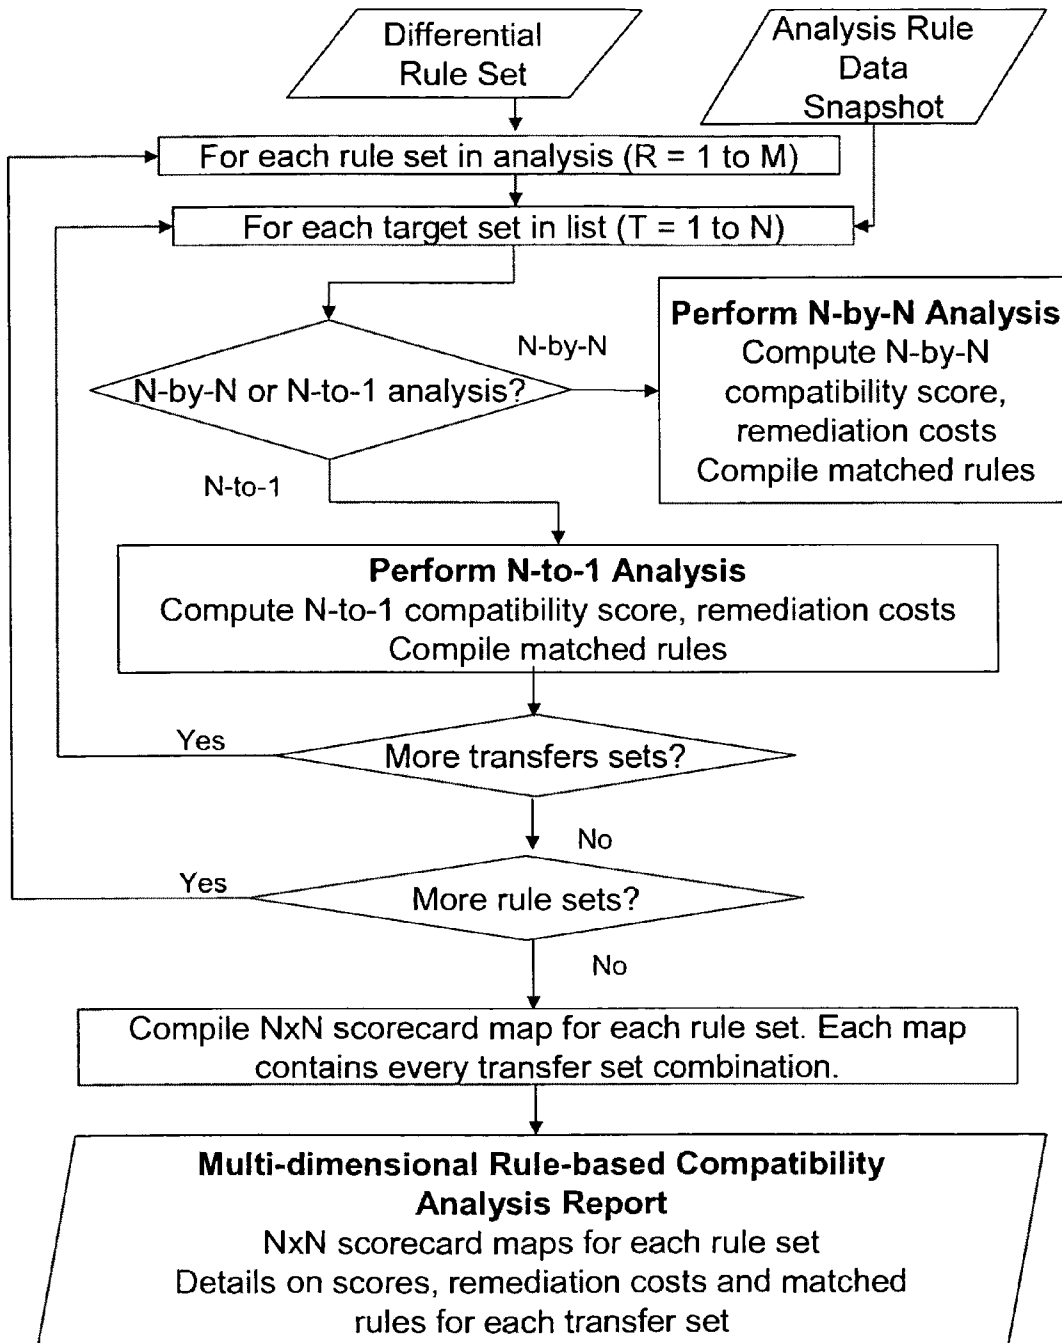


Figure 19

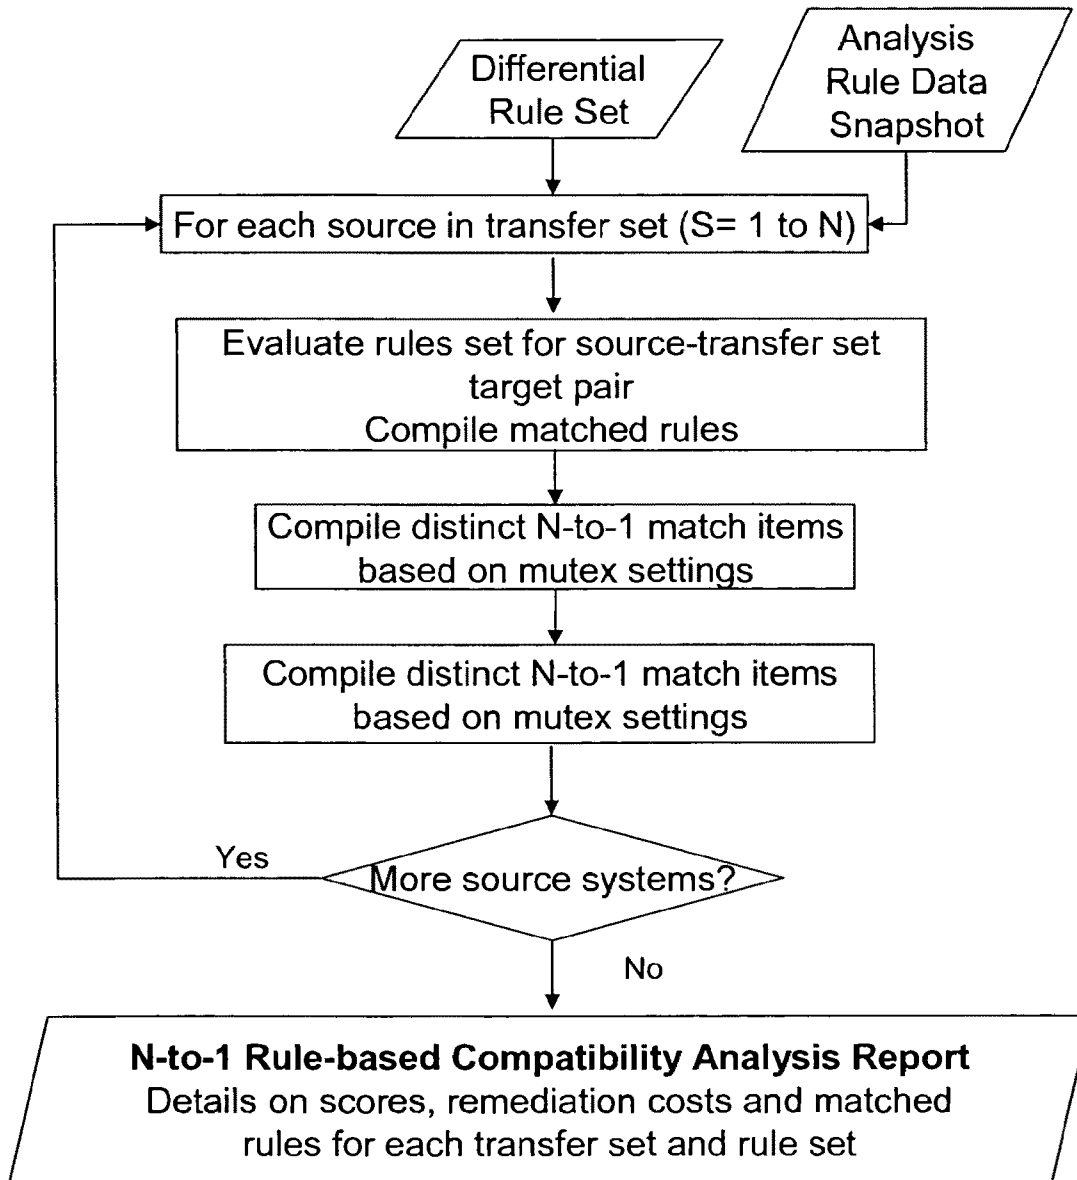


Figure 20

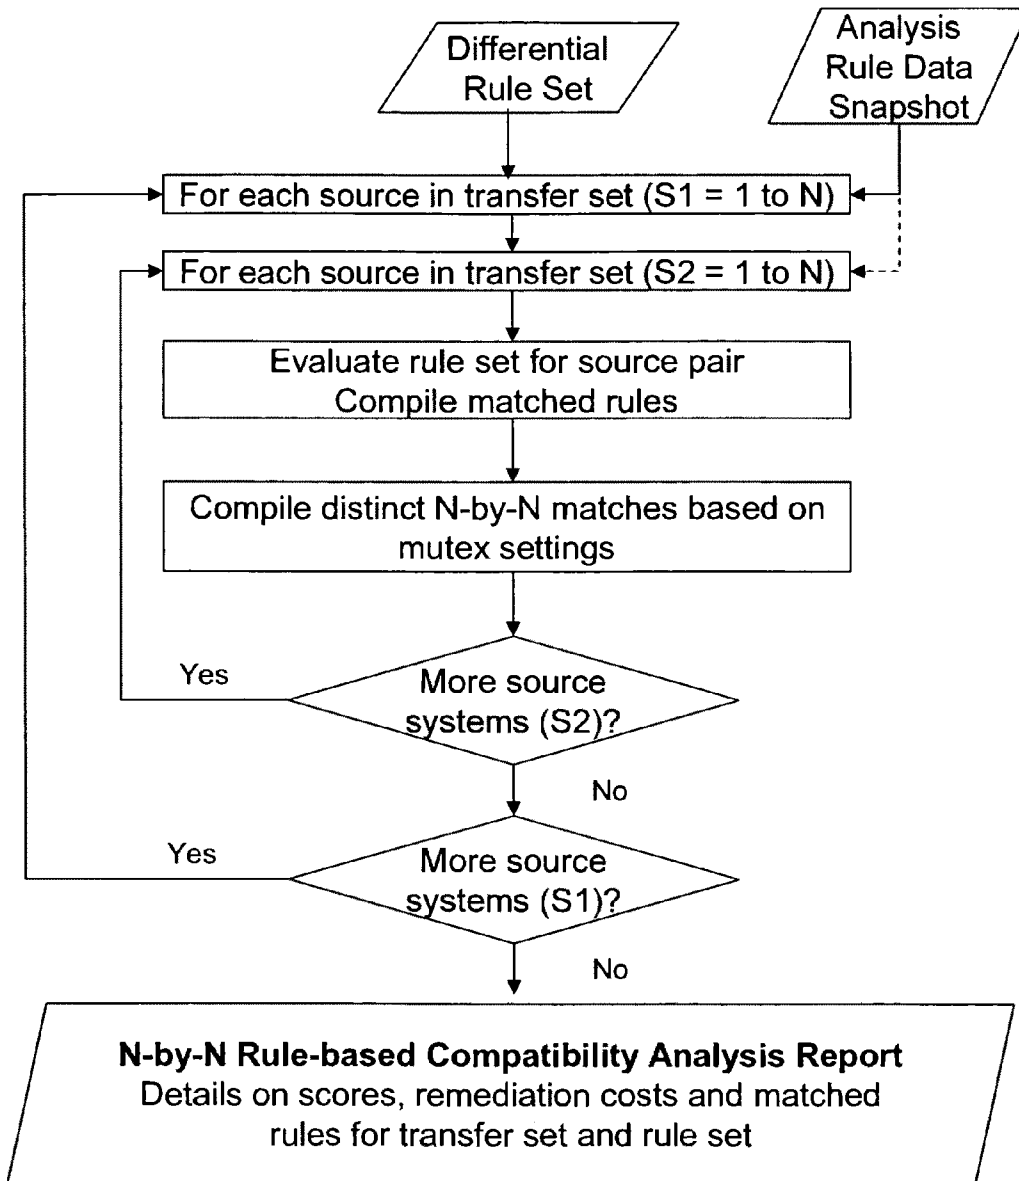


Figure 21

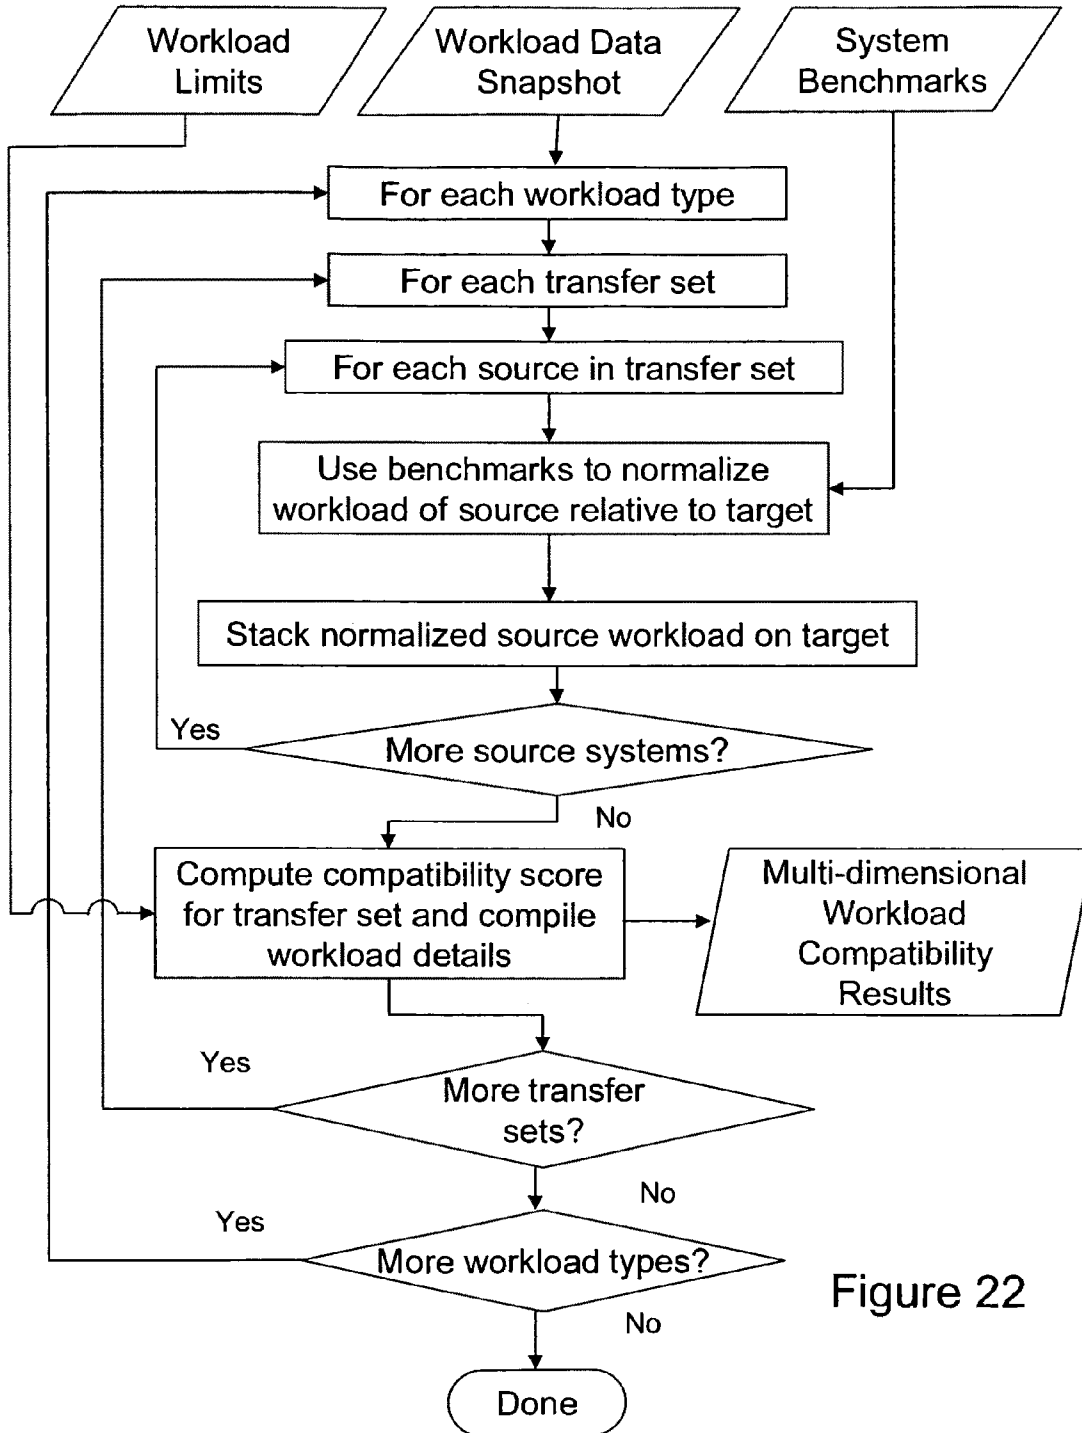


Figure 22

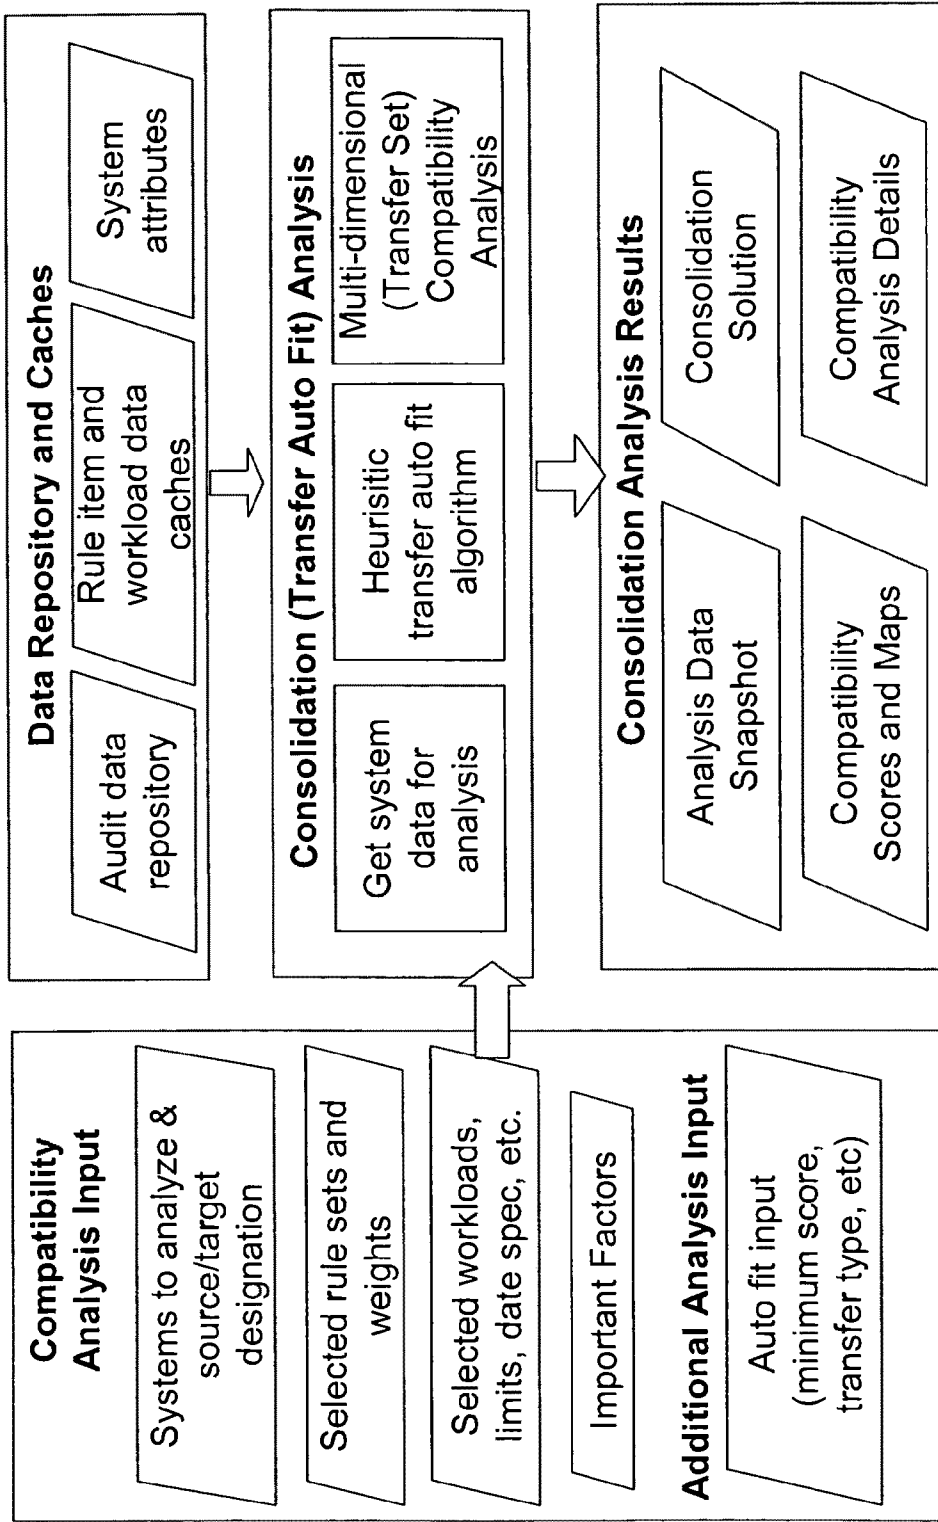


Figure 23

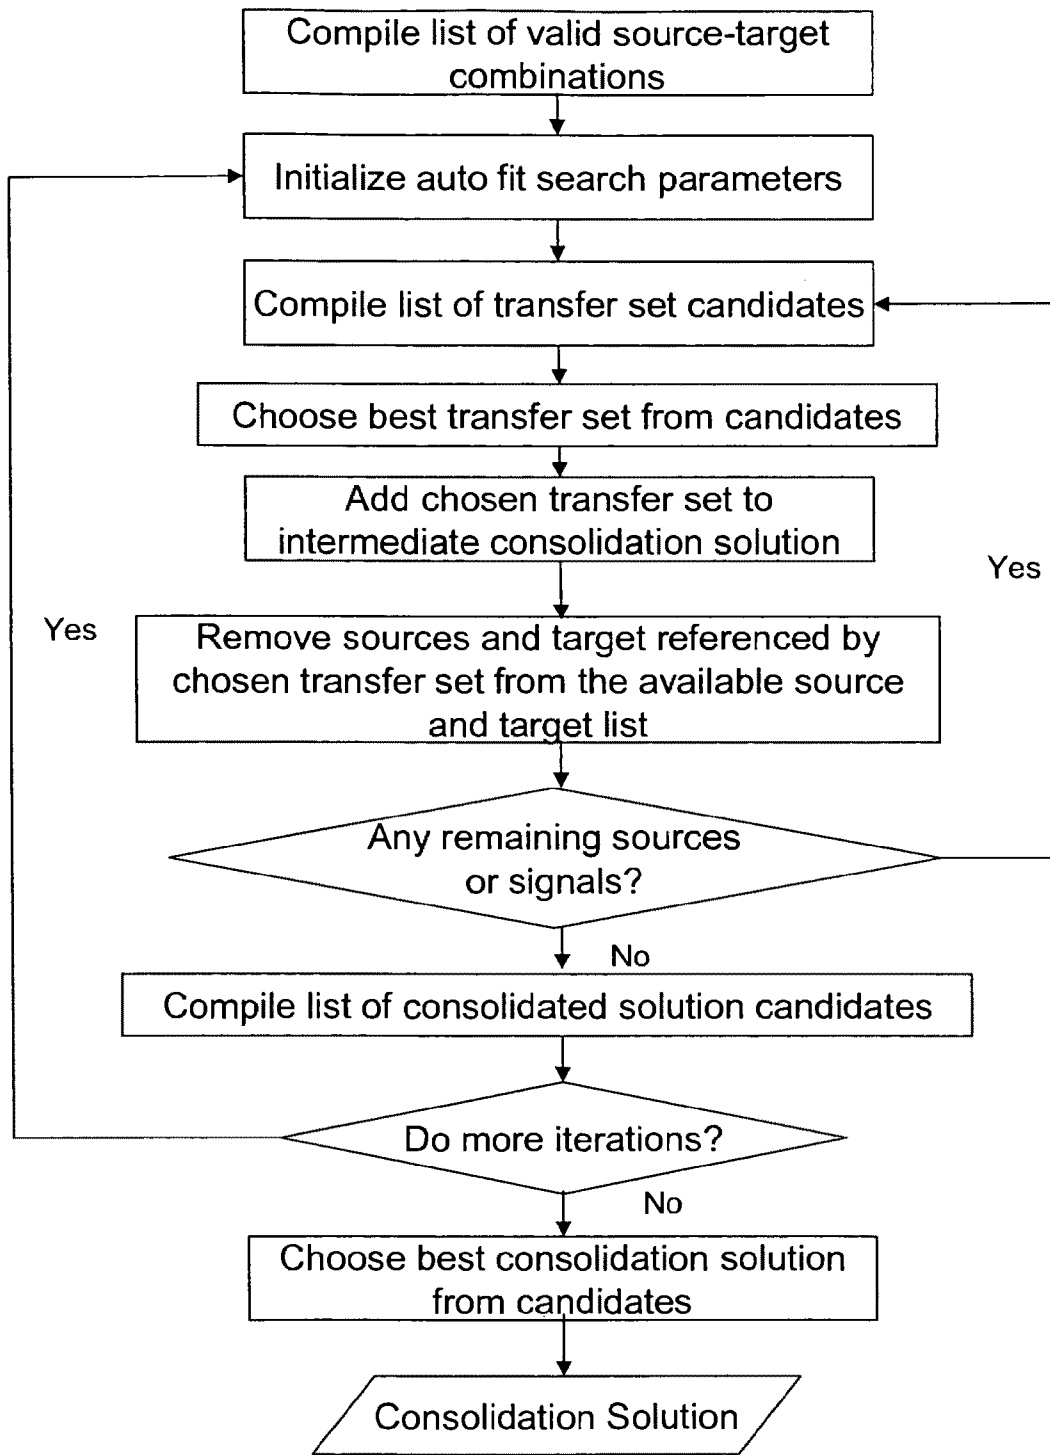


Figure 24

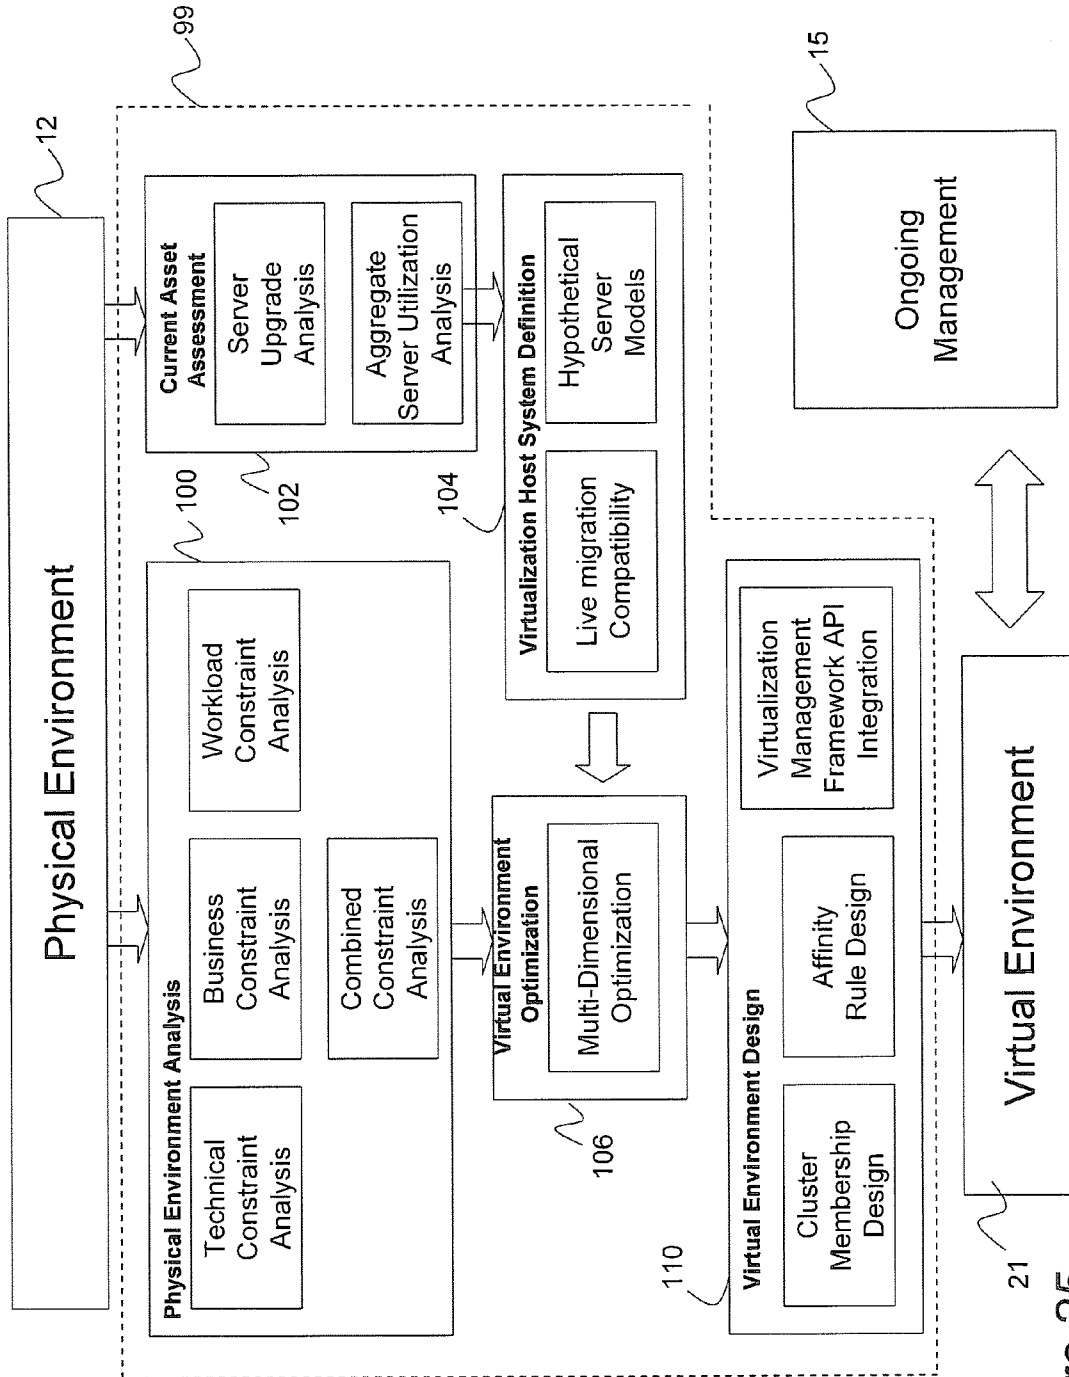


Figure 25

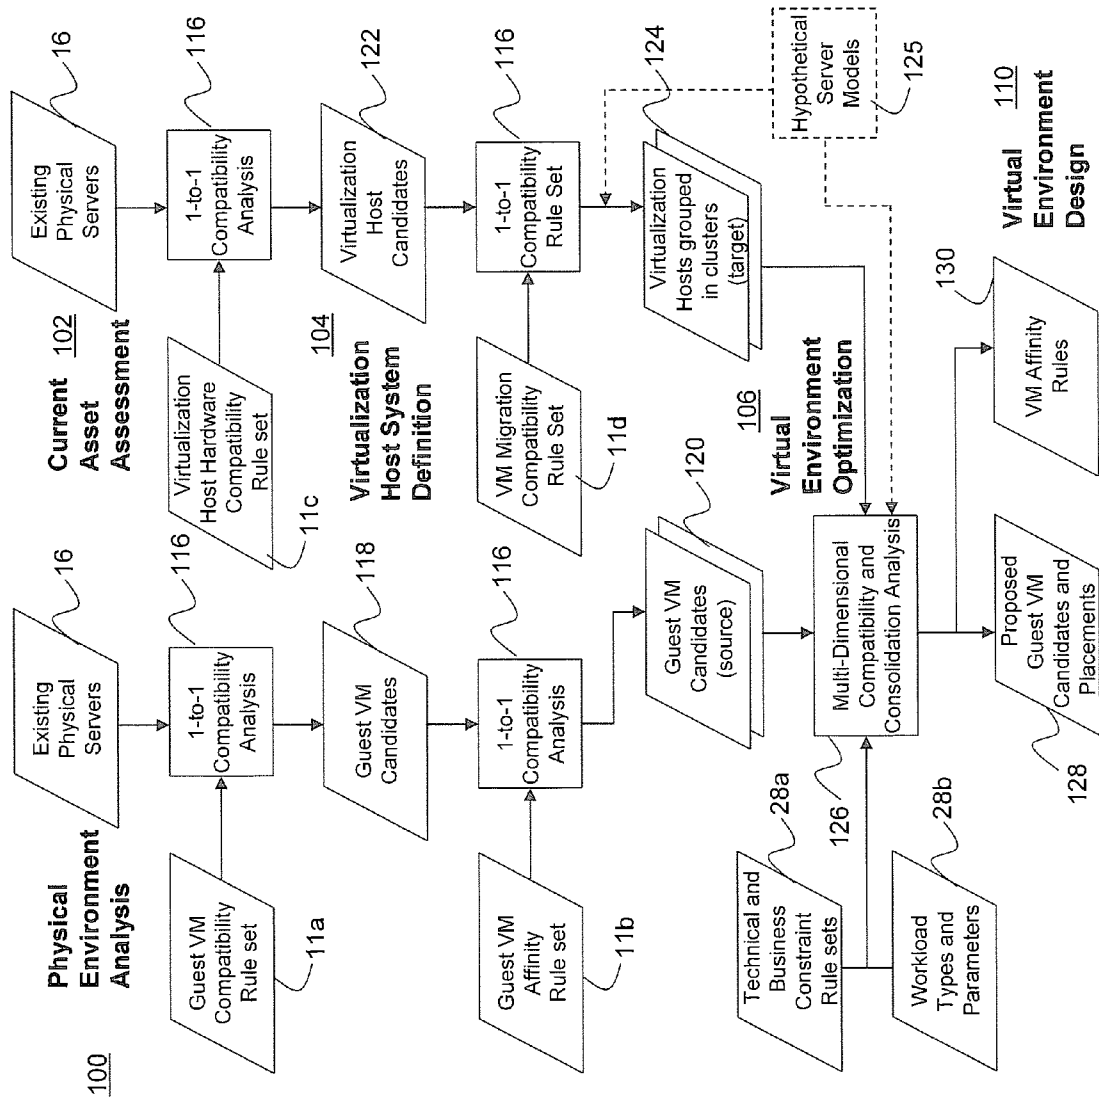


Figure 26

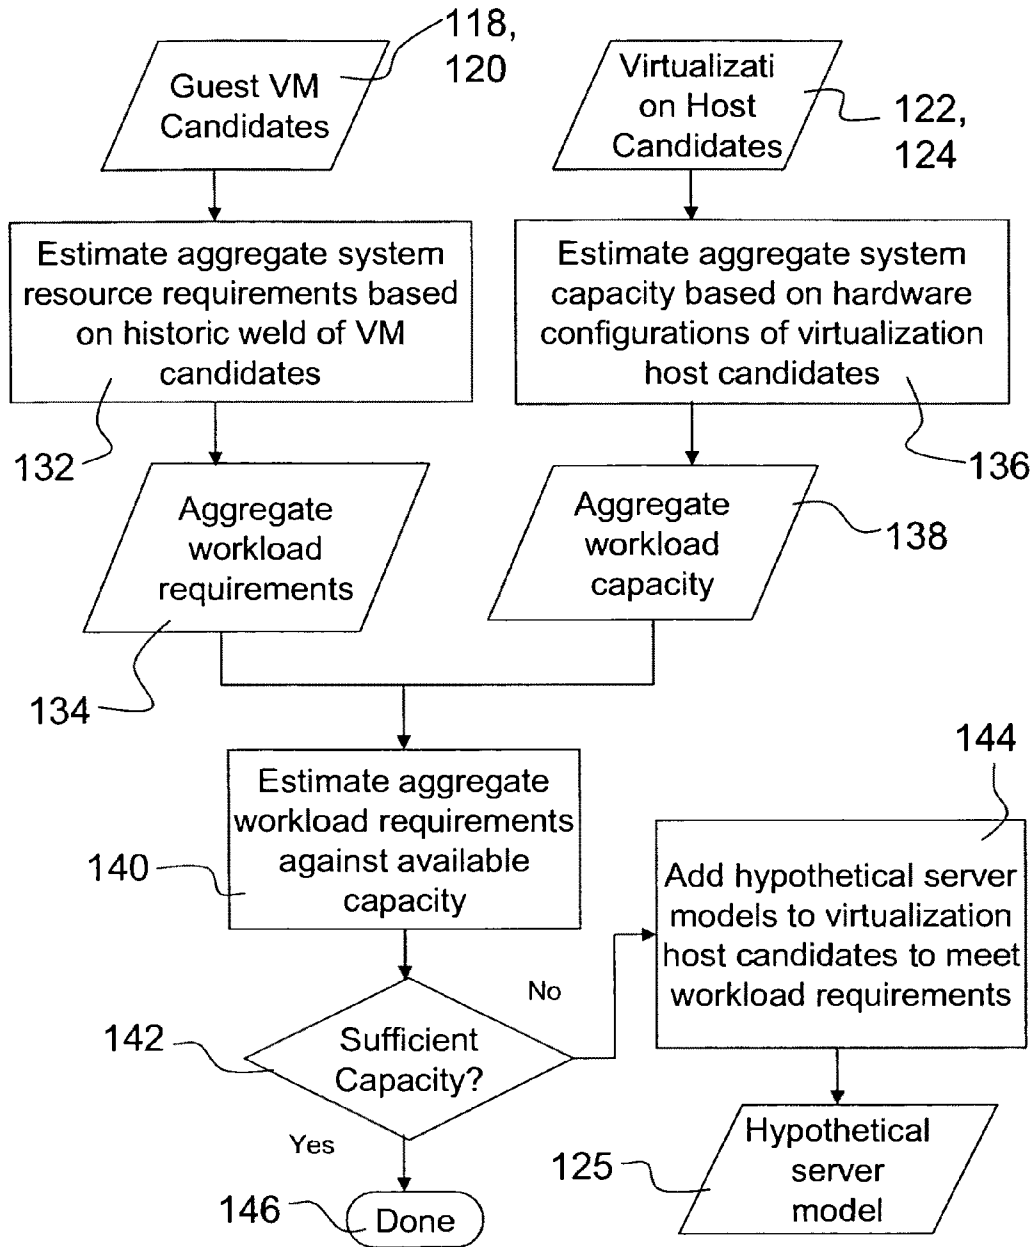


Figure 27

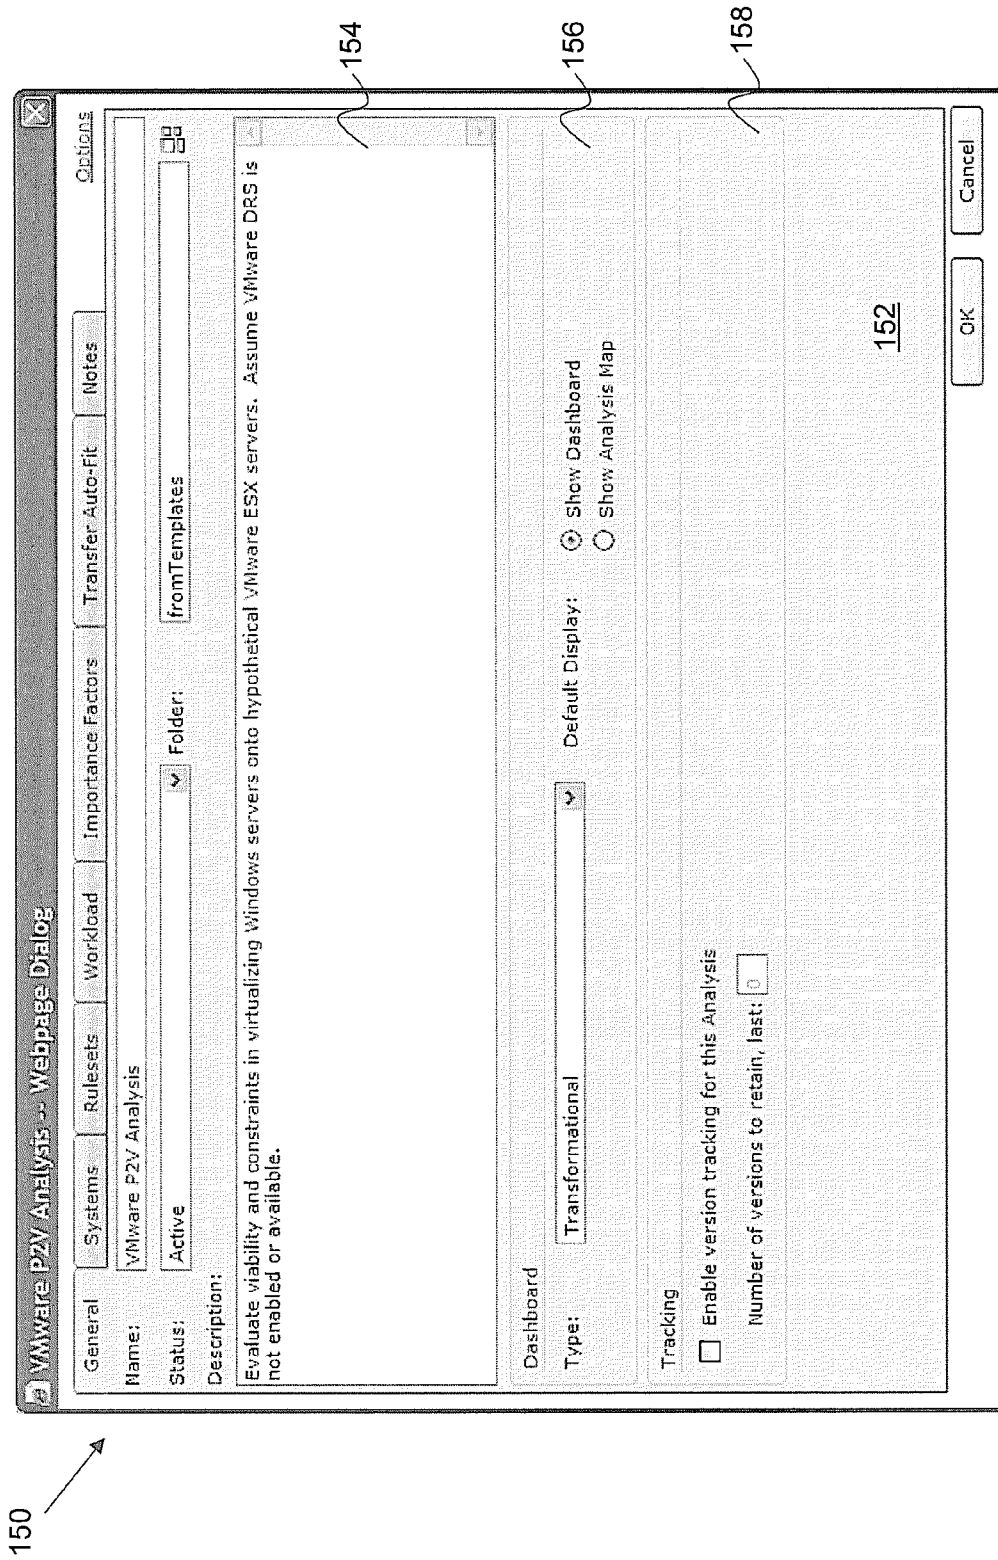


Figure 28

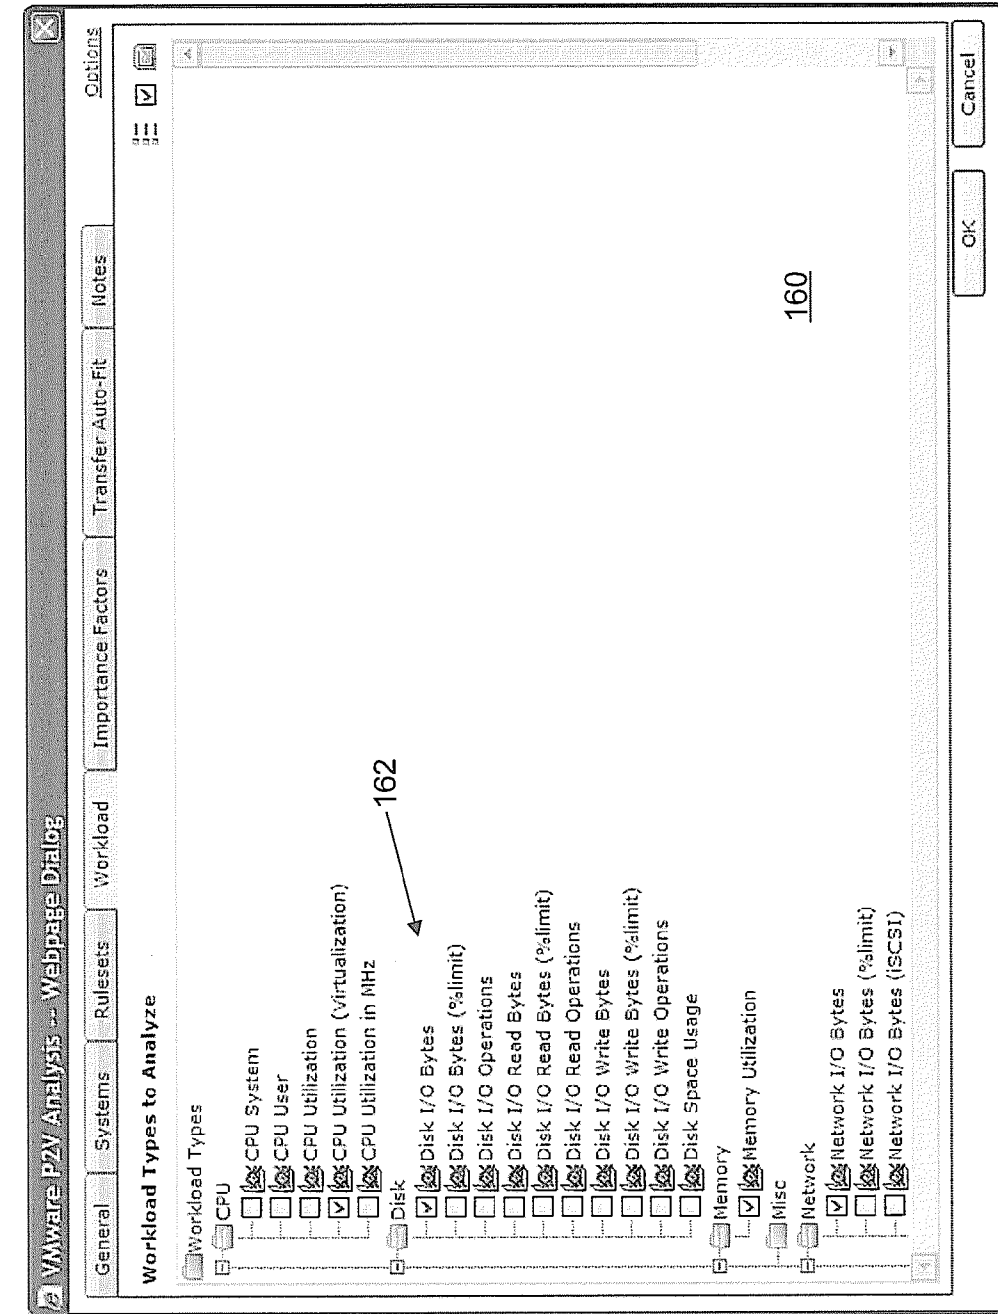
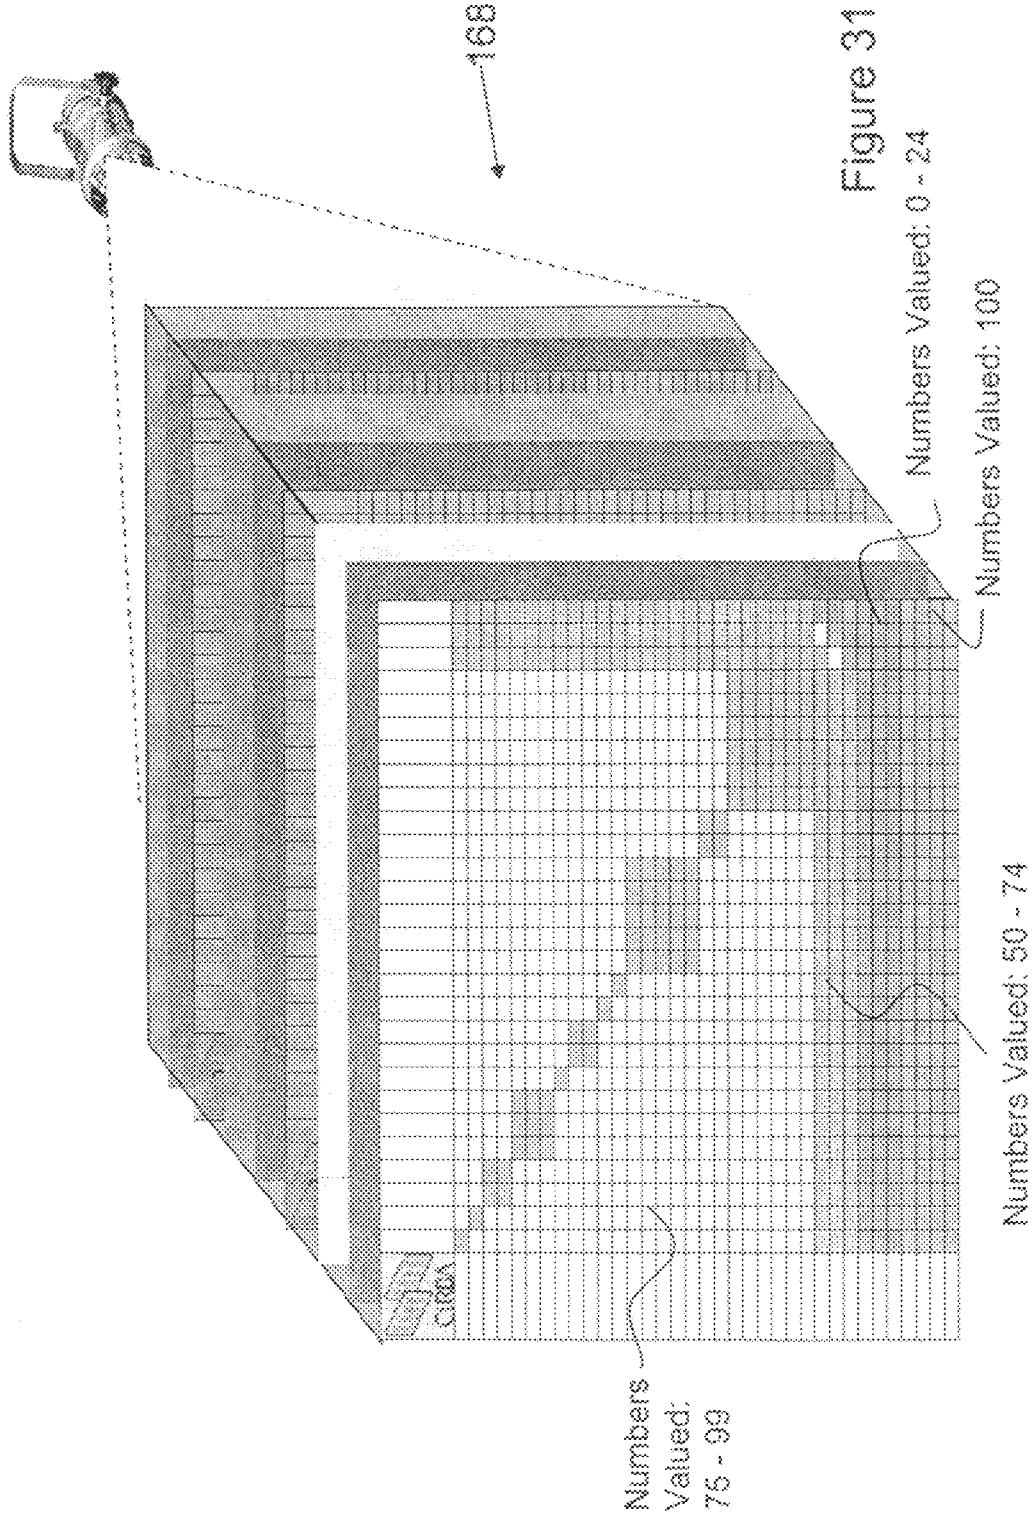


Figure 29

The figure displays a grid of patent classification codes (e.g., G06F 0013 00 00) grouped into columns labeled with patent numbers. A circular highlight is centered on a group of codes in the right-hand section of the grid. The label 'Figure 30' is located to the right of the grid, and the number '166' is positioned at the top right corner.

Figure 30 166 164



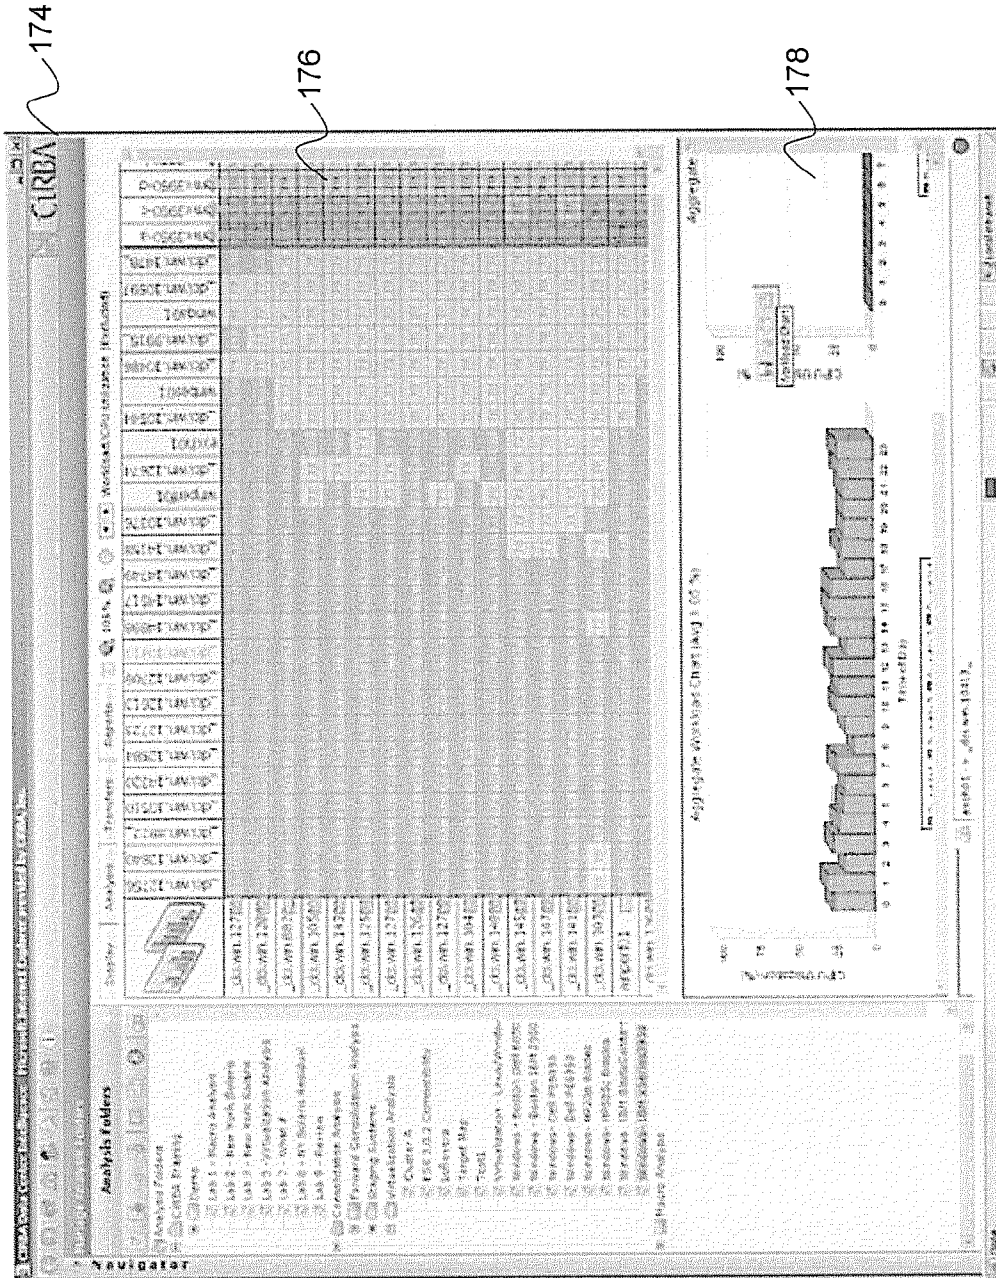


Figure 33

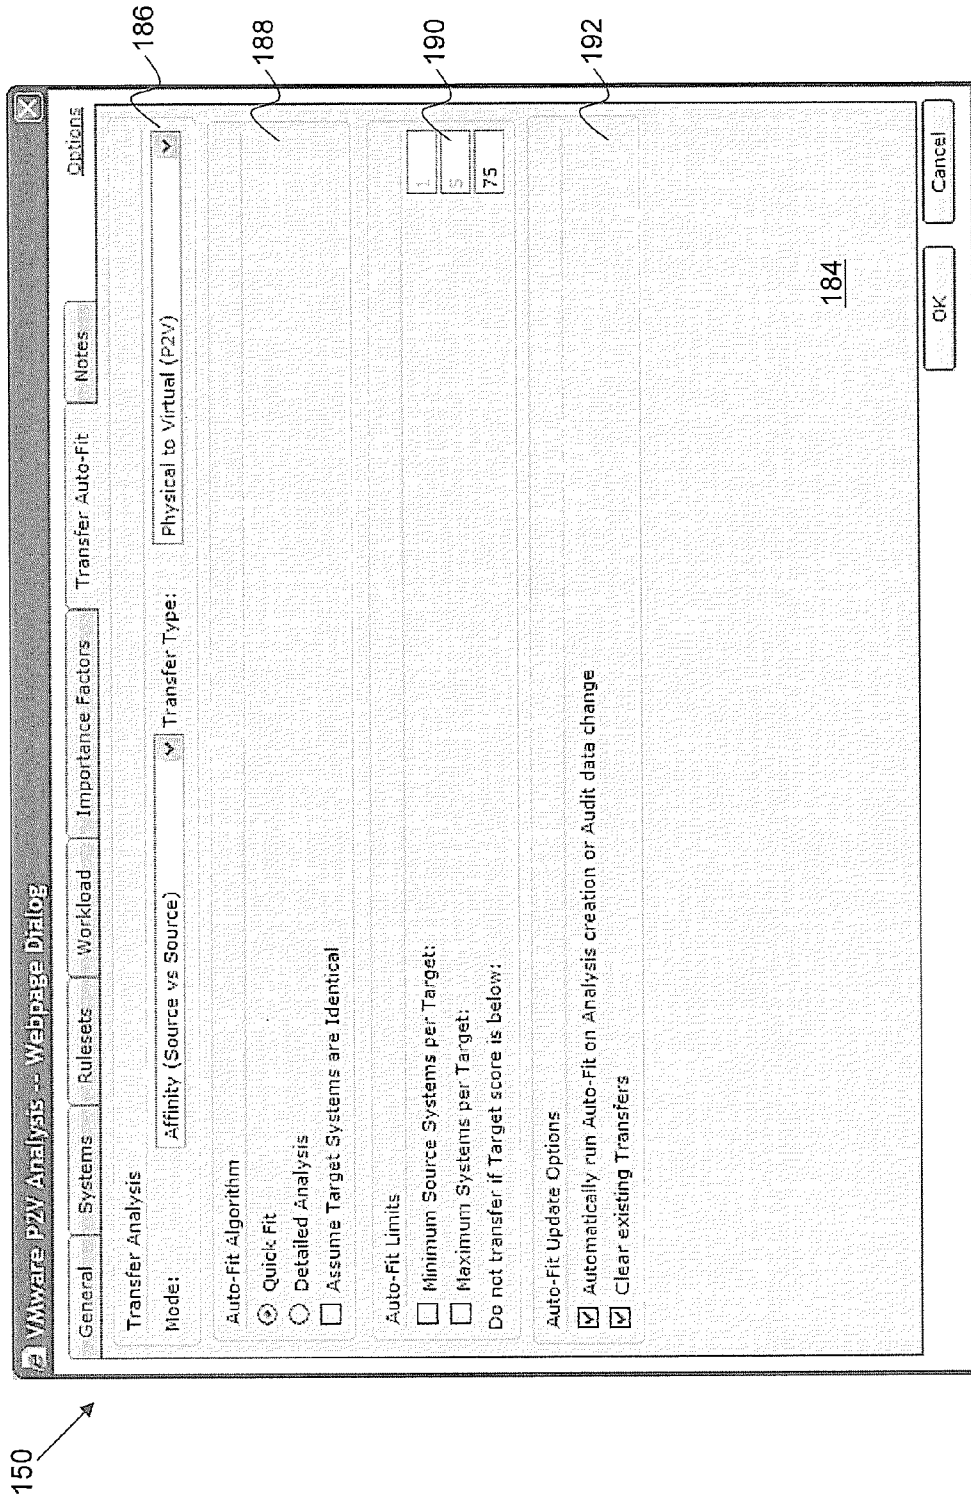


Figure 35

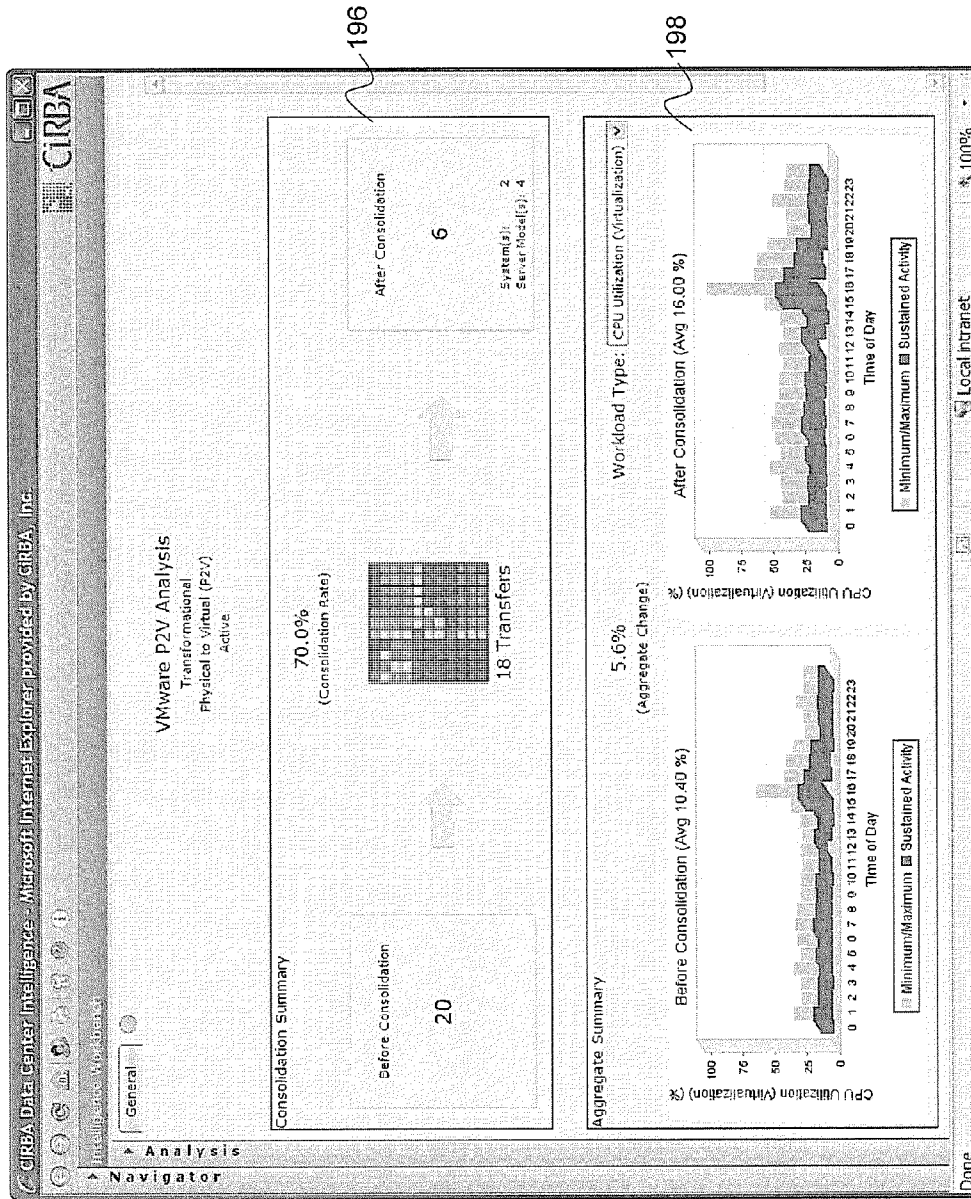
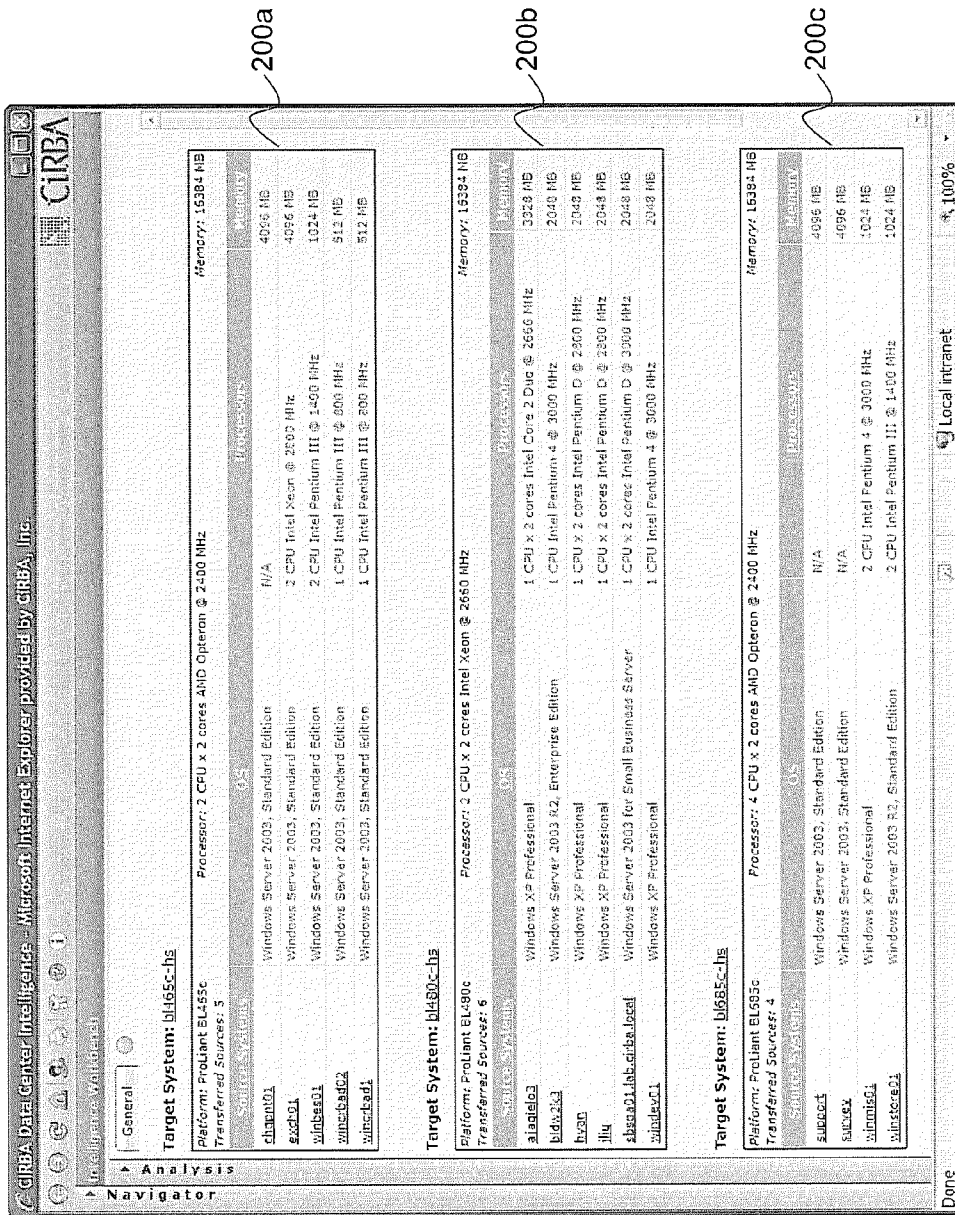


Figure 36



194

Figure 37




Figure 39




Figure 40

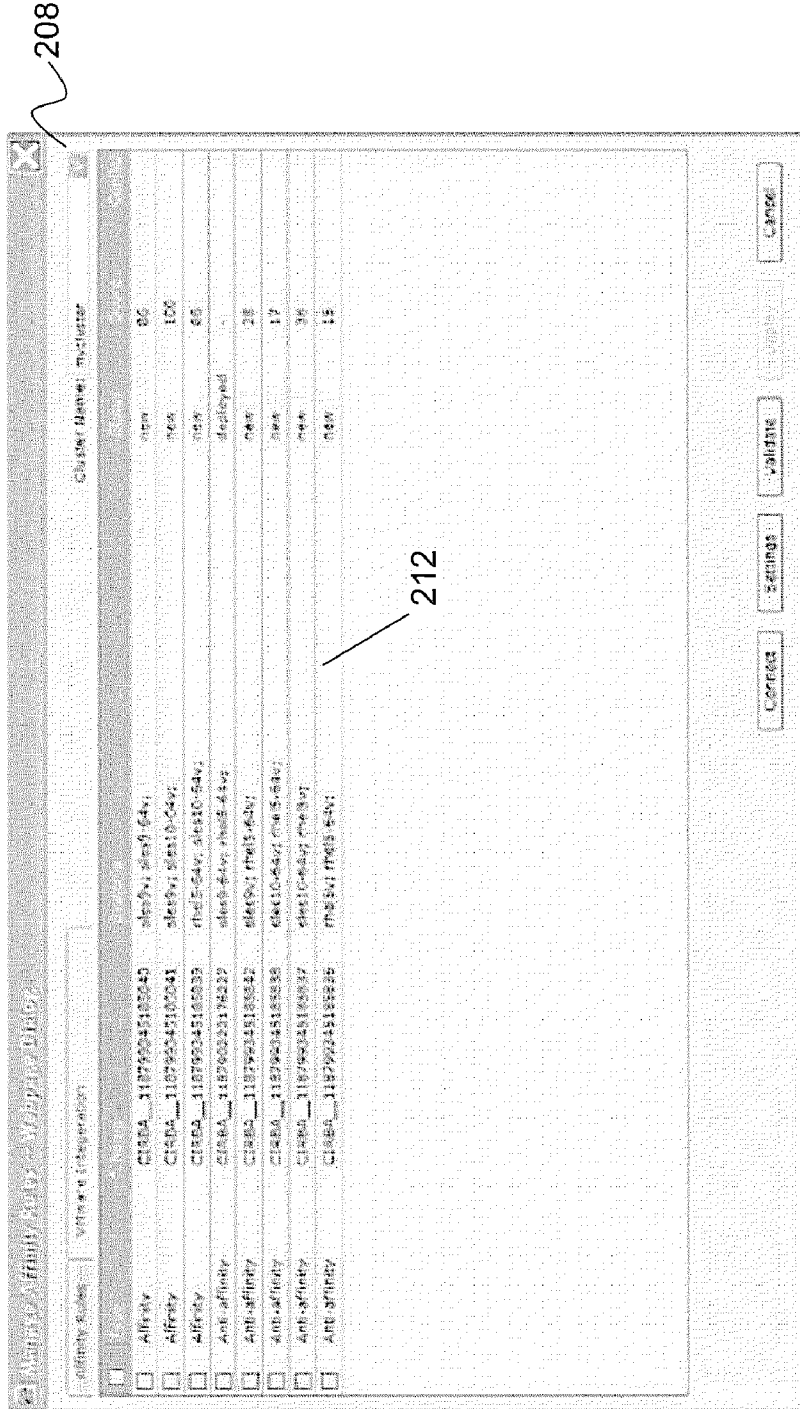


Figure 41

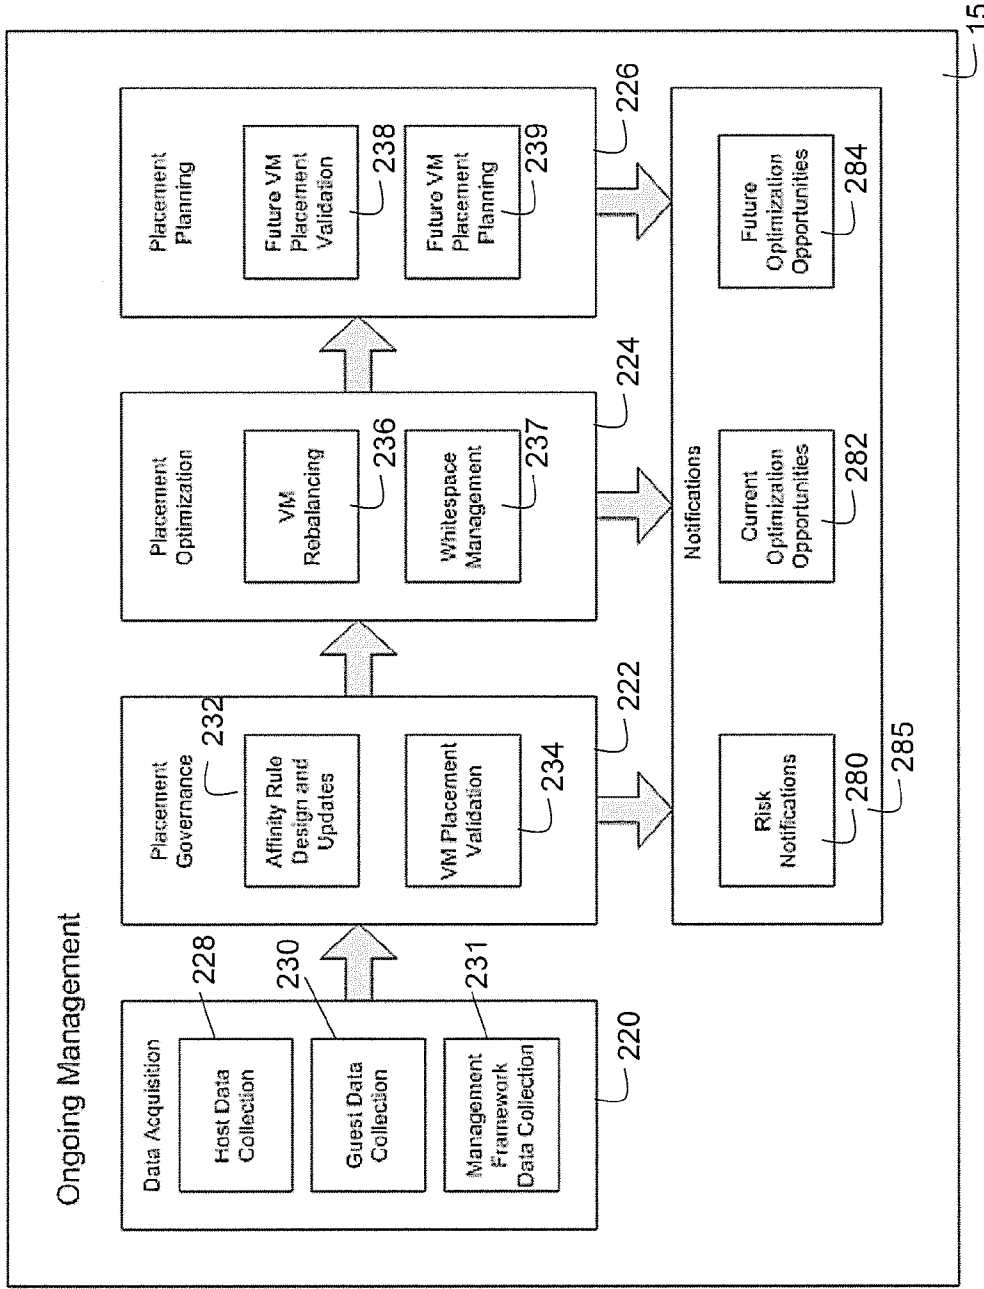


Figure 42

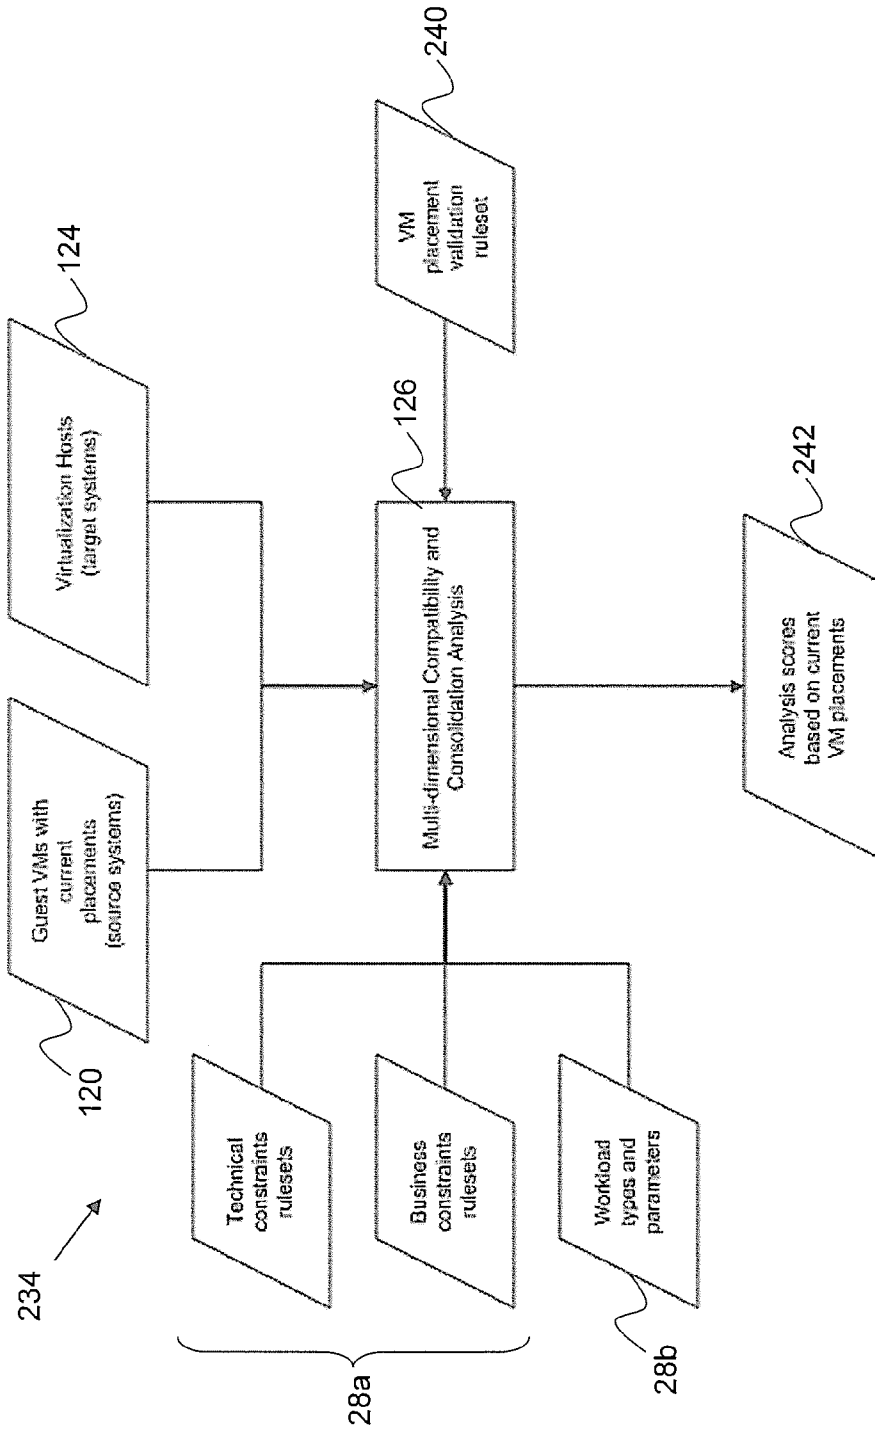


Figure 43

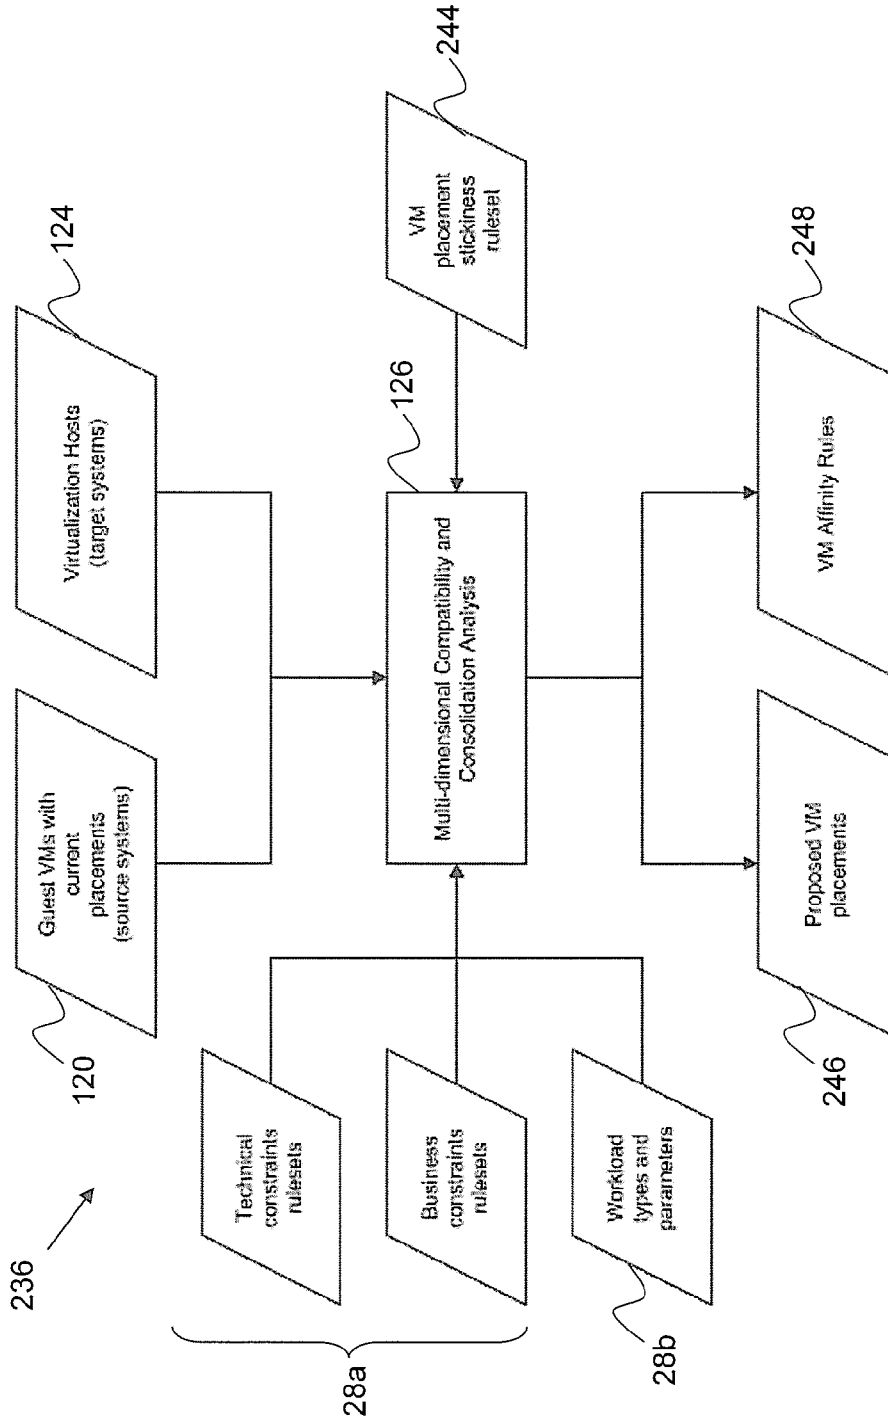


Figure 44

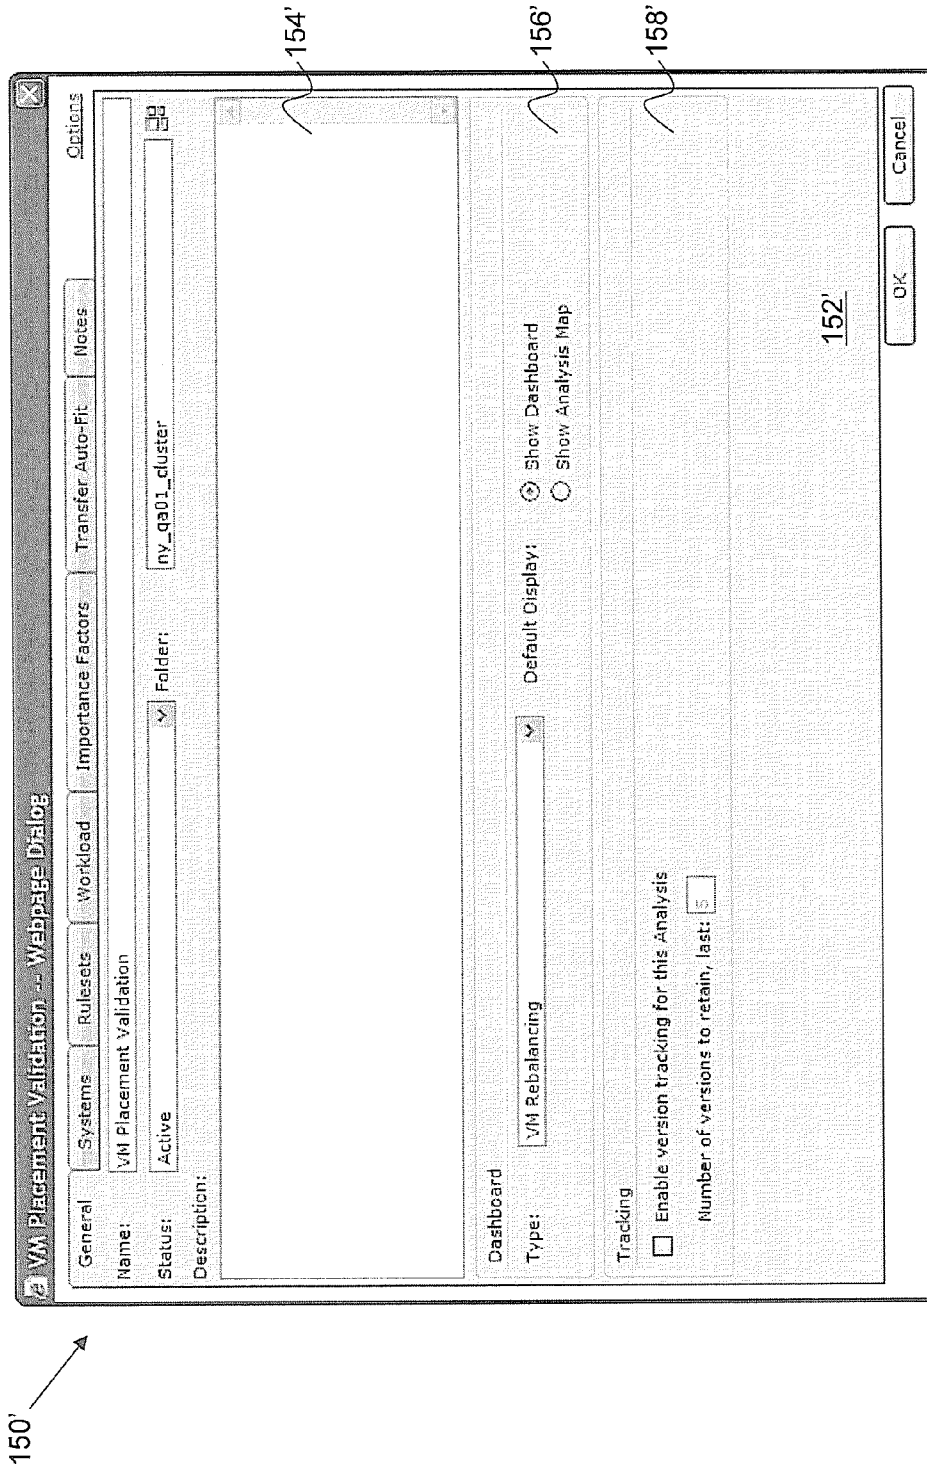


Figure 45

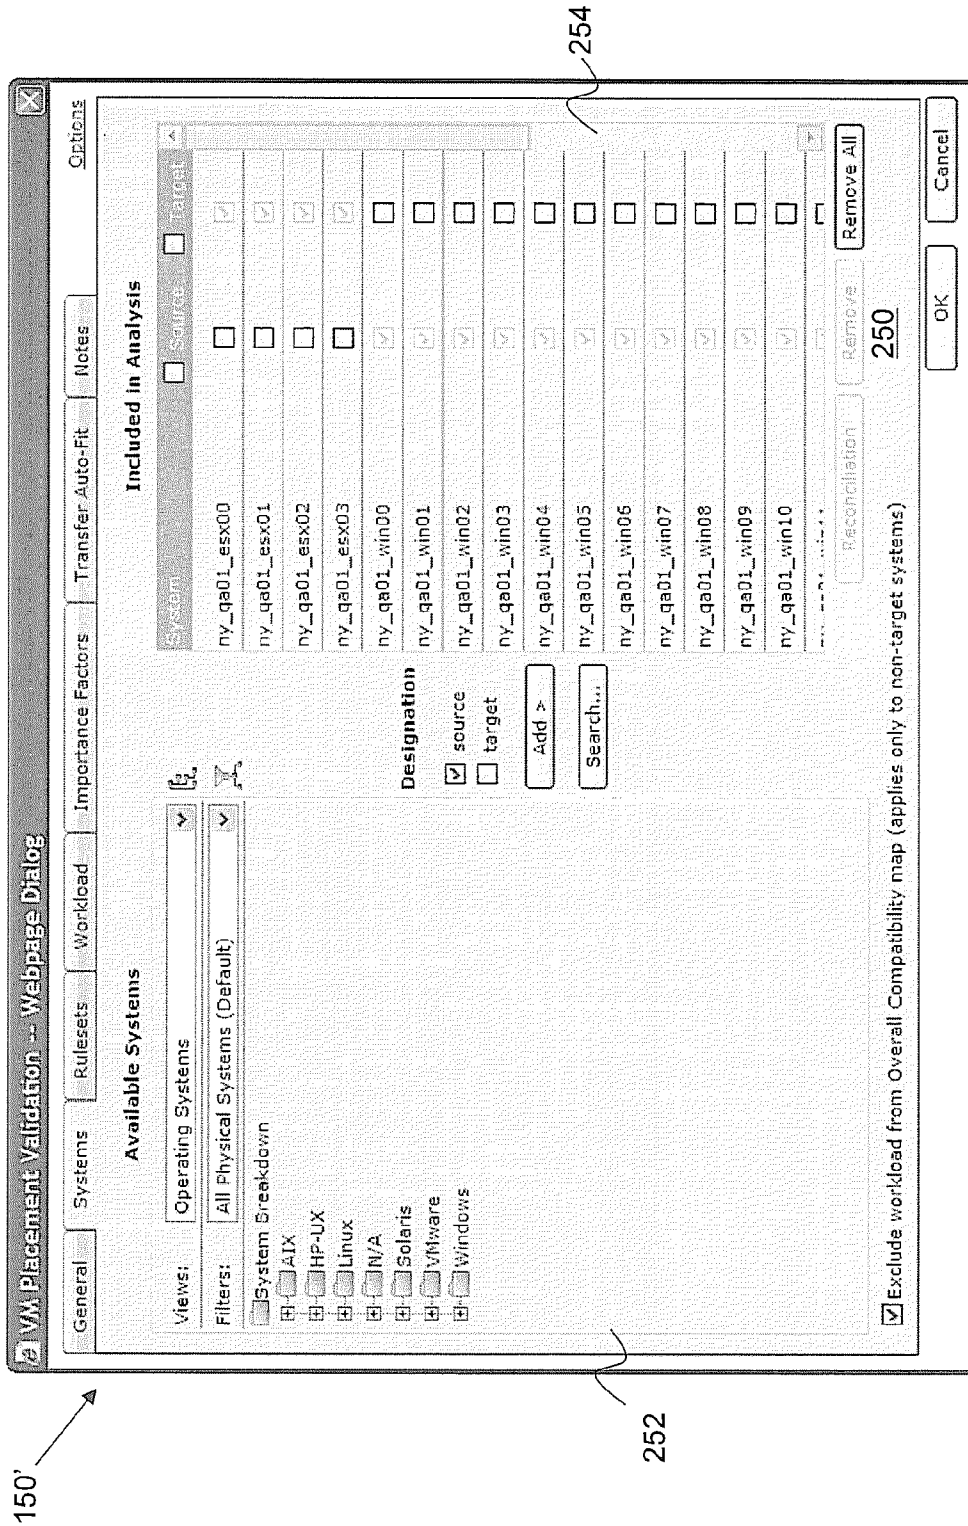


Figure 46




Figure 47




Figure 48

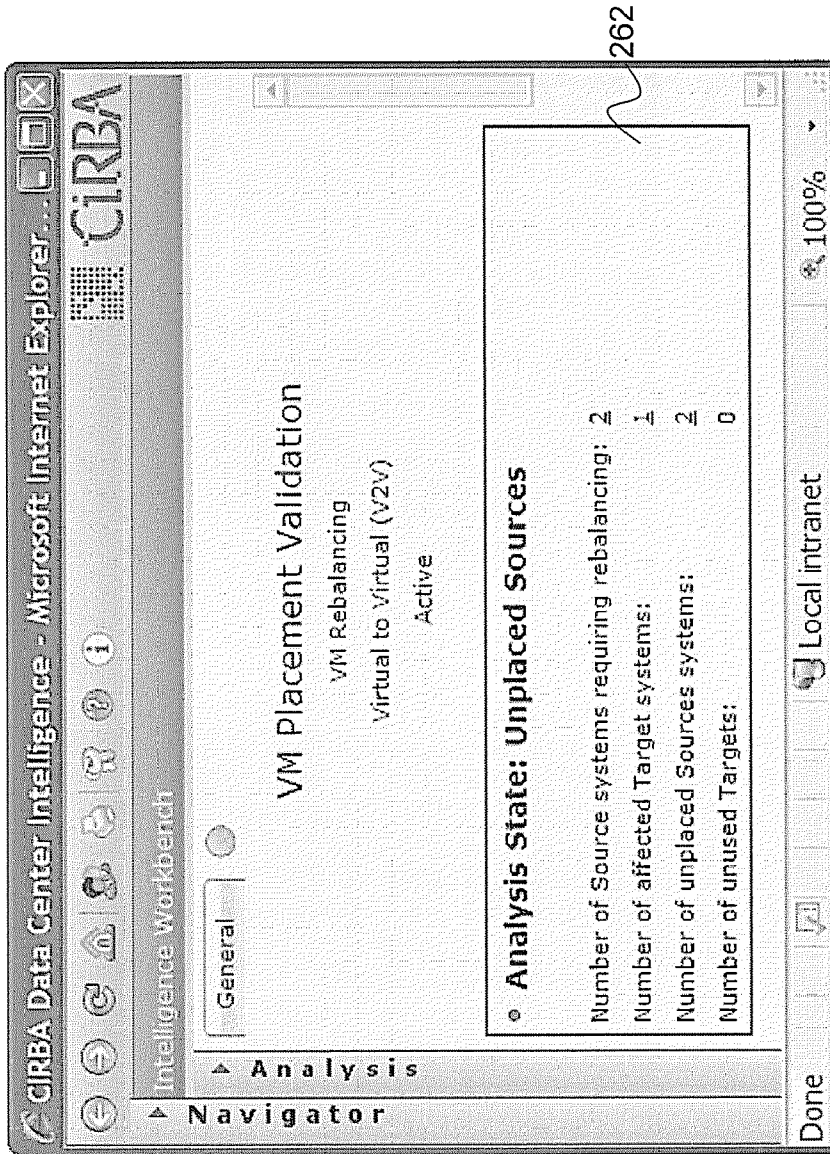
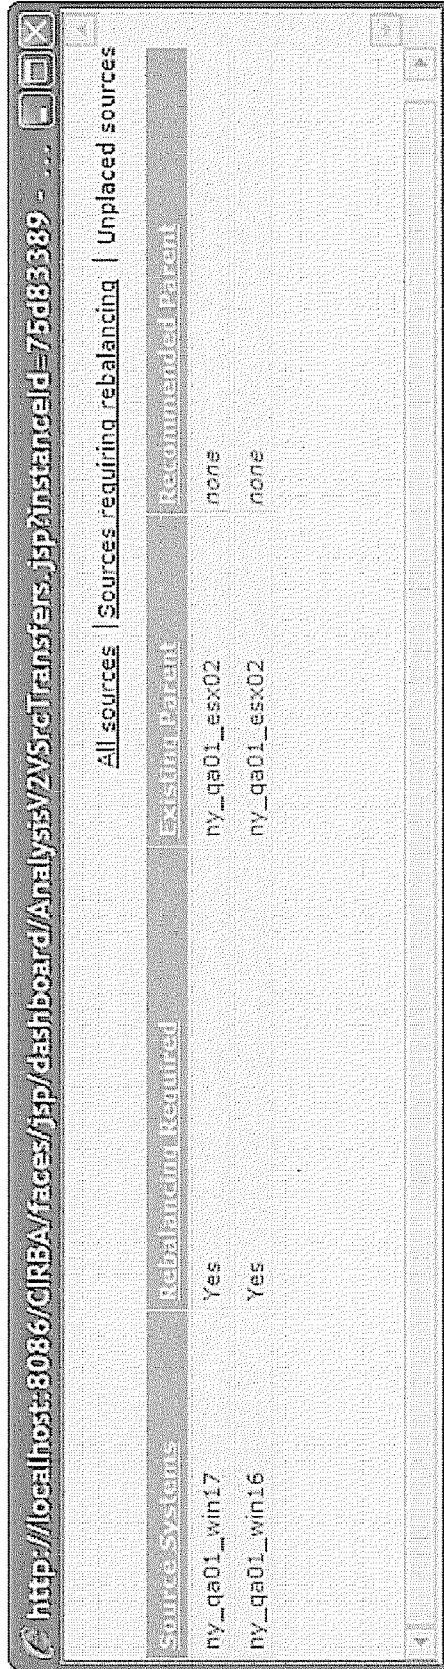


Figure 49

264 →



The screenshot shows a web browser window with the address bar containing the URL: <http://localhost:8086/CIRBA/faces/jsp/dashboard/Analysis/2V5rcTransfers.jsp?instanceId=75d83389>. The browser tabs include "All sources", "Sources requiring rebalancing", and "Unplaced sources". The main content area displays a table with the following data:

Source Systems	Rebalancing Required	Existing Parent	Recommended Parent
ny_qa01_win17	Yes	ny_qa01_esx02	none
ny_qa01_win16	Yes	ny_qa01_esx02	none

Figure 50

2618

Display	Analysis	Transitions	Reports	Overall: Compatibility Map
py_0001_wv001				
py_0001_wv002				
py_0001_wv003				
py_0001_wv004				
py_0001_wv005				
py_0001_wv006				
py_0001_wv007				
py_0001_wv008				
py_0001_wv009				
py_0001_wv010				
py_0001_wv011				
py_0001_wv012				
py_0001_wv013				
py_0001_wv014				
py_0001_wv015				
py_0001_wv016				
py_0001_wv017				
py_0001_wv018				
py_0001_wv019				
py_0001_wv020				
py_0001_wv021				
py_0001_wv022				
py_0001_wv023				
py_0001_wv024				
py_0001_wv025				
py_0001_wv026				
py_0001_wv027				
py_0001_wv028				
py_0001_wv029				
py_0001_wv030				
py_0001_wv031				
py_0001_wv032				
py_0001_wv033				
py_0001_wv034				
py_0001_wv035				
py_0001_wv036				
py_0001_wv037				
py_0001_wv038				
py_0001_wv039				
py_0001_wv040				
py_0001_wv041				
py_0001_wv042				
py_0001_wv043				
py_0001_wv044				
py_0001_wv045				
py_0001_wv046				
py_0001_wv047				
py_0001_wv048				
py_0001_wv049				
py_0001_wv050				
py_0001_wv051				
py_0001_wv052				
py_0001_wv053				
py_0001_wv054				
py_0001_wv055				
py_0001_wv056				
py_0001_wv057				
py_0001_wv058				
py_0001_wv059				
py_0001_wv060				
py_0001_wv061				
py_0001_wv062				
py_0001_wv063				
py_0001_wv064				
py_0001_wv065				
py_0001_wv066				
py_0001_wv067				
py_0001_wv068				
py_0001_wv069				
py_0001_wv070				
py_0001_wv071				
py_0001_wv072				
py_0001_wv073				
py_0001_wv074				
py_0001_wv075				
py_0001_wv076				
py_0001_wv077				
py_0001_wv078				
py_0001_wv079				
py_0001_wv080				
py_0001_wv081				
py_0001_wv082				
py_0001_wv083				
py_0001_wv084				
py_0001_wv085				
py_0001_wv086				
py_0001_wv087				
py_0001_wv088				
py_0001_wv089				
py_0001_wv090				
py_0001_wv091				
py_0001_wv092				
py_0001_wv093				
py_0001_wv094				
py_0001_wv095				
py_0001_wv096				
py_0001_wv097				
py_0001_wv098				
py_0001_wv099				
py_0001_wv100				

Figure 51

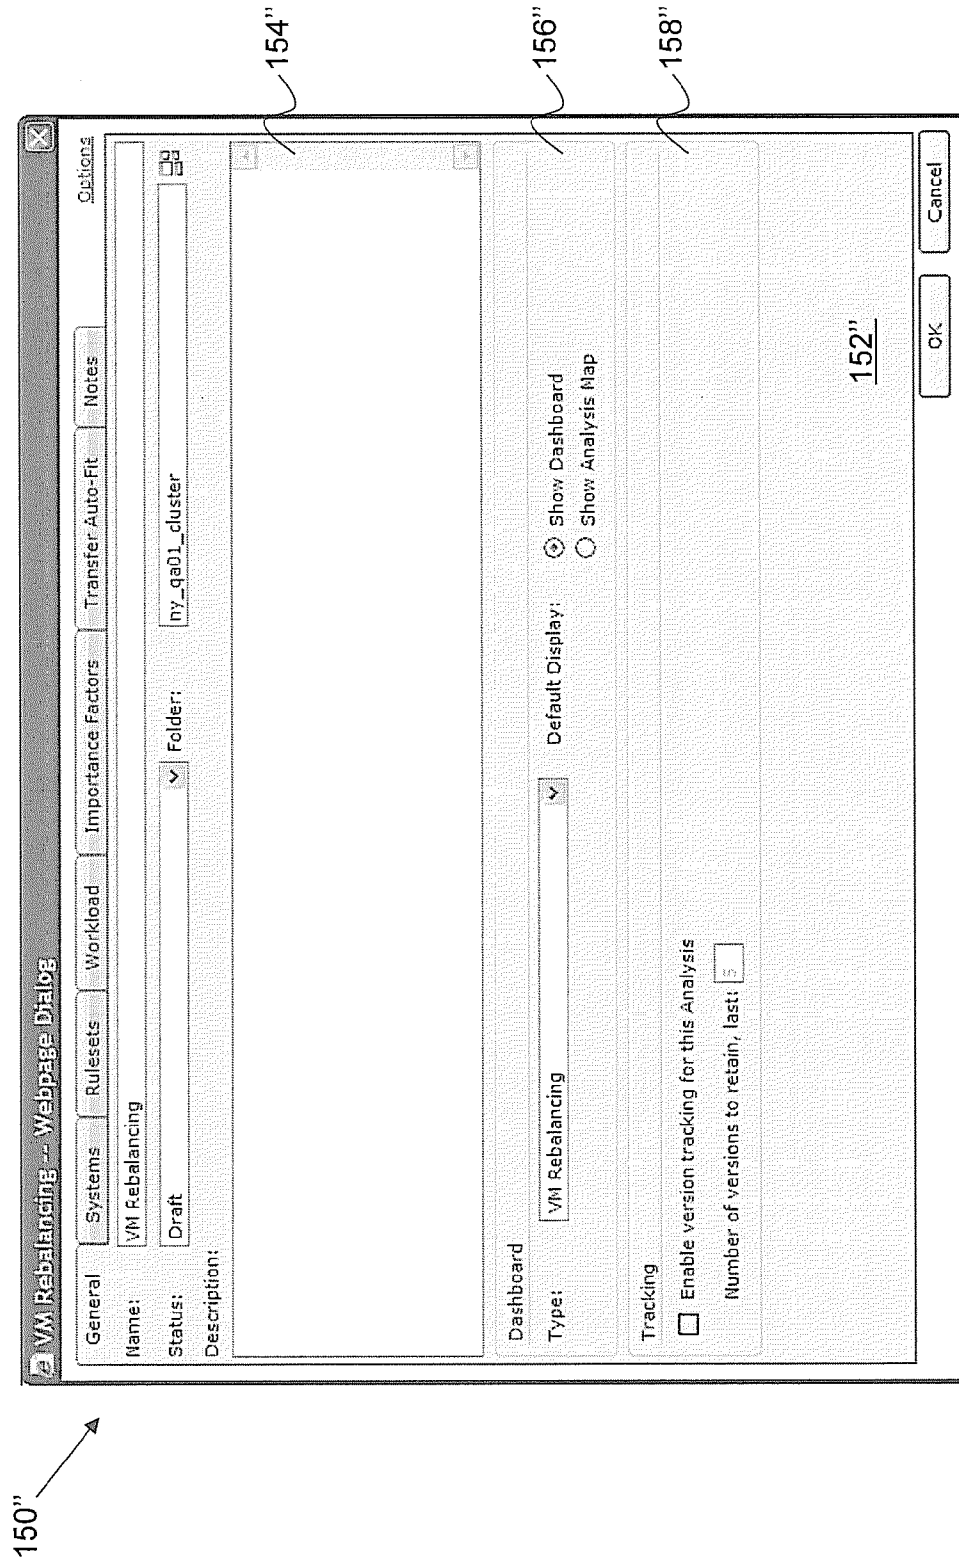


Figure 52




Figure 53

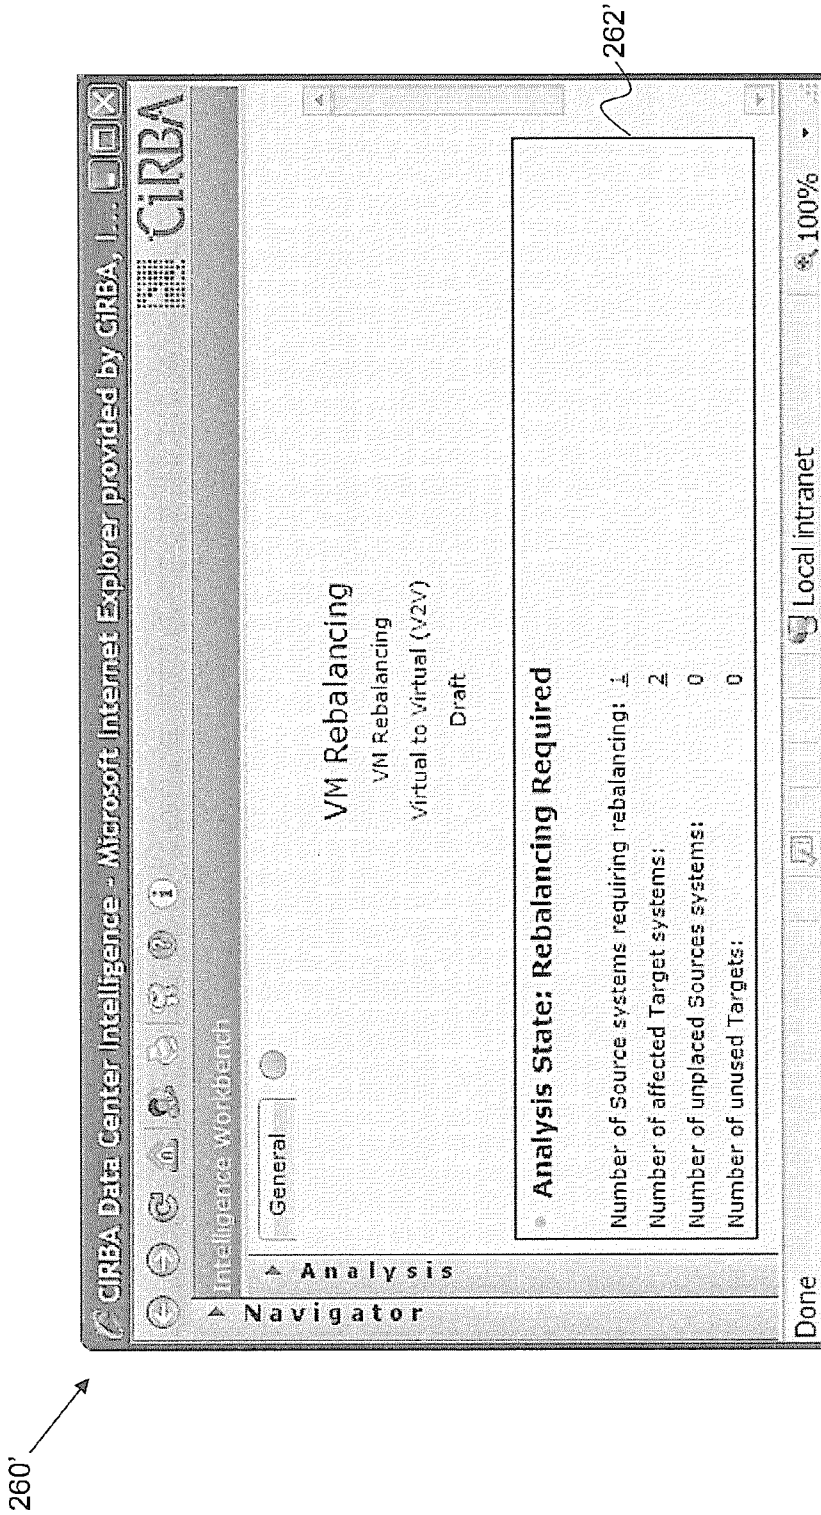


Figure 54

264' →




Figure 56

US 8,209,687 B2

1

METHOD AND SYSTEM FOR EVALUATING VIRTUALIZED ENVIRONMENTS

This application claims priority from U.S. Application No. 60/969,344 filed on Aug. 31, 2007, the contents of which are incorporated herein by reference. 5

TECHNICAL FIELD

The present invention relates generally to information technology infrastructures and has particular utility in designing and evaluating visualized environments. 10

BACKGROUND

As organizations have become more reliant on computers for performing day to day activities, so to has the reliance on networks and information technology (IT) infrastructures increased. It is well known that large organizations having offices and other facilities in different geographical locations utilize centralized computing systems connected locally over local area networks (LAN) and across the geographical areas through wide-area networks (WAN). 20

As these organizations grow, the amount of data to be processed and handled by the centralized computing centers also grows. As a result, the IT infrastructures used by many organizations have moved away from reliance on centralized computing power and towards more robust and efficient distributed systems. 25

While the benefits of a distributed approach are numerous and well understood, there has arisen significant practical challenges in managing such systems for optimizing efficiency and to avoid redundancies and/or under-utilized hardware. In particular, one challenge occurs due to the sprawl that can occur over time as applications and servers proliferate. Decentralized control and decision making around capacity, the provisioning of new applications and hardware, and the perception that the cost of adding server hardware is generally inexpensive, have created environments with far more processing capacity than is required by the organization. 30 35 40

When cost is considered on a server-by-server basis, the additional cost of having underutilized servers is often not deemed to be troubling. However, when multiple servers in a large computing environment are underutilized, having too many servers can become a burden. Moreover, the additional hardware requires separate maintenance considerations; separate upgrades and requires the incidental attention that should instead be optimized to be more cost effective for the organization. Heat production and power consumption can also be a concern. Even considering only the cost of having redundant licenses, removing even a modest number of servers from a large computing environment can save a significant amount of cost on a yearly basis. 45 50

As a result, organizations have become increasingly concerned with such redundancies and how they can best achieve consolidation of capacity to reduce operating costs. The cost-savings objective can be evaluated on the basis of consolidation strategies such as, but not limited to: virtualization strategies, operating system (OS) level stacking strategies, database consolidation strategies, application stacking strategies, physical consolidation strategies, and storage consolidation strategies. 55 60

Virtualization involves virtualizing a physical system as a separate guest OS instance on a host machine. This enables multiple virtualized systems to run on a single physical machine, e.g. a server. Examples of virtualization technolo-

2

gies include VMware®, Microsoft Virtual Server®, IBM LPAR®, Solaris Containers®, Zones®, etc.

The consolidation strategies to be employed, for virtualization or otherwise, and the systems and applications to be consolidated, are to be considered taking into account the specific environment. Consolidation strategies should be chosen carefully to achieve the desired cost savings while maintaining or enhancing the functionality and reliability of the consolidated systems. Moreover, multiple strategies may often be required to achieve the full benefits of a consolidation initiative.

Complex systems configurations, diverse business requirements, dynamic workloads and the heterogeneous nature of distributed systems can cause incompatibilities between systems. These incompatibilities limit the combinations of systems that can be consolidated successfully. In enterprise computing environments, the virtually infinite number of possible consolidation permutations which include suboptimal and incompatibility system combinations make choosing appropriate consolidation solutions difficult, error-prone and time consuming. 15 20

It is therefore an object of the following to address the above concerns.

SUMMARY

In one aspect, there is provided a method for designing a virtualized environment based on an existing physical environment comprising a plurality of systems, the method comprising: obtaining a data set for each of the plurality of systems, each data set comprising information pertaining to parameters associated with a corresponding system; performing a first compatibility analysis on the systems to determine candidate virtual guests; performing a second compatibility analysis on the systems to determine candidate virtual hosts; and performing a third compatibility analysis using the candidate virtual hosts, the candidate virtual guests and one or more rule sets pertaining to technical, business and workload constraints to generate a virtual environment design for virtualizing the plurality of systems. 25 30 35 40

In another aspect, there is provided a method for managing a virtualized environment, the method comprising: generating a virtual environment design for a plurality of existing physical systems using technical, business and workload constraints; facilitating the deployment of the virtualized environment according to the design; and on an ongoing basis: obtaining data pertaining to systems being used in the virtualized environment, validating placement of the systems in the virtualized environment, if necessary rebalancing the systems, and refining the virtualized environment. 45 50

In yet another aspect, there is provided a method for performing a virtual to virtual (V2V) transformation for a plurality of existing virtual servers, the method comprising: analyzing the existing virtual servers based on technical, business and workload constraints; based on the analyzing, determining which of the existing virtual servers are most suitable for conversion from one virtualized platform to another virtualized platform; and providing a mapping from the one platform to the another platform to facilitate the transformation. 55 60

In yet another embodiment, there is provided a method for determining a set of virtualization hosts for a virtualized environment based on an existing physical environment comprising a plurality of systems, the method comprising: obtaining a data set for each of the plurality of systems, each data set comprising information pertaining to parameters associated with a corresponding system; performing a first compatibility

US 8,209,687 B2

3

analysis of the plurality of systems using the data sets and a first rule set pertaining to virtualization specific constraints to determine an intermediate set of virtualization host candidates; and performing a second compatibility analysis of the intermediate set of candidates using a second rule set pertaining to migration specific constraints to determine the set of virtualization hosts.

In some embodiments, the method for determining the set of virtualized hosts comprises incorporating one or more hypothetical hosts into the set of virtualization hosts based on workload requirements for the virtualized environment.

In yet another aspect, there is provided a method for evaluating virtualization candidates to determine if additional systems are required to implement a desired virtualized environment, the method comprising: obtaining a set of virtualization guest candidates and determining aggregate workload requirements based on workload data pertaining to the guest candidates; obtaining a set of virtualization host candidates and determining aggregate workload capacity based on configuration data pertaining to the host candidates; comparing the workload requirements against the workload capacity to determine if sufficient capacity exists to satisfy the workload requirements; and if there is insufficient capacity, adding hypothetical server models to the host candidates to meet the workload requirements.

In yet another aspect, there is provided a method for validating an existing virtualized environment comprising a plurality of virtual machines placed on one or more virtual hosts, the method comprising: obtaining a data set for each of the plurality of virtual machines, each data set comprising information pertaining to technical, business and workload constraints associated with a corresponding virtual machine; evaluating the placement of the virtual machines in the virtualized environment using the data sets; and identifying the existence of virtual machines with suboptimal placements to enable replacement of the virtual machines.

In yet another aspect, there is provided a method for performing a power utilization analysis for a server, the method comprising: determining server load; determining power consumption for the server at idle and maximum loads; and estimating power utilization by combining the idle power consumption with a measurement based on a relationship between the maximum and idle power consumption.

In some embodiments, the method for performing a power utilization analysis comprises estimating the power utilization according to the following relationship: Estimated Power=Idle Power+Server Load*(Maximum Power-Idle Power).

BRIEF DESCRIPTION OF THE DRAWINGS

An embodiment of the invention will now be described by way of example only with reference to the appended drawings wherein:

FIG. 1 is a block diagram of a transformational physical to virtual (P2V) analytics system.

FIG. 2 is a flow diagram of a transformational P2V analysis process using the system shown in FIG. 1.

FIG. 3(a) is a block diagram of the analysis program depicted in FIG. 1.

FIG. 3(b) is a block diagram illustrating a sample consolidation solution comprised of multiple transfers.

FIG. 4 is an example of a compatibility analysis map.

FIG. 5 is a process flow diagram of the compatibility and consolidation analyses.

FIG. 6 is a process flow diagram illustrating the loading of system data for analysis.

4

FIG. 7 is a high level process flow diagram for a 1-to-1 compatibility analysis.

FIG. 8 is a process flow diagram for the 1-to-1 compatibility analysis.

FIG. 9 is a flow diagram illustrating operation of the rule engine analysis.

FIG. 10 is a flow diagram of the 1-to-1 rule-based compatibility analysis.

FIG. 11 is a flow diagram illustrating the evaluation of a rule set.

FIG. 12 is a flow diagram of workload data extraction process.

FIG. 13 is a flow diagram of the 1-to-1 workload compatibility analysis.

FIG. 14 is a screen shot of a date settings tab accessed through a workload settings page.

FIG. 15 is screen shot of an advanced workload settings page accessed through the 7 workload settings page shown in FIG. 14.

FIG. 16 is a screen shot of a limits tab accessed through the workload settings page.

FIG. 17 is a screen shot of a parameters tab accessed through the workload settings page.

FIG. 18 is a high level process flow diagram of the multi-dimensional compatibility 12 analysis.

FIG. 19 is a flow diagram showing the multi-dimensional analysis.

FIG. 20 is a flow diagram showing use of a rule set in an N-to-1 compatibility analysis.

FIG. 21 is a flow diagram showing use of a rule set in an N-by-N compatibility analysis.

FIG. 22 is a process flow diagram of the multi-dimensional workload compatibility analysis.

FIG. 23 is a process flow diagram of the consolidation analysis.

FIG. 24 is a process flow diagram of an auto fit algorithm used by the consolidation analysis.

FIG. 25 is process flow diagram showing further detail of the transformational P2V analysis process shown in FIG. 2.

FIG. 26 is a process flow diagram of an example implementation of the diagram shown in FIG. 25 using the analysis program illustrated in FIGS. 3 to 24.

FIG. 27 is a process flow diagram of an example aggregate workload sizing estimate process for evaluating resource capacity requirements.

FIG. 28 is a screen shot showing the main tab of an analysis editor program.

FIG. 29 is a screen shot showing the workload tab of the analysis editor program.

FIG. 30 is a compatibility map showing the result of a virtualization rule set applied against a set of physical systems.

FIG. 31 shows a net effect cube illustrating an N×N×M map for affinity and optimization analysis.

FIG. 32 is a target system compatibility map showing which systems in a current physical environment are candidates for being a virtualization host.

FIG. 33 is a screen shot showing an aggregate utilization report showing normalized utilization of an environment.

FIG. 34 is a live migration compatibility map showing the sets of systems that are compatible from a live migration perspective.

FIG. 35 is a screen shot showing a transfer auto-fit tab of the analysis editor of the analysis program.

FIG. 36 is a screen shot showing a dashboard summarizing the analysis results viewed through the analysis program.

US 8,209,687 B2

5

FIG. 37 is a screen shot of the proposed transfers from the analysis results viewed through the analysis-program.

FIG. 38 is a screen shot of a transfer map from the analysis results viewed through the analysis program.

FIG. 39 is a map showing a cluster-based view of virtual machines in a virtualized environment.

FIG. 40 is a screen shot of an affinity rule programming interface showing anti-affinity rules derived from the analysis results.

FIG. 41 is another screen shot of the affinity rule programming interface that supports the synchronization of affinity and anti-affinity rules with a third-party virtualization management framework.

FIG. 42 is a process flow diagram showing further detail of the ongoing management stage shown in FIG. 1.

FIG. 43 is a process flow diagram showing further detail of the VM placement validation step shown in FIG. 42.

FIG. 44 is a process flow diagram showing further detail of the VM rebalancing step shown in FIG. 42.

FIG. 45 is a screen shot of a main tab as viewed in the analysis editor program when used for a placement validation process.

FIG. 46 is a screen shot of a systems tab as viewed in the analysis editor program when used for a placement validation process.

FIG. 47 is a screen shot of a rule sets tab as viewed in the analysis editor program when used for a placement validation process.

FIG. 48 is a screen shot of the workload tab as viewed in the analysis editor program when used for a placement validation process.

FIG. 49 is a screen shot of a placement validation summary screen.

FIG. 50 is screen shot of a transfer summary produced during a placement validation process.

FIG. 51 is a screen shot of a placement validation compatibility map.

FIG. 52 is a screen shot of the main tab as viewed in the analysis editor program when used for a rebalancing process.

FIG. 53 is a screen shot of the rule sets tab as viewed in the analysis editor program when used for the rebalancing process.

FIG. 54 is a screen shot of the placement validation summary screen for the rebalancing process.

FIG. 55 is a screen shot of the placement validation compatibility map for the rebalancing process.

FIG. 56 is screen shot of the transfer summary produced during the rebalancing process.

DETAILED DESCRIPTION OF THE DRAWINGS

It has been recognized that virtualization often involves more than considering sizing, for example, it is beneficial to understand all the constraints that govern and impact a target environment and ensure that these constraints are taken into account when planning and managing a virtual environment. This has been found to be particularly true of virtualization infrastructures such as VMware Infrastructure®, where sophisticated features such as VMotion, distributed resource scheduling (DRS) and HA require careful planning and diligent administration of virtual environments. It has been found that to fully realize the capabilities of the virtualization infrastructure, the virtualization scheme being used should be combined with accurate intelligence and focused analytics in order to safely and effectively transform existing systems into a new virtual paradigm. In order to provide such intelligence and focused analytics, an analysis program for determining

6

compatibilities in a computing environment 12 can be utilized along with specific virtualization rule sets and user interfaces (UIs) to address the considerations of a virtualization infrastructure.

Virtualization Analysis and Optimization Overview

Turning now to FIG. 1, transformational physical-to-virtual (P2V) analytics system 9 can be implemented as noted above, by utilizing the principles and features provided by an analysis program 10 and incorporating virtualization rule sets 11 and a virtualization user interface 13, to transform an existing physical environment 12 comprising one or more systems 16 into a virtualized environment 21. As can also be seen in FIG. 1, the system 9 can be used on an ongoing basis once the virtualized environment 21 has been deployed to track the environment 21 as well as enable further analysis and optimization as servers and constraints change over time. It will be appreciated that although the examples provided herein are directed to P2V analyses, the principles and processes are equally applicable to transformational virtual-to-virtual (V2V) analysis, e.g. VMware® to Hyper-V® and vice versa.

Transformational P2V Analysis and Ongoing Management Overview

FIG. 2 provides a high-level process flow diagram 99 illustrating various stages performed by the transformational P2V analysis system 9. As will be explained in greater detail below, in order to intelligently analyze the physical environment 12 for virtualization, one or more data sets 18 are obtained, which pertain to information associated with parameters of the physical systems 16. These data sets 18 are used to perform a physical environment analysis 100 and a current asset assessment 102. The physical environment analysis 100 analyzes existing physical systems 16 in the current physical environment 12 to be virtualized to evaluate various technical, business and workload constraints and affinity considerations of the virtualization candidates. In this way, the suitability of each system 16 to be virtualized can be determined to identify suitable source systems for virtualization to facilitate the design of the virtual environment 21. The current asset assessment 102 evaluates the viability of repurposing existing physical systems 16 as virtualization hosts. A virtualization host refers to a target system that runs hypervisor software and is capable of hosting virtual machines. This allows for an assessment of the equipment that is currently available to minimize the amount of new equipment required for virtualization.

The outcome of the current asset assessment 102 can be used to perform a virtualization host system definition 104, which can incorporate an analysis of hypothetical systems used to model target systems that do not currently exist in the physical environment 12. This allows users to evaluate a wide range of scenarios. The virtualization host system definition 104 can also incorporate live migration compatibilities amongst a target solution (set of target systems based on current asset assessment 102 and hypothetical systems). In this way, a target solution can be defined to facilitate the design of the virtual environment 21, i.e. in conjunction with the outcome of the physical environment analysis 100.

The set of source systems and the set of target systems are then used to perform a virtual environment optimization 106, which determines the optimal layout of the source systems onto the target systems based on technical, business and workload constraints according to a multi-dimensional compatibility and consolidation analysis using the analysis program 10. It can be seen in FIG. 2 that virtualization rule sets 11 are used during stages 100-106 in various ways as will be explained below. The virtualization UI 13 can also be used

during these stages to permit a user to interact with the analysis program **10** and ultimately generate a virtual environment design **110**. It will be appreciated that the process flow shown in FIG. **2** is for illustrative purposes only and may proceed differently in different scenarios. For example, based on outcomes of the physical environment analysis **100** and virtualization host system definition **104**, various analyses may be conducted iteratively to narrow in on progressively more optimal solutions to accommodate both existing constraints and changing environments (both physical and virtual). The virtual environment design **110** can then be used to create a virtualization solution **112** that, when implemented, can be tracked, analyzed and refined over time by conducting ongoing management **15**.

As discussed above, the transformational P2V analysis **9** advantageously utilizes the components and principles of the analysis program **10**. As such, to assist in understanding the transformational P2V analytics **9**, an overview of an example of the analysis program **10** will be provided. It may be noted that additional detail pertaining to the analysis program is described in U.S. patent application Ser. No. 11/738,936 filed on Apr. 23, 2007 and published under U.S. 2007/0250829, the contents of which are incorporated herein by reference.

Analysis Program Overview

A block diagram of an analysis program **10** for determining compatibilities in computing environment **12** is provided in FIG. **3(a)**. The analysis program **10**, accessed through a computer station **14**, gathers data **18** pertaining to a collection of systems to be consolidated **16**. The analysis program **10** uses the gathered data **18** to evaluate the compatibility of the computer systems **28** and provide a roadmap **20** specifying how the original set of systems can be consolidated to a smaller number of systems **22**.

A distinct data set is obtained for each system **16** to contribute to the combined system data **18** shown in FIG. **3(a)**. Each data set comprises one or more parameters that relate preferably to technical **24**, business **26** and workload **28** characteristics or features of the respective system **16**. The parameters can be evaluated by scrutinizing program definitions, properties, objects, instances and any other representation or manifestation of a component, feature or characteristic of the system **16**. In general, a parameter is anything related to the system **16** that can be evaluated, quantified, measured, compared etc. Examples of technical parameters relevant of the consolidation analysis include the operating system, OS version, patches, application settings, hardware devices, etc. Examples of business parameters of systems relevant to the consolidation analysis include the physical location, organization department, data segregation requirements, owner, service level agreements, maintenance windows, hardware lease agreements, software licensing agreements, etc. Examples of workload parameters relevant to consolidation analysis include various resource utilization and capacity metrics related to the system processor, memory, disk storage, disk I/O throughput and network bandwidth utilization.

The system data parameters associated with a system **16** comprise the system model used in the analyses. In the following examples, a source system refers to a system from which applications and/or data are to be moved, and a target server or system is a system to which such applications and/or data are to be moved. For example, an underutilized environment having two systems **16** can be consolidated to a target system (one of the systems) by moving applications and/or data from the source system (the other of the systems) to the target system.

The computer systems **16** may be physical systems, virtual systems or hypothetical models. In contrast to actual physical

systems, hypothetical systems do not currently exist in the computing environment **12**. Hypothetical systems can be defined and included in the analysis to evaluate various types of “what if” consolidation scenarios. Hypothetical targets can be used to simulate a case where the proposed consolidation target systems do not exist in the environment **12**, e.g. for adding a system **16**. Similarly, 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**. Hypothetical systems can be created through data imports, cloning from actual systems models, and manual specification by users, etc. The system model can be minimal (sparse) or include as much data as an actual system model. These system models may also be further modified to address the analysis requirements.

The compatibility analysis can also be generalized to evaluate entities beyond physical, virtual or hypothetical systems. For example, entities can be components that comprise systems such as applications and database instances. By analysing the compatibility of database instances and database servers with database stacking rule sets, database consolidation can also be assessed. Similarly, application consolidation can be evaluated by analyzing application servers and instances with application stacking rules. The entity could also be a logical application system and technical data can pertain to functional aspects and specifications of the entity. It will therefore be appreciated that a “system” or “computer system” hereinafter referred, can encompass any entity which is capable of being analysed for any type of compatibility and should not be considered limited to existing or hypothetical physical or virtual systems etc.

Consolidation as described above can be considered to include one or more “transfers”. The actual transfer describes the movement of a single source entity onto a target, wherein the specification identifies the source, target and transfer type. The transfer type (or consolidation strategy) describes how a source entity is transferred onto a target, e.g. virtualization, OS stacking etc. A transfer set **23** (see FIG. **3(b)**) can be considered one or more transfers that involve a common target, wherein the set specifies one or more source entities, the target and a transfer type. A consolidation solution (or roadmap) is one or more transfer sets **23** based on a common pool of source and target entities. As can be seen in FIG. **3(a)**, the consolidation roadmap can be included in the analysis results **20**. Each source or target entity is referenced at most one time by the transfer sets that comprise the solution. FIG. **3(b)** shows how an example pool **24** of 5 systems (S1, S2, S3, S4 and S5) can be consolidated through 2 transfer sets **23**: stack S1 and S2 onto S3, and stack S4 onto S5. The transfer sets **23** include 3 transfers, and each system **16** is referenced by the transfer sets **23** only once. In the result, a consolidated pool **26** of 2 systems is achieved. It will be appreciated that the principles described herein support many transformation strategies and consolidation is only one example.

The following discusses compatibilities between systems **16** based on the parameters to determine if efficiencies can be realized by consolidating either entire systems **16** or aspects or components thereof. The analyses employ differential rule sets **28** to evaluate and quantify the compatibility of systems **16** with respect to technical configuration and business related factors comprised in the gathered system data **18**. Similarly, workload compatibility of a set of systems **16** is assessed using workload stacking and scoring algorithms **30**. The results of configuration (technical), business and workload compatibility analyses are combined to produce an overall compatibility score for a set of systems **16**. In addition to compatibility scores, the analysis provides details that

account for the actual scores. The scores can be presented in color coded maps **32** that illustrate patterns of the compatibility amongst the analyzed systems as shown generally in FIG. **4**.

The compatibility analysis map **32** provides an organized graphical mapping of system compatibility for each source/target system pair on the basis of configuration data. The map **32** shown in FIG. **4** is structured having each system **16** in the environment **12** listed both down the leftmost column and along the uppermost row. Each row represents a consolidation source system, and each column represents the possible consolidation target. Each cell **92** contains the score **36** corresponding to the case where the row system is consolidated onto the column (target) system **16**.

The output shown in FIG. **4** arranges the systems **16** in the map **32** such that a 100% compatibility exists along the diagonal where each system **16** is naturally 100% compatible with itself. The map **32** is preferably displayed such that each cell **92** includes a numerical score **36** and a shade of a certain colour **34**. As noted above, the higher the score (from zero (0) to one hundred (100)), the higher the compatibility. The scores are pre-classified into predefined ranges that indicate the level of compatibility between two systems **16**. Each range maps to a corresponding colour or shade for display in the map **32**. For example, the following ranges and colour codes can be used: score=100, 100% compatible, dark green; score=75-99, highly compatible, green; score=50-74, somewhat compatible, yellow; score=25-49, low compatibility, orange; and score=0-24, incompatible, red.

The above ranges are only one example. Preferably, the ranges can be adjusted to reflect more conservative and less conservative views on the compatibility results. The ranges can be adjusted using a graphical tool similar to a contrast slider used in graphics programs. Adjustment of the slider would correspondingly adjust the ranges and in turn the colours. This allows the results to be tailored to a specific situation. It is therefore seen that the graphical output of the map **32** provides an intuitive mapping between the source/target pairs in the environment **12** to assist in visualizing where compatibilities exist and do not exist. Detailed differences and other information can be viewed by selecting a relevant cell **92**, which accesses information such as differences tables showing the important differences between the two systems, the rules and weights that were applied and may even provide a remediation cost.

A collection of systems **16** to be consolidated can be analyzed in one of three modes: 1-to-1 compatibility, multi-dimensional compatibility and consolidation analyses. These analyses share many common aspects but can be performed independently.

The 1-to-1 compatibility analysis evaluates the compatibility of every possible source-target pair combination in the collection of systems **16** on a 1-to-1 basis. This analysis is useful in assessing single transfer consolidation candidates. In practice, it may be prudent to consolidate systems **16** incrementally and assess the impact of each transfer before proceeding with additional transfers. The multi-dimensional compatibility analysis evaluates the compatibility of transfer sets that can involve multiple sources being transferred to a common target. The analysis produces a compatibility score for each specified transfer set **23** by evaluating the compatibility of the systems **16** that comprise the transfer set **23**. The consolidation analysis searches for a consolidation solution that minimizes the number of remaining source and target entities after the proposed transfers are applied, while meeting requisite compatibility constraints. This analysis employs

the multi-dimensional compatibility analysis described above to evaluate the compatibility of postulated transfer sets.

The analysis program **10** performs consolidation analyses for virtualization and stacking strategies as will be explained in greater detail below, however, it will be appreciated that other consolidation strategies may be performed according to similar principles.

Referring now to FIG. **5**, a process flow diagram illustrates the data flow for performing the compatibility and consolidation analyses discussed above. The flow diagram outlines four processes: a data load and extraction process (A), a 1-to-1 compatibility analysis process (B), a multi-dimensional compatibility analysis process (C), and a consolidation analysis process (D).

In process A, the system data **18** collected via audits or imports as discussed above is prepared for use by the analyses. The compatibility and consolidation analyses processes B, C and D can be performed independently. The analyses share a common analysis input specification and get system data **18** from the data repository **54** and caches **56** and **58**. The multi-dimensional compatibility and consolidation analyses take additional inputs in the form of a consolidation solution and auto fit input parameters **84** and **86** respectively.

The 1-to-1 compatibility analysis process B evaluates the compatibility of each system pair on a 1-to-1 basis. In contrast, the multi-dimensional analysis process C evaluates the compatibility of each transfer set **23** in the consolidation solution that was specified as part of its input.

The consolidation analysis process D searches for the best consolidation solution that fulfills the constraints defined by the auto fit input **86**. The consolidation analysis employs the multi-dimensional compatibility analysis C to assess potential transfer set candidates.

A process flow diagram for the data load and extraction process A is illustrated in FIG. **6**. System data including technical configuration, business related and workload collected through audits, data import and user input are prepared for use by the analyses processes B, C and D.

When system data **18** and attributes are loaded into the analysis program **10**, they are stored in the audit data repository **54** and system attribute table **55**, respectively. As well, system data **18** referenced by rule set items **28** (see FIG. **9**), workload types **30** and benchmarks are extracted and loaded into their respective caches **56**, **58**. Alias specifications **60** describe how data can be extracted and if necessary, normalized from a variety of data sources.

The data repository **54** and caches **56** and **58** thus store audited data **18**, system attributes, the latest rule set data, historical workload data and system workload benchmarks.

A high level flow diagram of the 1-to-1 compatibility analysis is shown in FIG. **7**. The 1-to-1 compatibility analysis can take into account analysis input, including input regarding the systems **16** to be analyzed, rule set related parameters, workload related parameters, workload benchmarks and importance factors **88** used to compute overall scores.

The compatibility analysis evaluates the compatibility of every specified system as source-target pairs on a 1-to-1 basis. This analysis produces a compatibility score for each system pair so that analyzing a collection often (10) systems **16** produces 10x10 scores. The compatibility analysis is based on the specified rule sets and workload types. An analysis may be based upon zero or more rule sets and zero or more workload types, such that at least one rule set or workload type is selected.

The selection of rule sets **28** and workload types **30** for an analysis depends on the systems **28** and the consolidation strategy to analyze. For example, to assess the consolidation

11

of a set of UNIX™ systems 16, an analysis may employ the UNIX™ application stacking, location, maintenance window and ownership rule sets 28, and CPU, memory, disk space, disk I/O and network I/O workload types 30.

A process flow diagram of the 1-to-1 compatibility analysis is shown in FIG. 8. The analysis generally comprises four stages. In the first stage, data referenced by the selected rule sets 28 and workload types 30 for the specified date range are retrieved from the data repository 54 and caches 56, 58 for each system 16 to be analyzed. This analysis data is saved as a snapshot and can be used for subsequent analyses. In the second stage, technical and business related compatibility may be analyzed using the specified rule sets 28 and weights. Next, workload compatibility is evaluated based on the specified workload types 30 and input parameters. Finally, the overall compatibility scores are computed for each pair of systems 16. Upon completion of the compatibility analysis, the results 20 are provided to the user. The results 20 include rule item and workload data snapshots, 1-to-1 compatibility score maps for each rule set 28 and workload type 30 as well as an overall score map. Analysis details for each map may also be provided.

As noted above, the differential rule sets 28 are used to evaluate the compatibility of systems as they relate to technical and business related constraints. The rule set 28 defines which settings are important for determining compatibility. The rule set 28 typically defines a set of rules which can be revised as necessary based on the specific environment 12. The rule set 28 is thus preferably compiled according to the systems 16 being analysed and prior knowledge of what makes a system 16 compatible with another system 16 for a particular purpose. As will be discussed below, the rule sets 28 are a form of metadata 62.

Further detail regarding the differential rules and differential rule sets 28 is now described making reference to FIG. 9. Additional detail regarding the differential rules and rule sets 28 is also described in co-pending U.S. patent application Ser. No. 11/535,308 filed on Sep. 26, 2006, and entitled "Method for Evaluating Computer Systems", the contents of which are incorporated herein by reference.

With respect to the following description of the rule sets 28 and the general application of the rule sets 28 for detecting system incompatibilities by evaluating differences between data parameters of systems 16, the following alternative nomenclature may be used. A target system refers to a system being evaluated, and a baseline system is a system to which the target system is being compared. The baseline and target systems may be the same system 16 at different instances in time (baseline=prior, target=now) or may be different systems 16 being compared to each other. As such, a single system 16 can be evaluated against itself to indicate changes with respect to a datum as well as how it compares to its peers. It will be appreciated that the terms "source system" and "baseline system" are herein generally synonymous, whereby a source system is a type of baseline system.

FIG. 3(a) illustrates the relationships between system data 18 and the analysis program 10. Data 18 is obtained from the source and target computer systems 16 and is used to analyze the compatibility between the systems 16. In this example, the parameters are evaluated to determine system compatibilities for a consolidation strategy. A distinct data set 18 is preferably obtained for each system 16 (or instance in time for the same system 16 as required). Rule sets 28 are computer readable and storable so that they may be accessed by the program 10 and modified if necessary, for use in evaluating the computer systems 16.

12

Rule sets 28 are groupings of rules that represent higher-level considerations such as business objectives or administrative concerns that are taken into account when reporting on or analysing the systems 16. In FIG. 9, six rules 43, A, B, C, D, E and F are grouped into three rule sets 28, Rule Set 1, 2 and 3. It will be appreciated that there may be any number of rules in any number of rule sets 28 and those shown in FIG. 9 are for illustrative purposes only.

Rules evaluate data parameters according to rule definitions to determine incompatibilities due to differences (or contentious similarities) between the baseline and target systems. The rule definitions include penalty weights that indicate the importance of the incompatibility as they relate to the operation of the systems 16. The penalty weights are applied during an evaluation if the incompatibility is detected. The evaluation may include the computation of a score or generation of other information indicative of nature of the incompatibilities between the baseline and target systems.

Rules comprised by a rule set 28 may reference common parameters but perform different tests to identify different forms of incompatibilities that may have different levels of importance. For example a version four operating system versus a version three operating system may be considered less costly to remedy and thus less detrimental than a version five operating system compared to a version one operating system. As can be seen, even though the operating systems are different in both cases, the nature of the difference can also be considered and different weights and/or remedies applied accordingly.

Rules can also test for similarities that indicate contentions which can result in incompatibilities between systems. For example, rules can check for name conflicts with respect to system names, database instance names, user names, etc.

The flow of data for applying exemplary rule sets 28 is shown in FIG. 9. In this example, the system data gathered from a pair of systems 16 are evaluated using three rule sets. A rule engine or similar device or program evaluates the data parameters of the systems 16 by applying rule sets 1, 2 and 3 which comprise of the exemplary rules A, B, C, D, E and F. The evaluation of the rules results in compatibility scores and zero or more matched rule items for each rule set 28. These results can be used for subsequent analyses, such as combining with workload compatibility results to obtain overall compatibility scores.

The system consolidation analysis computes the compatibility of a set of systems 16 based not only on technical and workload constraints as exemplified above, but also business constraints. The business constraints can be expressed in rule sets 28, similar to the technical constraints discussed above.

It may be appreciated that basic and advanced rule sets 28 can be created. Where basic and advanced rule sets 28 are available for the same analysis program 10, there are a number of options for providing compatibility. The rule set specification can be extended to include a property indicating the minimum required rule engine version that is compatible with the rule set. In addition, the basic rule sets can be automatically migrated to the advanced rule set format since the advanced specification provides a super set of functionality relative to the basic rule set specification. It will be appreciated that as new rules and rule formats are added, compatibility can be achieved in other ways so long as legacy issues are considered where older rule versions are important to the analysis.

An exemplary process flow for a rule-based compatibility analysis is shown in greater detail in FIGS. 10 and 11. When analyzing system compatibility, the list of target and source systems 16 are the same. The compatibility is evaluated in two

directions, e.g. for a Server A and a Server B, migrating A to B is considered as well as migrating B to A.

Turning first to FIG. 10, for each rule set R (R=1 to M where M is the number of rule sets) and for each target system T (T=1 to N where N is the number of systems), the rule engine 90 first looks at each source system S (S=1 to N). If the source=target then the configuration compatibility score for that source is set to 100, no further analysis is required and the next pair is analyzed. If the source and target are different, the rules are evaluated against the source/target pair to compute the compatibility score, remediation cost and to compile the associated rule details. Estimated remediation costs are optionally specified with each rule item. As part of the rule evaluation and subsequent compatibility score calculation, if a rule is true, the corresponding cost to address the deficiency is added to the remediation cost for the pair of systems 16 being analysed.

The evaluation of the rules is shown in FIG. 11. The evaluation of the rules considers the snapshot data 18 for the source system and the target system, as well as the differential rule set 28 that being applied. For each rule in the set 28, the data referenced by the rule is obtained for both the target and source. The rule is evaluated by having the rule engine 90 compare the data. If the rule is not true (i.e. if the systems 16 are the compatible according to the rule definition) then the data 18 is not considered in the compatibility score and the next rule is evaluated. If the rule is true, the rule details are added to an intermediate result. The intermediate result includes all true rules.

Preferably, a suppression tag is included with each rule. As discussed above, the suppression tag indicates other rules that are not relevant if that rule is true. The suppression flag allows the program 10 to avoid unnecessary computations. A mutex flag is also preferably used to avoid unfairly reducing the score for each true rule when the rules are closely affected by each other.

Once each rule has been evaluated, a list of matched rules is created by removing suppressed rule entries from the intermediate results based on rule dependencies, which are defined by rule matching and suppression settings (e.g. match flags and suppression tags). The compatibility score for that particular source/target pair is then computed based on the matched rules, weights and mutex settings. Remediation costs are also calculated based on the cost of updating/upgrading etc. and the mutex settings.

Turning back to FIG. 10, the current target is then evaluated against all remaining sources and then the next target is evaluated. As a result, an N×N map 32 can be created that shows a compatibility score for each system against each other system. The map 32 can be sorted by grouping the most compatible systems. The sorted map 32 is comprised of every source/target combination and thus provides an organized view of the compatibilities of the systems 16.

Preferably, configuration compatibility results are then generated for each rule set 28, comprising the map 32 (e.g. FIG. 4) and for each source-target pair details available pertaining to the configuration compatibility scoring weights, remediation costs and applicable rules. The details can preferably be pulled for each source/target pair by selecting the appropriate cell 92.

The workload compatibility analysis evaluates the compatibility of each source-target pair with respect to one or more workload data types 30. The analysis employs a workload stacking model to combine the source workloads onto the target system. The combined workloads are then evaluated using threshold and a scoring algorithm to calculate a compatibility score for each workload type.

System workload constraints must be assessed when considering consolidation to avoid performance bottlenecks. Workload types representing particularly important system resources include % CPU utilization, memory usage, disk space used, disk I/O throughput and network I/O throughput. The types of workload analyzed can be extended to support additional performance metrics. Workload values can be represented as percentages (e.g. % CPU used) or absolute values (e.g. disk space used in MB, disk I/O in MB/sec).

The term workload benchmark refers to a measure of a system's capability that may correspond to one or more workload types. Workload benchmarks can be based on industry benchmarks (e.g. CINT2000 for processing power) or the maximum value of a system resource (e.g. total disk space, physical memory, network I/O bandwidth, maximum disk I/O rate). Benchmarks can be used to normalize workload types that are expressed as a percentage (e.g. % CPU used) to allow direct comparison of workloads between different systems 16. Benchmarks can also be used to convert workload types 30 that are expressed as absolute values (e.g. disk space used in MB) to a percentage (e.g. % disk space used) for comparison against a threshold expressed as a percentage.

System benchmarks can normalize workloads as follows. For systems X and Y, with CPU benchmarks of 200 and 400 respectively (i.e. Y is 2× more powerful than X), if systems X and Y have average CPU utilizations of 10% and 15% respectively, the workloads can be normalized through the benchmarks as follows. To normalize X's workload to Y, multiply X's workload by the benchmark ratio X/Y, i.e. $10\% \times 200 / 400 = 5\%$.

Stacking X onto Y would then yield a total workload of $5\% + 15\% = 20\%$. Conversely, stacking Y onto X would yield the following total workload: $10\% + 15\% \times 400 / 200 = 40\%$.

As discussed above, workload data is collected for each system 16 through various mechanisms including agents, standard instrumentation (e.g. Windows Performance Monitor™, UNIX™ System Activity Reporter), custom scripts, third party performance monitoring tools, etc. Workload data is typically collected as discrete time series data. Higher sample frequencies provide better accuracy for the analysis (5 minute interval is typical). The workload data values should represent the average values over the sample period rather than instantaneous values.

Data from different sources may need to be normalized to common workload data types 30 to ensure consistency with respect to what and how the data is measured. For example, CPU usage may be reported as Total % CPU utilization, % CPU idle, % CPU system, % CPU user, % CPU I/O, etc. Disk utilization may be expressed in different units such as KB, MB, blocks, etc.

The time series workload data can be summarized into hourly quartiles. Specifically, the minimum, 1st quartile, median, 3rd quartile, maximum, and average values are computed for each hour. The compatibility analysis for workload uses the hourly quartiles. These statistics allow the analysis to emphasize the primary operating range (e.g. 3rd quartile) while reducing sensitivity to outlier values.

Workload data is typically collected and stored in the workload data cache 58 for each system 16 for multiple days. At least one full day of workload data should be available for the analysis. When analyzing workloads, users can specify a date range to filter the workload data under consideration. A representative day is selected from this subset of workload data for the analysis. The criteria for selecting a representative day should be flexible. A preferable default assessment of the workload can select the worst day as the representative day based on average utilization. A less conservative assessment

15

may consider the N^{th} percentile (e.g. 95th) day to eliminate outliers. Preferably, the worst days (based on daily average) for each system and for each workload type are chosen as the representative days.

The data extraction process flow for the workload compatibility analysis is shown in FIG. 12. Preferably, the workload data cache 58 includes data obtained during one or more days. For each system 16 in the workload data set, for each workload data type 30, get the workload data for the specified date range, determine the most representative day of data, (e.g. if it is the worst day) and save it in the workload data snapshot. In the result, a snapshot of a representative day of workload data is produced for each system 16.

To evaluate the compatibility of one or more systems with respect to server consolidation, the workloads of the source systems are combined onto the target system. Some types of workload data are normalized for the target system. For example, the % CPU utilization is normalized using the ratio of target and source CPU processing power benchmarks. The consolidated workload for a specific hour in the representative day is approximated by combining the hourly quartile workloads.

There are two strategies for combining the workload quartiles, namely original and cascade. The original strategy simply adds like statistical values (i.e. maximum, third quartile, medians, etc.) of the source systems to the corresponding values of the target system. The cascade strategy processes the statistical values in descending order, starting with the highest statistical value (i.e. maximum value). The strategy adds like statistical values as with original, but may clip the resulting sums if they exceed a configurable limit and cascades a portion of the excess value to the next statistic (i.e. the excess of sum of the maximum values is cascaded to 3rd quartile).

Workload compatibility scores quantify the compatibility of consolidating one or more source systems onto a target system. The scores range from 0 to 100 with higher scores indicating better compatibility. The scores are computed separately for each workload type 30 and are combined with the system configuration and business-related compatibility scores to determine the overall compatibility scores for the systems 16. The workload scores are based on the following: combined system workload statistics at like times and worst case, user-defined workload thresholds, penalty calculation, score weighting factors, and workload scoring formula.

Workloads are assessed separately for two scenarios: like-times and worst case. The like times scenario combines the workload of the systems at like times (i.e. same hours) for the representative day. This assumes that the workload patterns of the analyzed systems are constant. The worst case scenario time shifts the workloads for one or more systems 16 to determine the peak workloads. This simulates the case where the workload patterns of the analyzed systems may occur earlier or be delayed independently. The combined workload statistics (maximum, 3rd quartile, median, 1st quartile and minimum) are computed separately for each scenario.

For a specific analysis, workload thresholds are specified for each workload type. The workload scores are penalized as a function of the amount the combined workload exceeds the threshold. Through the workload type definition, the workload data and corresponding thresholds can be specified independently as percentages or absolute values. The workload data type 30 is specified through the unit property and the threshold data type is specified by the test as percent flag. The common workload/threshold data type permutations are handled as follows.

16

If the workload is expressed as a percentage and test as percent is true (e.g. % CPU), normalize workload percentage using the benchmark and compare as percentages.

If the workload is expressed as an absolute value and test as percent is true (e.g. disk space), convert the workload to a percentage using benchmark and compare as percentages.

If workload unit is expressed as an absolute value and test as percent if false (e.g. network I/O), compare workload value against threshold as absolute values.

A penalty value ranging from 0 to 1 can be calculated for each workload statistic and for each scenario as a function of the threshold and the clipping level. The penalty value is computed as follows:

$$\begin{aligned} & \text{If Workload} \leq \text{Threshold,} \\ & \text{Penalty} = 0 \\ & \text{If Workload} \geq \text{Clipping Level,} \\ & \text{Penalty} = 1 \\ & \text{If Threshold} < \text{Workload} < \text{Clipping Level,} \\ & \text{Penalty} = (\text{Workload Value} - \text{Threshold}) / (\text{Clipping level} - \text{Threshold}) \end{aligned}$$

The workload score is composed of the weighted penalty values. The weights are used to compute the workload score from the penalty values. If the sum of the weights exceeds 1, the weights should be normalized to 1. The actual score is computed for a workload type by subtracting the sum of the weighted penalties from 1 and multiplying the result by 100:

$$\text{Score} = 100 * (1 - \text{Sum}(\text{Weight} * \text{Penalty}))$$

Using the previous example and assuming that the like times are the same as the worst times, the score is calculated as follows:

$$\begin{aligned} \text{Score} &= 100 * (1 - (\text{Weight}_{\text{MaxWorst}} * \text{Penalty}_{\text{MaxWorst}} + \text{Weight}_{\text{MaxLike}} * \\ & \text{Penalty}_{\text{MaxLike}} + \text{Weight}_{\text{Q3Worst}} * \text{Penalty}_{\text{Q3Worst}} + \text{Weight}_{\text{Q3Like}} * \\ & \text{Penalty}_{\text{Q2Like}} + \text{Weight}_{\text{Q2Worst}} * \text{Penalty}_{\text{Q2Worst}} + \text{Weight}_{\text{Q2Like}} * \\ & \text{Penalty}_{\text{Q2Like}} + \\ & \text{Weight}_{\text{Q1Worst}} * \text{Penalty}_{\text{Q1Worst}} + \text{Weight}_{\text{Q1Like}} * \text{Penalty}_{\text{Q1Like}} + \\ & \text{Weight}_{\text{MinWorst}} * \text{Penalty}_{\text{MinWorst}} + \text{Weight}_{\text{MinLike}} * \text{Penalty}_{\text{MinLike}})) \\ &= 100 * (1 - (0.1 * 1 + 0.2 * 1 + 0.3 * 0.5 + 0.4 * 0.5)) \\ &= 30 \end{aligned}$$

A flow chart illustrating a workload compatibility analysis is shown in FIG. 13. When analyzing 1-to-1 workload compatibility, the list of target and source systems 16 is the same. The compatibility is evaluated in two directions, e.g. for Server A and Server B, migrating A to B is considered as well as migrating B to A.

The workload analysis considers one or more workload types, e.g. CPU busy, the workload limits 94, e.g. 75% of the CPU being busy, and the system benchmarks 96, e.g. relative CPU power. Each system 16 in the workload data set is considered as a target (T=1 to N) and compared to each other system 16 in the data set 18 as the source (S=1 to N). The analysis engine 64 first determines if the source and target are the same. If yes, then the workload compatibility score is set to 100 and no additional analysis is required for that pair. If the source and target are different, the system benchmarks are then used to normalize the workloads (if required). The normalized source workload histogram is then stacked on the normalized target system.

System benchmarks can normalize workloads as follows. For systems X and Y, with CPU benchmarks of 200 and 400 respectively (i.e. Y is 2x more powerful than X), if systems X and Y have average CPU utilization of 10% and 15% respectively, the workloads can be normalized through the bench-

17

marks as follows. To normalize X's workload to Y, multiply X's workload by the benchmark ratio X/Y, i.e. $10\% \times 200/400 = 5\%$. Stacking X onto Y would then yield a total workload of $5\% + 15\% = 20\%$. Conversely, stacking Y onto X would yield the following total workload: $10\% + 15\% \times 400/200 = 40\%$.

Using the stacked workload data, the workload compatibility score is then computed for each workload type as described above.

Each source is evaluated against the target, and each target is evaluated to produce an N×N map **32** of scores, which can be sorted to group compatible systems (e.g. see FIG. 4). Preferably, a workload compatibility results is generated that includes the map **32** and workload compatibility scoring details and normalized stacked workload histograms that can be viewed by selecting the appropriate cell **92**. The workload compatibility results are then combined with the rule-based compatibility results to produce the overall compatibility scores, described below.

FIGS. 14 to 17 illustrate a workload settings page **42** which can be used with the analysis program **10** in performing a workload analysis. FIG. 14 illustrates a date settings tab in the settings page **42**. The audit date range specification allows users to choose the appropriate range of workload data to be considered for the analysis. Users can choose data based on the last N days of available data, the last N calendar days, specific date ranges or all available data. An advanced settings page **44** can be launched from the workload settings page **42**. The advanced settings page **44** is shown in FIG. 15.

The advanced settings for workload selection allows users to filter specific days of the week or based on basic weekly or monthly patterns. The specification also lets users exclude outlier days using on percentiles based on the daily average or busiest average hour of the day. Users can also exclude specific hours of the day. After filtering undesired days of workload, users can finally choose a representative day based on the busiest, least busy, typical or average values. Users can also choose a predicted workload in the future based on an expected growth rate or based on projected trends to some date in the future.

FIG. 16 illustrates a limits tab accessed from the workload settings page **42**. The analysis program **10** allows user to specify workload limits when evaluating the workload types to be analyzed. These limits are used when computing the workload scores.

FIG. 17 illustrates a parameters tab accessed from the workload settings page **42**. The analysis program **10** allows users to specify workload type specific parameters. For example, the virtual CPU utilization can be used to model the virtualization overhead based on CPU utilization, disk I/O rates and network I/O rates. The confidence limit value can range between 0 and 100% and allows users to adjust the workload computation based on the probability of outcomes when combining the workload of multiple systems. A confidence limit of 100% indicates that the workload computation is based on the worst case scenario where the maximum values of every system **16** are assumed to coincide. A 99% confidence limit effectively discards 1% of the worst possible cases, resulting in less conservative workload stacking results. The strategy name specifies the workload scoring strategy to employ when computing the workload score. Pre-defined scoring strategies such as Peak and Sustained emphasize peak (maximum) and sustained (third quartile) workloads, respectively. Peak scoring is useful for performance sensitive applications whose performance should not be degraded. Sustained scoring is appropriate for less perfor-

18

mance sensitive applications such as batch jobs where slight performance degradations may be acceptable.

The results of the rule and workload compatibility analyses are combined to compute an overall compatibility score for each server pair. These scores preferably range from 0 to 100, where higher scores indicate greater compatibility and 100 indicating complete or 100% compatibility.

As noted above, the analysis input can include importance factors. For each rule set **28** and workload type **30** included in the analysis, an importance factor **88** can be specified to adjust the relative contribution of the corresponding score to the overall score. The importance factor **88** is an integer, preferably ranging from 0 to 10. A value of 5 has a neutral effect on the contribution of the component score to the overall score. A value greater than 5 increase the importance whereas a value less than 5 decreases the contribution.

The overall compatibility score for the system pair is computed by combining the individual compatibility scores using a formula specified by an overlay algorithm which performs a mathematical operation such as multiply or average, and the score is recorded.

Given the individual rule and workload compatibility scores, the overall compatibility score can be calculated by using the importance factors as follows for a "multiply" overlay:

$$O = 100 * \frac{100 - (100 - S_1) * \frac{F_1}{5}}{100} * \frac{100 - (100 - S_2) * \frac{F_2}{5}}{100} * \dots * \frac{100 - (100 - S_n) * \frac{F_n}{5}}{100}$$

where O is the overall compatibility score, n is the total number of rule sets **28** and workload types **30** included in the analysis, S_i is the compatibility score of the i^{th} rule set **28** or workload type **30** and F_i is the importance factor of the i^{th} rule set **28** or workload type **30**.

It can be appreciated that setting the importance factor **88** to zero eliminates the contribution of the corresponding score to the overall score. Also, setting the importance factor to a value less than 5 reduces the score penalty by 20% to %100 of its original value.

For example, a compatibility score of 90 implies a score penalty of 10 (i.e. $100 - 90 = 10$). Given an importance factor of 1, the adjusted score is 98 (i.e. $100 - 10 * 1/5 = 100 - 2 = 98$). On the other hand, setting the importance factor to a value greater than 5 increases the score penalty by 20% to 100% of its original value. Using the above example, given a score of 90 and an importance factor of 10, the adjusted score would be 80 (i.e. $100 - 10 * 10/5 = 100 - 20 = 80$).

If more systems **16** are to be examined, the above process is repeated. When overall compatibility analysis scores for all server pairs have been computed, a map **32** is displayed graphically and each cell **92** is linked to a scorecard that provides further information. The further information can be viewed by selecting the cell **92**. A sorting algorithm is then preferably executed to configure the map **32**. The maps **32** can be sorted in various ways to convey different information. For example, sorting algorithms such as a simple row sort, a simple column sort and a sorting by group can be used.

A simple row sort involves computing the total scores for each source system (by row), and subsequently sorting the rows by ascending total scores. In this arrangement, the highest total scores are indicative of source systems that are the

best candidates to consolidate onto other systems. A simple column sort involves computing the total scores for each target system (by column) and subsequently sorting the columns by ascending total score. In this arrangement, the highest total scores are indicative of the best consolidation target systems. Sorting by group involves computing the difference between each system pair, and arranging the systems to minimize the total difference between each pair of adjacent systems in the map. The difference between a system pair can be computed by taking the square root of the sum of the squares of the difference of a pair's individual compatibility score against each other system in the analysis. In general, the smaller the total difference between two systems, the more similar the two systems with respect to their compatibility with the other systems. The group sort promotes the visualization of the logical breakdown of an environment by producing clusters of compatible systems **18** around the map diagonal. These clusters are indicative of compatible regions in the environment **12**. In virtualization analysis, these are often referred to as "affinity regions."

The high level process flow of the multi-dimensional compatibility analysis is illustrated in FIG. **18**. In addition to the common compatibility analysis input, this analysis takes a consolidation solution as input. In contrast to the 1-to-1 compatibility analysis that evaluates the compatibility of each system pair, this multi-dimensional compatibility analysis evaluates the compatibility of each transfer set **23** specified in the consolidation solution.

The multi-dimensional compatibility analysis extends the original 1-to-1 compatibility analysis that assessed the transfer of a single source entity to a target. As with the 1-to-1 compatibility analysis, the multi-dimensional analysis produces an overall compatibility scorecard **98** based on technical, business and workload constraints. Technical and business compatibility are evaluated through one or more rule sets **28**. Workload compatibility is assessed through one or more workload types **30**.

This produces multi-dimensional compatibility analysis results, which includes multi-dimensional compatibility scores, maps and details based on the proposed transfer sets **23**.

For each transfer set **23**, a compatibility score is computed for each rule set **28** and workload type **30**. An overall compatibility score for the transfer set **23** is then derived from the individual scores.

In addition to evaluating the compatibility of the specified transfer sets, the compatibility analysis can evaluate the incremental effect of adding other source systems (specified in the analysis input) to the specified transfer sets. Similar to the 1-to-1 compatibility analysis, this analysis involves 4 stages. The first stage is gets the system data **18** required for the analysis to produce the analysis data snapshot. The second stage performs a multi-dimensional compatibility analysis for each rule set **28** for each transfer set **23**. Next, the workload compatibility analysis is performed for each workload type **30** for each transfer set **23**. Finally, these analysis results are combined to determine overall compatibility of each transfer set. The multi-dimensional rule-based compatibility analysis differs from the 1-to-1 compatibility analysis since a transfer set can include multiple sources (N) to be transferred to the target, 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) as will be explained in greater detail below. The multi-dimensional workload and overall compatibility analysis algorithms are analogous to their 1-to-1 analysis counterparts.

To assess the compatibility of transferring multiple source entities (N) to a target (1), the rule-based analysis can compute a compatibility score based on a combination of N-to-1 and N-by-N compatibility analyses. An N-to-1 intercompatibility analysis assesses each source system against the target. An N-by-N intracompatibility analysis evaluates each source system against each of the other source systems. This is illustrated in a process flow diagram in FIG. **19**.

Criteria used to choose when to employ an N-to-1, N-by-N or both compatibility analyses depend upon the target type (concrete or malleable), consolidation strategy (stacking or virtualization), and nature of the rule item.

Concrete target models are assumed to be rigid with respect to their configurations and attributes such that source entities to be consolidated are assumed to be required to conform **8** to the target. To assess transferring source entities onto a concrete target, the N-to-1 inter-compatibility analysis is performed. Alternatively, malleable target models are generally adaptable in accommodating source entities to be consolidated. To assess transferring source entities onto a malleable target, the N-to-1 inter-compatibility analysis can be limited to the aspects that are not malleable.

When stacking multiple source entities onto a target, the source entities and targets coexist in the same operating system environment. Because of this inherent sharing, there is little flexibility in accommodating individual application requirements, and thus the target is deemed to be concrete. As such, the multi-dimensional analysis considers the N-to-1 inter-compatibility between the source entities and the target as the primary analysis mechanism, but, depending on the rule sets in use, may also consider the N-by-N intracompatibility of the source entities amongst each other.

When virtualizing multiple source entities onto a target, the source entities are often transferred as separate virtual images that run on the target. This means that there is high isolation between operating system-level parameters, and causes virtualization rule sets to generally ignore such items. What is relevant, however, is the affinity between systems at the hardware, storage and network level, and it is critical to ensure that the systems being combined are consistent in this regard. In general, this causes the multi-dimensional analysis to focus on the N-to-N compatibility within the source entities, although certain concrete aspects of the target systems (such as processor architecture) may still be subjected to (N-to-1) analysis.

N-to-1 intercompatibility scores reflect the compatibility between N source entities and a single target as defined by a transfer set **23** as shown in FIG. **20**. This analysis is performed with respect to a given rule set and involves: 1) Separately evaluate each source entity against the target with the rule set to compile a list of the union of all matched rule items; 2) For each matched rule item, use the rule item's mutex (mutually exclusive) flag to determine whether to count duplicate matched rule items once or multiple times; and 3) Compute the score based on the product of all the penalty weights associated with the valid matched rule items:

$$S=100*(1-w_1)*(1-w_2)*(1-w_3)*\dots*(1-w_n);$$

where S is the score and w_i is the penalty weight of the i^{th} matched item.

N-by-N intracompatibility scores reflect the compatibility amongst N source entities with respect to a given rule set as shown in FIG. **21**. This analysis involves: 1) Separately evaluate each source entity against the other source entities with the rule set to compile a list of the union of all matched rule items; 2) For each matched rule item, use the rule item's mutex (mutually exclusive) flag to determine whether to count

duplicate matched rule items once or multiple times; and 3) Compute the score based on the product of all the penalty weights associated with the valid matched rule items:

$$S=100*(1-w_1)*(1-w_2)*(1-w_3)*\dots*(1-w_n);$$

where S is the score and w_i is the penalty weight of the i^{th} matched item.

A procedure for stacking the workload of multiple source systems on a target system is shown in FIG. 22. The multi-stacking procedure considers the workload limits that is specified using the program 150, the per-system workload benchmarks (e.g. CPU power), and the data snapshot containing the workload data for the source and target systems 16 that comprise the transfer sets 23 to analyze. The analysis may evaluate transfer sets 23 with any number of sources stacked on a target for more than one workload type 30.

For each workload type 30, each transfer set 23 is evaluated. For each source in the transfer set 23, the system benchmarks are used to normalize the workloads as discussed above, and the source workload is stacked on the target system. Once every source in the set is stacked on the target system, the workload compatibility score is computed as discussed above. The above is repeated for each transfer set 23. A multi-stack report may then be generated, which gives a workload compatibility scorecard for the transfer sets along with workload compatibility scoring details and normalized multi-stacked workload charts.

The consolidation analysis process flow is illustrated as D in FIG. 5. Using the common compatibility analysis input and additional auto fit inputs, this analysis seeks the consolidation solution that maximizes the number of transfers while still fulfilling the several pre-defined constraints. The consolidation analysis repeatedly employs the multi-dimensional compatibility analysis to assess potential transfer set candidates. The result of the consolidation analysis comprises of the consolidation solution and the corresponding multi-dimensional compatibility analysis.

A process flow of the consolidation analysis is shown in FIG. 23.

The auto fit input includes the following parameters: transfer type (e.g. virtualize or stacking), minimum allowable overall compatibility score for proposed transfer sets, minimum number of source entities to transfer per target, maximum number of source entities to transfer per target, and quick vs. detailed search for the best fit. Target systems can also be designated as malleable or concrete models.

As part of a compatibility analysis input specification, systems can be designated for consideration as a source only, as a target only or as either a source or a target. These designations serve as constraints when defining transfers in the context of a compatibility analysis. The analysis can be performed on an analysis with pre-existing source-target transfers. Analyses containing systems designated as source or target-only (and no source or target designations) are referred to as “directed analysis.”

The same transfer type may be assumed for all automatically determined transfers within an analysis. The selected transfer type affects how the compatibility analysis is performed. The minimum overall compatibility score dictates the lowest allowable score (sensitivity) for the transfer sets to be included in the consolidation solution. Lowering the minimum allowable score permits a greater degree of consolidation and potentially more transfers. The minimum and maximum limits for source entities to be transferred per target (cardinality) define additional constraints on the consolidation solution. The quick search performs a simplified form of the auto fit calculation, whereas the detailed search performs

a more exhaustive search for the optimal solution. This distinction is provided for quick assessments of analyses containing a large numbers of systems to be analyzed.

The transfer auto fit problem can be considered as a significantly more complex form of the classic bin packing problem. The bin packing problem involves packing objects of different volumes into a finite number of bins of varying volumes in a way that minimizes the number of bins used. The transfer auto fit problem involves transferring source entities onto a finite number of targets in a way that maximizes the number of transfers. The basis by which source entities are assessed to “fit” onto targets is based on the highly nonlinear compatibility scores of the transfer sets. As a further consideration, which can increase complexity, some entities may be either source or targets. The auto fit problem is a combinatorial optimization problem that is computationally expensive to solve through a brute force search of all possible transfer set permutations. Although straightforward to implement, this exhaustive algorithm is impractical due to its excessive computational and resource requirements for medium to large data sets. Consequently, this class of problem is most efficiently solved through heuristic algorithms that yield good but likely suboptimal solutions.

There are four variants of the heuristic auto fit algorithm that searches for the best consolidation solution:

Quick Stack—quick search for a stacking-based consolidation solution;

Detailed Stack—more comprehensive search for a stacking-based consolidation solution;

Quick Virtualization—quick search for a virtualization-based consolidation solution; and

Detailed Virtualization—more comprehensive search for a virtualization-based consolidation solution.

The auto fit algorithms are iterative and involve the following common phases:

The initial phase filters the source and target lists by eliminating invalid entity combinations based on the 1-to-1 compatibility scores that are less than the minimum allowable compatibility score. It also filters out entity combinations based on the source-only or target-only designations. The auto fit algorithm search parameters are then set up. The parameters can vary for each algorithm. Example search parameters include the order by which sources and targets are processed and the criteria for choosing the best transfer set 23. The next phase compiles a collection of candidate transfer sets 23 from the available pool of sources and targets. The candidate transfer sets 23 fulfill the auto fit constraints (e.g. minimum allowable score, minimum transfers per transfer set, maximum transfers per transfer set). The collection of candidate transfer sets may not represent a consolidation solution (i.e. referenced sources and targets may not be mutually exclusive amongst transfer sets 23). The algorithms vary in the criteria employed in composing the transfer sets. In general, the detailed search algorithms generate more candidate transfer sets than quick searches in order to assess more transfer permutations.

The next phase compares the candidate transfer sets 23 and chooses the “best” transfer set 23 amongst the candidates. The criteria employed to select the best transfer set 23 varies amongst the algorithms. Possible criteria include the number of transfers, the compatibility score, general compatibility of entities referenced by set and whether the transfer set target is a target-only.

Once a transfer set is chosen, it is added to the intermediate consolidation solution. The entities referenced by the transfer set are removed from the list of available sources and targets

and the three preceding phases are repeated until the available sources or targets are consumed.

Once all the sources or targets are consumed or ruled out, the consolidation solution is considered complete and added to a list of candidate solutions. Additional consolidation solutions can be compiled by iterating from the second phase with variations to the auto fit parameters for compiling and choosing candidate transfer sets. The criteria used to stop compiling additional solutions can be based on detecting that the solution is converging on a pre-defined maximum number of iterations. Finally, the best candidate consolidation solution can be selected based on some criteria such as the largest reduction of systems with the highest average transfer set scores. The general algorithm is shown in the flow diagram depicted in FIG. 24.

Accordingly, the compatibility and consolidation analyses can be performed on a collection of system to 1) evaluate the 1-to-1 compatibility of every source-target pair, 2) evaluate the multi-dimensional compatibility of specific transfer sets, and 3) to determine the best consolidation solution based on various constraints including the compatibility scores of the transfer sets. Though these analyses share many common elements, they can be performed independently. These analyses are based on collected system data related to their technical configuration, business factors and workloads. Differential rule sets and workload compatibility algorithms are used to evaluate the compatibility of systems. The technical configuration, business and workload related compatibility results are combined to create an overall compatibility assessment. These results are visually represented using color coded scorecard maps.

It will be appreciated that although the system and workload analyses are performed in this example to contribute to the overall compatibility analyses, each analysis is suitable to be performed on its own and can be conducted separately for finer analyses. The finer analysis may be performed to focus on the remediation of only configuration settings at one time and spreading workload at another time. As such, each analysis and associated map may be generated on an individual basis without the need to perform the other analyses.

It will be appreciated that each analysis and associated map discussed above may instead be used for purposes other than consolidation such as capacity planning, regulatory compliance, change, inventory, optimization, administration etc. and any other purpose where compatibility of systems is useful for analyzing systems 16. It will also be appreciated that the program 10 may also be configured to allow user-entered attributes (e.g. location) that are not available via the auditing process and can factor such attributes into the rules and subsequent analysis.

It will further be appreciated that although the examples provided above are in the context of a distributed system of computer servers, the principles and algorithms discussed are applicable to any system having a plurality of sub-systems where the sub-systems perform similar tasks and thus are capable theoretically of being consolidation. For example, a local network having a number of personal computers (PCs) could also benefit from a consolidation analysis.

Power Utilization Analysis

It has also been recognized that the analysis program 10 can be used to estimate the power utilization of existing source and proposed target servers to compare the power utilization before and after the transformation. This information can be very useful with the high cost of energy, more power hungry servers and the power and cooling constraints of data centers.

If no actual server power utilization data is available, the analysis program 10 estimates the power for each server based the server utilization level, the estimated power at idle and power at maximum utilization.

The power utilization of servers can be analyzed as a workload type. This is especially useful when comparing the aggregate power utilization of a set of servers before and after consolidation.

While some modern server models support the measurement of its power utilization, the majority of servers do not support its measurement. As a result, the analysis program must estimate power utilization. The power utilization is computed according to the server load, the power consumption at idle and maximum loads.

The server load can be approximated through server activity such as CPU, memory and disk activity. The power consumption at idle and maximum loads can be measured empirically or through various power calculators provided by server vendors.

When estimating the power utilization as a function of the server load, a simplifying assumption could be to assume a linear relationship between the server load and power consumption. Thus, if the server load is zero, the power consumption is equal to the estimated power level corresponding to idle load. Similarly, if the server load is at 100%, the power consumption is equal to the estimated power at maximum load. Finally, if the server load is between 0 and 100%, it is estimated based on a linear relationship between the idle and maximum power loads.

$$\text{Estimated Power} = \text{Idle Power} + \text{Pct. Server Load} * (\text{Max Power} - \text{Idle Power})$$

For example, assume the estimated power utilization of a server at idle and maximum loads are 300 and 600 watts, respectively. If the server is at 50% load, the power utilization would be estimated as 450 watts.

$$\text{Power}@50\% = 300 + 50\% * (600 - 300) = 450 \text{ watts}$$

Transformational P2V Analytics Using Analysis Program

FIG. 25 provides further detail for the process 99 shown in FIG. 2 to illustrate conceptually the various steps that may be performed in designing the virtual environment 21. In general, the analysis process 99 begins with the gathering of highly detailed configuration and workload data from systems 16 in an existing physical environment 12. The systems 16 of interest for this data acquisition include the systems to be virtualized as well as those that may be converted to virtual hosts (target servers running hypervisor software) to form part of the new virtual environment 21. The analysis program 20 can be used to automate the data collection from the systems 16, using either agent-based or agentless means, in order to ensure that all analyses are based on up-to-date data. This data is combined with business attributes and process-related information related to the systems 16 to form a complete set of analysis inputs.

From this collected data, the current asset assessment 102 utilizes virtualization rule sets 11 to identify the physical systems 16 that can be converted into virtual hosts, allowing existing systems to be repurposed as virtual servers (e.g. ESX Servers for VMware®) without buying new hardware. The virtualization host system definition 104 can also estimate the aggregate resource capacity of the existing server hardware and compare it against the expected resource requirements of the virtual environment 21. This allows analysts to specify hypothetical server models 125 (i.e. candidates for purchase) that can be used to make up the shortfall.

The analysis program 10 may then group the target system candidates based on live migration compatibility or other

logical grouping criteria, which defines the clusterable pools of systems **16** from which the new virtual environment **21** will be constructed. The physical environment analysis **100**, as discussed above, evaluates technical, business and workload constraints against the systems **16** to be virtualized using advanced rule sets and workload algorithms. The constraints are combined to produce an overall system affinity map that indicates the systems which should be kept together and which ones should be separated when they are virtualized.

The virtual environment optimization **106** determines the optimal mapping of physical servers onto virtual environments and clusters, and allows for “what-if” analyses to determine the optimal cluster design that minimizes the server count while addressing the server and virtual machine compatibility constraints. The resulting analysis maps define the cluster memberships of the servers and virtual machines as well as affinity and anti-affinity rules (e.g. DRS in VMware™).

The automated generation of the cluster membership design accelerates the implementation of virtualized environments **21** while at the same time reducing risk in implementation and subsequent operation. After the virtualized environment **21** is deployed, the analysis program **10** and virtualization UI **13** can be used to provide decision support for ongoing management **15** by gathering configuration and workload data from the virtualization hosts and virtual machines on an ongoing basis and using this to both track the environments as well as enable further analysis and optimization as servers and constraints change over time. Further detail regarding the ongoing management **15** will be provided later.

As can be seen in FIG. **25**, the physical environment analysis **100** comprises individual constraint analyses related to technical, business and workload constraints that affect virtualization and consolidation strategies and an overall combined constraint analysis using the individual constraint analyses.

A technical constraint analysis is performed by the analysis program **10** using technical constraint rulesets. Technical constraints are constraints that affect “what can go together”, and typically include configuration-oriented elements such as version compatibilities, environmental settings, patch requirements, security configurations, etc. In a virtualization analysis, the technical constraint models employed typically focus on virtual host and live migration compatibilities, storage configurations, unsupported or non-standard hardware, network connectivity, and other considerations that may impact the viability of and/or path to virtualization. The technical analysis identifies the physical systems that can be virtualized by considering virtualized system constraints including guest operating system support, maximum limits on virtual processors, memory and swap. In addition, the analysis highlights constraints that can impact the compatibility of virtualized systems including unique legacy devices, and uncommon network connectivity or storage requirements. The technical constraint analysis also evaluates the sameness of guest system images to assess the potential to take advantage of the virtualization package’s transparent page sharing capabilities (if applicable). The resulting technical affinity map illustrates groups of systems that must be kept together or apart, as well as groups that are ideally kept together or apart.

In general, guest candidates (i.e. those being considered for conversion to virtual machines) must be physical systems **16** and not already virtual machines. The technical constraint analysis should check for potentially incompatible hardware such as fax boards, token ring cards etc. There are various

technical constraints for guest candidates that are hypervisor-specific. For example: ensuring that the operating system is supported by the hypervisor; or constraints based on OS type, OS version, kernel bits and service pack/patch levels, e.g. Microsoft® Hyper-V 1.0 supports the following server operating systems as guests: Windows® Server 2008 x86 and x64, Windows® Server 2003 x86 and x64 SP2, Windows® Server 2000 SP4, and SUSE® Linux Enterprise Server 10 SP1/2 x64 and x86. Another constraint may be guest resource configuration limits such as maximum memory, virtual processors, number of network interfaces etc. Other hypervisor-specific constraints can be based on hypervisor-specific P2V rulesets, e.g. rulesets for VMware® ESX, Microsoft® Hyper-V and Citrix® XenServer.

There are also various server affinity considerations that should be made during the technical constraint analysis, including checking for network affinity (i.e. servers with common networking configurations are more suited to be clustered) and checking for network communications (i.e. servers that communicate with each other may be suited to run on the same host to take advantage of lower network latency).

A business constraint analysis is performed by the analysis program using business constraint rulesets. Business constraints are more concerned with “what should go together”, both from a business and a process perspective. Criteria such as maintenance windows, system availability targets, application owners, locations, departments, and other non-technical criteria are analyzed to ensure that there is consistency in the virtual environment and to prevent any production problems post-virtualization. This analysis focuses on business factors that impact the compatibility of the source systems. Other factors considered may include such things as service chargeback models, service levels and regulatory requirements. As with the technical constraint analysis, the business affinity map can be generated that reflects groups of systems to keep together or apart. The business constraints are typically used to organize the guest virtual machines into affinity groups, e.g. group systems from the same department, service level, environment etc. It may be noted that the business constraints can also be used to disqualify certain systems, e.g. do not virtualize systems from specific locations, departments etc.

A workload constraint analysis is based on workload constraints, answers the question “what fits together”, and looks at the utilization levels and patterns of servers to determine what the optimal combinations may be (both onto existing hardware as well as new servers). The workload analyses that can be performed by the analysis program **10** uses quartile-based representations of CPU, disk I/O, network I/O, memory utilization and disk utilization in order to build out a comprehensive scorecard-based view of the workload affinities in an environment. The workload analysis evaluates the combination of one or more source workloads onto the target servers to evaluate processor utilization, memory, disk I/O and network I/O. The analysis employs the workload normalization and virtualization overhead models described below to predict workloads with better accuracy. The workload analysis can consider sustained and peak system workloads at like times and at offset times to consider the normal and worst case scenarios. Workload analysis parameters can be specified to adjust the conservativeness or aggressiveness of the constraints. In general, systems with lower workloads are better virtualization candidates than those with very high workloads and resource requirements.

When analyzing workloads, an analyst can specify various configuration parameters including resource thresholds on

target systems to define desired workload headroom; scoring strategies to emphasize the importance of peak vs. sustained workloads as well as analyzing workload based on like times or offset times; workload contention confidence limits allows analyst to adjust risk tolerance related to likelihood of peak workload contention among multiple systems; and workload data date range, filters, trends or assumed growth rates. In addition, the CPU utilization of the virtualized server can be better estimated with a virtualization overhead model based on measured physical CPU utilization, disk I/O and network I/O rates. CPU utilization can be normalized between different server models using processor benchmarks. Different processor benchmarks can be employed, depending on the personality of the system workload. Examples of processor benchmarks that may be employed include CINT2000 and CINT2006 rate from SPEC (Standard Performance Evaluation Corporation).

Capacity planning for high available clusters can be readily performed through a what-if workload analysis by adjusting the workload headroom thresholds or excluding target servers from the analysis to simulate host failures.

As discussed above, a compatibility analysis performed by the analysis program 10 can generate a compatibility map 32 as shown in FIG. 4. FIG. 30 illustrates a compatibility map 164 showing the result of applying a virtualization rule set 11 against a set of physical systems 16. As per FIG. 4, the systems 16 are listed in the map 164 along the left side of the matrix as well as along the top of the matrix thus producing a cross-correlation of the compatibilities of the listed systems 16. In this example, it can be appreciated that the similarly shaded regions comprising a score 36 of “100” and normally shaded 34 green (as identified by the circle 166 in FIG. 30), represent affinity regions where the systems 16 are generally self-consistent. Those regions showing as darker or lighter than those in the circle 166 (typically yellow, orange, red etc.), on the other hand, represent system combinations where important constraints may be violated if they are virtualized onto the same infrastructure. The set of four systems 16 to the far right and bottom in this example are hypothetical targets that the environment is being analyzed onto. Similar maps 164 can be generated for technical, business and workload constraints, which are then used to conduct a combined constraint analysis.

A combined constraint analysis looks at the net-effect combining the technical, business and workload constraints to provide an overall affinity map. The analysis program 10 can analyze multiple constraint maps using a 3-dimensional data structure as illustrated conceptually in FIG. 31 that enables simultaneous assessment of all constraints. The overall affinity map defines regions of compatible source systems that can be assigned to common clusters. The compatibility scores would then reflect the degree of compatibility/incompatibility between systems 16.

Turning back to FIG. 25, the current asset assessment 102 generally comprises the steps of a server upgrade analysis and an aggregate server utilization analysis.

The server upgrade analysis assesses the viability of repurposing existing physical servers to serve as virtualization hosts (i.e. to run hypervisor software). This analysis can involve checking to see if hardware is compatible with specific hypervisor software (some hypervisors such as VMware® ESX support specific hardware server manufacturers and models) and checking whether a system 16 has sufficient resources (CPU, memory, network interfaces, etc.) to support virtualization software and guests. The analysis may assume that hardware in the existing system 16 can be upgraded to meet certain hardware requirements (e.g.

memory, HBA, network interface cards). The upgrade analysis can be performed by creating an analysis comprised of virtualization host candidates and applying the applicable hypervisor-specific host compatibility rule set 11 (e.g. VMware® ESX Hardware Compatibility). In general, the host compatibility rule set rules out any servers that are fully incompatible (e.g. unsupported processor architecture, virtual machine, etc.), and applies varying penalty levels based on correctable and less severe incompatibilities (e.g. insufficient memory, number of network adapters, etc.). In addition, the upgrade analysis can validate the various target server and live migration requirements such as minimum CPU clock speeds, maximum RAM, maximum CPU cores and multiple GB Ethernet network interface cards.

FIG. 32 illustrates a target compatibility map 170 showing which systems 16 in the current physical environment 12 are candidates for upgrading to run as a target (with hypervisor software). The large region 172 shown in FIG. 32 identifies systems 16 that are unconditionally supported (with “100” and normally shaded green), and the lighter regions (normally yellow) show those systems 16 that can become target servers with some qualifications.

The aggregate server utilization analysis combines resource utilization data of physical server source candidates to obtain a high level estimate of aggregate resource requirements of the target server environment. This analysis also determines whether existing physical servers are sufficient to support virtualization resource requirements or whether new servers need to be acquired to meet virtualization requirements; determines storage requirements (storage area networks—e.g. SAN); and determines network bandwidth requirements. Important system resources for sizing target servers are CPU and memory utilization, storage and disk and network I/O rates. The aggregate resource utilization of the source candidates is compared against the capacity of the target candidates to thus determine the additional server hardware, if any, that is required to support the planned virtualized environment 21.

To accurately combine the processor utilization of the systems based on different processors, industry benchmarks can be employed by the analysis program 10 to normalize the CPU workload data. Processor benchmarks such as SPEC CINT2000 or SPEC CINT2006 rate are better suited than basic processor speeds (MHz) since benchmarks account for different processor architectures that affect performance significantly. The analysis program 10 can be configured to use a variety of comprehensive CPU benchmark tables to determine the appropriate benchmark value of the physical systems based on the server model, processors and type of workload (e.g. CPU intensive, web, Java application, database, etc.).

As an additional software layer, virtualization software such as VMware® often adds a performance overhead. As such, when modeling the resource utilization of physical systems in the virtualized target environment, a virtualization overhead is added to the source system workloads. The analysis program 10 can use an advanced virtualization overhead model to estimate CPU utilization of physical systems when virtualized on a virtualization host. The CPU overhead is modeled for each guest as a function of the CPU utilization, network I/O and disk I/O rates. Similarly, the memory overhead is comprised of the service console memory (e.g. default 272 MB, maximum 800 MB) and guest system contributions. The memory overhead of each guest system is affected by its memory allocation, the number of virtual CPUs, and whether it is a 32 or 64 bit operating system. It may be noted that the memory overhead of similar virtualized systems can be offset

by the memory saved through features such as the transparent page sharing feature provided by VMware®.

By normalizing workloads and accounting for virtualization overhead, the projected resource requirements of the physical systems can be modeled with higher accuracy. The aggregate resource requirements are adjusted further to include the desired headroom to account for future growth and high availability requirements. Similarly, the aggregate resource capacity of the virtualization host candidates can be calculated by the analysis program 10.

FIG. 33 illustrates an aggregate utilization UI 174, showing the normalized utilization of an entire environment. Utilization is reported in this example as a rolled-up average 176 as well as a time-of-day curve 178 showing peak and sustained activity throughout the daily cycle. This can be an important measure of the consolidation potential of an environment, and gives an initial estimate of the CPU, I/O, disk and network capacity required in the virtualization host environment.

The virtualization host system definition 104 generally comprises a determination of hypothetical server models and live migration compatibilities.

Hypothetical servers can be used to model target servers that do not currently exist in the computing environment, which allows users to evaluate a wide range of scenarios. Predefined hypothetical servers are based on popular server models with typical hardware configurations (processor type, number of processors, memory, storage, network adapters). Analysts can define custom server models with specific hardware configurations. Hypothetical servers can be based on sparse models (hardware and operating systems configurations) and can also be based on more detailed models derived from existing servers. The projected aggregate workloads of the source and target systems are compared to determine whether additional computing resources are required. If there is insufficient capacity, the amount of hypothetical virtualization host hardware is estimated.

FIG. 27 illustrates an exemplary process flow diagram for determining initial high-level requirements regarding hypothetical server models 125. It can be seen that the process begins with guest candidates 118, 120 (see also FIG. 26 to be explained below) and virtualization host candidates 122, 124. At 132, the aggregate system resource requirements are estimated based on the historical workload of the candidates 118, 120 thereby producing aggregate workload requirements at 134. At 136, the aggregate system capacity is estimated based on hardware configurations of the virtualization host candidates thereby producing a measure of aggregate workload capacity 138. The aggregate workload requirements 134 and aggregate workload capacity 138 may then be compared at 140 to determine if there is sufficient capacity at 142 based on the proposed virtualization solution. If not, hypothetical server models are added at 144 to the virtualization host candidates to meet the workload requirements thereby generating the appropriate hypothetical server models 125. If the capacity is sufficient to meet the requirements, the process ends at 146.

Live migration compatibility can be assessed for hypervisors that support live migration of guest virtual machines between virtualization hosts that are part of the same cluster. Examples of live migration include VMotion for VMware® ESX and XenMotion for Citrix® XenServer®. This analysis assesses compatibility of existing and or hypothetical virtualization host candidates to determine which set of target hosts can be grouped into a cluster that supports live migration. An important aspect of live migration compatibility between virtualization hosts is processor architecture compatibility. The live migration analysis can be performed by

creating an analysis comprised of the virtualization hosts only, and applying the appropriate VM migration compatibility map ruleset (e.g. VMotion® Compatibility Map). The resulting map defines regions of compatible virtualization hosts.

FIG. 34 illustrates a live migration compatibility map 180 showing the sets 182 (example identified by circle in FIG. 34) of servers that are compatible from a live migration perspective. This can be an important step in defining a go-forward environment since many incompatibilities exist between server platforms including those from the same manufacturer. Since clusters rely on the live migration software, the map 180 effectively sculpts out the pools of servers from which clusters can be built.

The virtual environment optimization 106 analyzes the virtualization candidates and virtualization hosts to determine recommended cluster configurations, cluster memberships of guest systems and affinity/anti-affinity rules. The analysis program 10 can be used to employ heuristic optimization algorithms (referred to above as the auto-fit process) to automatically determine the virtualization solution that eliminates the largest number of systems 16 with the highest set of compatibility scores. Additional what-if scenarios can be readily modeled by modifying constraints, adding systems etc. to the analysis. As can be seen in FIG. 25, the virtual environment optimization 106 performs a multi-dimensional analysis, e.g. according to the processes described in FIGS. 18 to 24.

The multi-dimensional analysis employs the auto-fit analysis to determine the optimal layout of the source systems onto the target systems based on the technical, business and workload constraints. The analysis considers the combined constraint and affinity analysis of the physical source systems with the existing and hypothetical target systems. If live migration is to be supported, the target systems included in the auto-fit analysis should be compatible with respect to live migration. The optimization criteria can be based on searching for a solution that minimizes the number of target servers required to accommodate the source systems, or a solution that attempts to balance the load across a specific number of target servers. An example of the virtual environment optimization 106 will be provided later.

The end result of the transformational P2V analysis 99 is the virtual environment design 110, which provides the blueprint for creating a new virtual environment 21 or, as explained below, to refine or upgrade an existing virtual environment 21. The virtual environment design 110 comprises a cluster membership design, an affinity rule design and avirtualization management framework API integration as shown in FIG. 25.

Most virtualization technologies support grouping of the target hosts into a cluster thus the implementation of a cluster membership design. Within a cluster, guest virtual machines may then be migrated between target hosts. The VM-cluster assignments can be constrained by the clusterability of the target servers, the affinity of the source systems, workload requirements of the source systems and resource capacities of the target servers. The virtual environment optimization 106 considers all these constraints and recommends the placement of the source systems on the set of clusterable targets. Additional considerations for defining clusters are: the maximum allowable servers per cluster, the sharing of common storage and networking in clusters, the similarity of hardware specifications among the servers in the cluster and sharing common resources (e.g. blade servers are suitable for this reason). The virtualization rule sets 11 enable the analysis program 10 to account for many of the above considerations

31

and the optimal cluster size is typically considered to decide between when to make a separate cluster and when to employ affinity and anti-affinity rules within a single cluster.

FIG. 39 shows a cluster-based view 204 of a set of guest OSs. In FIG. 39, the larger areas of non-zero scores (i.e. non-dark) represent recommended cluster membership and in this example, there are 5 distinct clusters 206 emerging from the analysis. The clusters 206 may be separated by different colours (not shown) to identify anti-affinity regions within a cluster and the appropriate rules can then be generated by the analysis program 10 to ensure that constraints are honoured at runtime.

The affinity rule design is performed to specify which systems should be assigned to the same clusters 206. For virtualization technologies that support the migration of virtual machines among target hosts, there are cases where it is better to keep certain virtual machines together or apart. For example, a virtual machine that serves an active backup/failover for another virtual machine should not be deployed on the same target host (anti-affinity). Similarly, there are cases when virtual machines that transfer a high volume of data with each other may be optimally deployed on the same target host to reduce network latency in their communications (affinity). The affinity and anti-affinity rules are based on the technical and business compatibility analysis scores among the source systems. In general, systems with high compatibility scores can co-exist on the same target host while systems with very poor compatibility scores should be kept apart.

Most virtualization technologies support varying levels of integration with third-party applications to improve the management of the virtual environment. Some virtualization technologies support a mechanism to balance the load among virtualization hosts in a cluster. This is accomplished monitoring the current workload of the virtual machines and respective hosts and automatically moving virtual machines from heavily loaded hosts to less busy hosts as required. For example, VMware® Virtual Center supports DRS which provides such functionality. VMware® DRS also supports affinity and anti-affinity rules that allow users to define which virtual machines should be kept together and apart when virtual machines are automatically migrated. Based on the VM affinity rule design described earlier, DRS affinity and anti-affinity rules can be programmatically defined in the VMware® Virtual Center application.

FIG. 40 illustrates a rule-programming interface 208, in this example configured for DRS rule programming. FIG. 40 shows anti-affinity rules that have been automatically derived from an analysis map. By using threshold-based generation of rules, both affinity and anti-affinity rules can be established and maintained. A settings box 210 can be used to enable anti-affinity and affinity rules as well as to set thresholds.

Administrators can choose to synchronize the affinity rules directly with a central service, e.g. Virtual Center. For VMware®, the API-level integration can be made bi-directional and all cluster membership information and manually programmed rules can be automatically synchronized with the DRS to enable long-term management of virtual environments. As well, the synchronization operation ensures that there are no rule conflicts prior to applying the new rules. FIG. 41 shows an example user interface 208 that directly integrates with VMware® Virtual Center. From this UI, users can synchronize affinity and anti-affinity rules between the analysis program and the third-party application, in this example through a selectable list of entries 212.

Turning now to FIG. 26 an example process flow is shown that utilizes various capabilities of the analysis program 10, details of which have been described above and are shown in

32

FIGS. 3 to 24. The physical environment analysis 100, as discussed above, obtains data from the existing physical servers 16 and uses virtualization rule sets 11 to evaluate the compatibility of those systems 16 in the physical environment 12 with respect to their candidacy for being virtualized. In this example, the physical environment analysis 102 uses a guest VM compatibility rule set 11a and performs a 1-to-1 compatibility analysis 116 of the systems 16 to determine guest VM candidates 118. It will be appreciated that the 1-to-1 compatibility analysis 116 can be performed according to the principles discussed above, i.e. using the analysis program 10, and as shown in FIGS. 7 to 11 and thus further detail thereof need not be reiterated. This allows the analyst to filter out unsuitable candidates for the optimization stage, which utilizes the more comprehensive multi-dimensional compatibility and consolidation analysis 126. To further filter the candidates 118, another 1-to-1 compatibility analysis 116 can be performed using a guest VM affinity rule set 11b, which enables a more finely filtered set of guest VM candidates or sources 120 to be defined.

The current asset assessment 102 also utilizes data obtained from the existing physical servers 16 and in this example utilizes a virtualization host hardware compatibility rule set 11c to generate a first set of virtualization host candidates 122. The virtualization host system definition process 104 is then performed on the first set of host candidates 122 by performing another 1-to-1 compatibility analysis 116 according to a VM migration compatibility rule set 11d to generate a more refined set of virtualization hosts or targets 124, which would be grouped into clusters. It may be noted that at this stage, if there are insufficient hardware resources for virtualization hosts from the existing physical environment 12, additional servers may be acquired and modeled using hypothetical server models 125 as exemplified in greater detail in FIG. 27 (and discussed above). As can be appreciated from FIG. 26, the hypothetical server models 125 can be introduced not only at the virtualization host system definition 104 stage but also during the multi-dimensional compatibility and consolidation analysis 126 to fine tune the aggregate sizing estimate.

The virtualization environment optimization 106 can then be performed using the set of sources 120 and the set of targets 124. The optimization 106 uses technical and business constraint rule sets 28a and workload types and parameters 28b to determine guest VM candidates and placements 128 as well as VM affinity rules 130 for the virtual environment design 110. It will be appreciated that the multi-dimensional compatibility and consolidation analysis 126 can be performed using the analysis program 10 as discussed above and shown in FIGS. 18 to 24 which includes the application of a transfer auto-fit routine. The multi-dimensional compatibility and consolidation analysis 126 is performed separately for each group of guest VM candidates 120 and cluster of virtualization host candidates 224.

FIGS. 28, 29 and 35 to 38 illustrate example screen shots that can be provided by the virtualization user interface 13 to enable an analyst to perform the transformational P2V process 99. FIG. 28 shows a main or general tab 152 for an analysis editor 150, which provides a mechanism for the analyst to choose settings and generate a set of results that can be used to provide a virtual environment design 110. The description field 154 allows the user to specify a detailed description of the purpose of the analysis. The Dashboard specification 156 allows the user to choose the appropriate dashboard for presenting the analysis results. The Tracking specification 158 allows the user to specify whether multiple versions of the analysis results are to be automatically maintained.

FIG. 29 shows the workload tab **160** in the analysis editor **150**, which is used to select the desired workload types **162** to evaluate in the analysis. In this example, CPU utilization with virtualization overhead, the disk I/O rate in bytes/second, the memory utilization and network I/O in bytes/second are to be evaluated.

FIG. 35 shows the transfer auto-fit tab **184** in the analysis editor **150** which is used, once the initial compatibility analyses have been conducted and the source and target sets **120**, **124** chosen, to apply the auto-fit routine. When performing the transfer auto-fit analysis, users can specify the transfer analysis mode and transfer type **186**. The transfer analysis mode defines the manner in which the multi-dimensional compatibility analysis is performed. The possible modes are affinity, compatibility or both. The affinity mode involves comparing the source systems against the other source systems involved in transfers to a common target. The compatibility mode compares each source systems against their target. The “both” mode applies both the affinity and compatibility comparisons. The transfer type specifies the type of transformation being analyzed—this includes Physical to Virtual (P2V), Virtual to Virtual (V2V), OS Stacking and Application Stacking. The auto-fit algorithm specification **188** allows users to choose between a quick and a comprehensive search for the optimal consolidation solution. The auto-fit limits **190** specify the constraints for the auto-fit solution search. The auto-fit update options **192** allow users to specify whether the auto-fit is performed automatically and whether existing transfers should be removed when the auto-fit is executed.

Upon executing the auto-fit routine, a dashboard summary **194** of the transformational P2V analysis **99** results can be generated and displayed as shown in FIG. 36. A consolidation summary **196** is displayed, which summarizes the number of systems **16** before and after the consolidation and the total number of transfers involved. An aggregate workload summary is also displayed, which shows in this example CPU utilization over the course of a day at minimum/maximum and sustained activity levels both before and after consolidation. The transfers can be displayed in greater detail as shown in FIG. 37 wherein in this example, three target system data sets **200a**, **200b** and **200c** are shown that provide details regarding each target and the transfers involved for virtualization.

A detailed map **202** of the transfers can then be displayed as shown in FIG. 38. This example analysis map **202** shows the P2V transfers based on an auto-fit. In this example, all source systems are placed onto four (4) target systems.

Ongoing Management

After the virtual environment **21** is deployed, the analysis program **10** can be used to collect detailed configuration and workload data from the virtualization hosts and virtual machines (sources) for virtual environment tracking. The data collected from the virtual environment **21** is analyzed to detect outliers and non-compliant guest and virtualization host settings such as the installation of tools on guest systems, service console security settings, etc. The support for live migration between specific virtualization hosts and virtual machines is to be evaluated on an ongoing basis by considering the network and storage dependencies, live migration CPU compatibility, and relevant guest configuration settings (e.g. CPU affinity, connected drives, raw device mappings, internal virtual switch connections). It is typically important that compatibility between servers be maintained to maximize the reliability and optimal operation of the virtualized environment. As the virtual environment changes over time, the analysis program **10** and virtualization UI **13** can be used

to re-analyze the environment based on latest configuration, business and workload data to determine actions to improve compatibility and load balancing. Recommended actions may be to move existing virtual machines and/or virtualization hosts to different clusters, update affinity or anti-affinity rules, or update virtual machine resource reservations. When introducing new virtualization host servers and/or virtual machines to the virtualized environment, an optimization analysis **106** can be performed to determine the recommended assignments based on the compatibility and workload constraints.

Turning now to FIG. 42, a process flow for implementing the ongoing monitoring **15** to achieve the above is shown. After all or part of the physical environment **12** has been transformed, the ongoing analysis involves the management and maintenance of the new virtual environment **21**. Specifically, the analyses can be performed and scheduled to assist in governing, optimizing and planning the placements of virtual machines in the virtual environment. The ongoing management **15** as depicted in FIG. 42 comprises ongoing data acquisition **220**, placement governance **222**, placement optimization **224**, placement planning **226**, and user notifications **285** which is repeated at periodic or predetermined intervals on an ongoing basis. The analysis program can be configured to automatically notify the analyst of key results from the schedule tasks and analyses. Notifications can come in the form of dashboards or be forwarded to the analysts through various mechanisms such as email.

To manage the virtual machines and virtualization hosts, up-to-date data is collected on an ongoing basis, this involves host data collection **228**, guest data collection **230** and virtualization management framework data collection **231**. The majority of the data regarding the virtual machines is collected directly from the virtual machines. Specifically, detailed system configuration information such as operating system configuration, installed applications and workload are collected from the virtual machine. Data regarding the virtualization hosts, current placement of virtual machines and the configuration of the virtual environment such as cluster memberships is collected from the virtualization hosts and/or the virtualization management framework. Examples of virtualization management frameworks include Virtual Center for VMware® V13, System Center Virtual Machine Manager for Microsoft® Hyper-V or XenCenter for Citrix® XenServer. Some performance data such CPU utilization of VMs is collected from the virtualization host or management framework since the CPU utilization measurements from inside the virtual machine can be inaccurate. Virtualization hosts and management frameworks typically provide APIs to collect the required configuration and workload data (e.g. V13 SDK for VMware®, WMI for Microsoft® Hyper-V, Xen API for XenServer, etc.).

The placement governance **222** comprises affinity rule design and updates **232** and VM placement validation **234**. As aspects of the virtual machines change over time, the affinity and anti-affinity rules may need to be updated to reflect the latest conditions. When appropriate, these updated rules should be applied to the virtualization management framework (e.g. VMware® DRS).

The placements of virtual machines often need to be updated over time to reflect changes in the technical, business and workload constraints. The placement validation **234** involves re-analyzing the guest systems based on their current placements on the target hosts using the latest available data. If one or more guests are found to be deployed on inappropriate hosts, the VM layout may be adjusted by migrating

35

VMs as required. Further detail concerning the placement validation **234** is shown in FIG. **43**.

As can be seen in FIG. **43**, the set of source systems **120** and target systems **124** that have been deployed are input to a multi-dimensional compatibility and consolidation analysis **126** as before, utilizing the technical and business constraints rule sets **28a** and the workload types and parameters **28b**. Also input to the analysis **126** is a VM placement validation rule set **240**, which forces guest virtual machines (sources) to remain on their current host (target) by applying a significant penalty if it moves from its current placement. The analysis **126** performs the consolidation auto-fit analysis and generates analysis scores **242** based on the current VM placements. If the analysis results find that all source systems can be placed on their current virtualization hosts, this indicates that the guest VMs continue to meet the technical, business and workload constraints. If the analysis results find that one or more source systems are unplaced, it implies that the constraints are not met with the current placements and that some action is required to ensure operations at the desired levels for performance and risks. Possible actions can include relaxing constraints, moving guest VMs to a different hosts, not running some guest VMs or adding more virtualization hosts to the pool.

Turning back to FIG. **42**, the placement optimization **224** comprises the processes of VM rebalancing **236** and whitespace management **237**. VM rebalancing involves analyzing technical, business and workload constraints of virtual machines and hosts to determine the optimal placement of the virtual machines on an ongoing basis. The frequency of the rebalancing analysis can vary, depending on the volatility of the system workloads and changes in technical and business constraints. There are several variants for the VM rebalancing analysis. One variant places no considerations on the current placements of the virtual machines. This type of analysis searches for the optimal VM placements and assumes virtually no cost in moving the VMs from their current placements. This analysis is applicable for initial VM placements where the environment is being restarted. Another variant considers the current placement of the virtual machines and attempts to eliminate migrations that provide limited benefits. This is accomplished by employing the "VM stickiness" rule set **244** (see FIG. **44**) that penalizes any VM move, ensuring that a move is proposed only if there are significant benefits. FIG. **44** shows further detail of the rebalancing step **236**, which is similar to the placement validation **234** but as noted, uses the VM placement stickiness rule set **244** to determine proposed VM placements **246** and VM affinity rules **248** rather than only analysis scores. It may be noted that by performing the placement validation **234** and rebalancing **236** separately, the validation **234** can be used to indicate whether any of the current VM placements do not meet the analysis constraints and the rebalancing **236** used to indicate where to move VM to enhance load balancing, etc.

Whitespace management tracks the historical and recent server utilization levels against the VM placement constraints to determine if the available host capacity exceeds or falls short of application demands. This analysis can be effectively performed through consolidation analyses on one or more groups of servers in the existing virtual environment. If the analysis results find that the guests do not fit on the existing set of hosts, it indicates that there is a shortfall of capacity. Alternatively, if the analysis results find that there are unused host servers, it indicates a possible excess in capacity.

The placement planning **226** comprises a process of future VM placement validation **238** and planning **239**. Based on historical workload patterns, a model can be defined to predict future workload operation cycles, patterns and trends.

36

For example, when analyzing workload data, analysts can choose to validate current VM placements against these predicted trends to identify potential risks or inefficiencies. The placement planning comprises enabling the generation of future VM placement plans based on predicted operational patterns and trends.

FIGS. **45** to **56** illustrate a series of screen shots provided by the virtualization UI **13** to enable an analyst to perform the placement validation **234** and rebalancing **236** processes. In FIGS. **45** to **51**, like elements with respect to FIGS. **28**, **29** and **35** are given like numerals with a single prime ('). In FIGS. **52** to **56**, like elements with respect to FIGS. **28**, **29** and **35** are given like numerals with a double prime (") and like elements with respect to FIGS. **45** to **51** are given like numerals with a single prime (').

In FIG. **45**, it can be seen that when performing placement validation **234**, a similar analysis editor program **150'** is used wherein the dashboard settings **156'** are set to VM rebalancing since, in this example, the placement validation and **234** and rebalancing **236** utilize the same dashboard.

FIG. **46** shows a systems tab **250** in the analysis editor **150'**, which lists the available systems in a left pane **252** and list what is included in the analysis in a right pane **254**. The right hand pane **254** lists the source and target systems included in the analysis. In this example, the source systems correspond to the guest VMS and the targets are the virtualization hosts.

FIG. **47** shows a rule sets selection tab **256**, which provides a tree mechanism **258** for selecting applicable rule sets. In this example, the static VM placement ruleset is selected to perform the VM placement validation analysis.

FIG. **48** shows the workload tab **160'** when used during the placement validation **234**. In this example, the selected workload types **162** reflect the key resources for analyzing the utilization constraints on the virtualization hosts.

FIG. **49** is a placement validation dashboard page **260** which summarizes the results of the analysis. This page is displayed after the analysis is run and provides an overall status of the analysis and lists various metrics such as the number of source and target systems requiring rebalancing, number of unplaced sources and the number of unused targets **262**. If no actions are required, these metrics should all be zero. In this example, two (2) source systems are found to not fit on their current target host. FIG. **50** shows a page **264** listing the source systems that do not fit on their current host. FIG. **51** shows the analysis results in the form of an analysis map **266**. In the map **266**, the two (2) source systems are shown to be un-transferred and their lower scores of "68" are below the specific auto-fit score limit of "75".

FIGS. **52** to **56** illustrate yet another similar analysis editor **150"** when used for performing the rebalancing, which can be used in manner similar to FIGS. **45** to **51** thus many details thereof need not be reiterated. However, it may be noted that in FIG. **53**, the VM Rebalancing Stickiness ruleset **258** is used in place of the Static VM Placement ruleset. In FIG. **54**, the analysis results **262** indicate that all source systems have been placed, but that one source system was moved to a different target host to meet the auto-fit analysis score constraints. The specific source system that required a transfer is listed in a table **264** in FIG. **56**.

It will be appreciated that although the configuration and workload analyses are performed in this example to contribute to the overall compatibility analyses, each analysis is suitable to be performed on its own and can be conducted separately for finer analyses at any time using the analysis program **10**. The finer analysis may be performed to focus on the remediation of only configuration settings at one time and

spreading workload at another time. As such, each analysis and associated map may be generated on an individual basis without the need to perform the other analyses.

It will also be appreciated that each analysis and associated map discussed above may instead be used for purposes other than consolidation such as capacity planning, regulatory compliance, change, inventory, optimization, administration etc. and any other purpose where compatibility of systems is useful for analyzing systems **16**. It will also be appreciated that the program **10** may also be configured to allow user-entered attributes (e.g. location) that are not available via the auditing process and can factor such attributes into the rules and subsequent analysis.

It will further be appreciated that although the examples provided above are in the context of a distributed system of computer servers, the principles and algorithms discussed are applicable to any system having a plurality of sub-systems where the sub-systems perform similar tasks and thus are capable theoretically of being consolidated and/or virtualized. For example, a local network having a number of personal computers (PCs) could also benefit from a consolidation analysis.

Although the invention has been described with reference to certain specific embodiments, various modifications thereof will be apparent to those skilled in the art as outlined in the claims appended hereto.

The invention claimed is:

1. A method for designing a virtualized environment based on an existing physical environment comprising a plurality of systems, said method comprising:

obtaining a data set for each of said plurality of systems, each data set comprising information pertaining to parameters associated with a corresponding system;
performing a first compatibility analysis on said systems to determine candidate virtual guests;
performing a second compatibility analysis on said systems to determine candidate virtual hosts; and
performing a third compatibility analysis using said candidate virtual hosts, said candidate virtual guests and one or more rule sets pertaining to technical, business and workload constraints to generate a virtual environment design for virtualizing said plurality of systems by evaluating each candidate virtual guest against each candidate virtual host and other candidate virtual guests using said one or more rule sets to determine guest-host placements based on compatibilities of a plurality of virtual design scenarios.

2. A method for managing a virtualized environment, said method comprising:

generating a virtual environment design for a plurality of existing physical systems using technical, business and workload constraints;

facilitating the deployment of said virtualized environment according to said design; and

on an ongoing basis:

obtaining data pertaining to systems being used in said virtualized environment;

validating placement of said systems in said virtualized environment by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to technical, business and workload constraints;

if necessary, rebalancing said systems by determining guest-host placements based on compatibilities of a plurality of virtual design scenarios; and

refining said virtualized environment according to said one of said plurality of virtual design scenarios.

3. A method for performing a virtual to virtual (V2V) transformation for a plurality of existing virtual guests and hosts, said method comprising:

analyzing said existing virtual guests and hosts based on technical, business and workload constraints by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to said technical, business and workload constraints to determine guest-host placements;

based on said analyzing, determining which of said existing virtual servers are most suitable for conversion from one virtualized platform to another virtualized platform; and

providing a mapping from said one platform to said another platform to facilitate said transformation.

4. A method for determining a set of virtualization hosts for a virtualized environment based on an existing physical environment comprising a plurality of systems, said method comprising:

obtaining a data set for each of said plurality of systems, each data set comprising information pertaining to parameters associated with a corresponding system;

performing a first compatibility analysis of said plurality of systems using said data sets and a first rule set pertaining to virtualization specific constraints by evaluating each system against said first rule set to determine an intermediate set of virtualization host candidates that are qualified to be virtual hosts; and

performing a second compatibility analysis of said intermediate set of virtual host candidates using a second rule set pertaining to migration specific constraints by evaluating each intermediate candidate against each other to determine which of said intermediate candidates are compatible with each other and form one or more groups of compatible hosts to be used as said set of virtualization hosts.

5. The method according to claim **4** further comprising incorporating one or more hypothetical hosts into said set of virtualization hosts based on workload requirements for said virtualized environment.

6. A method for evaluating virtualization candidates to determine if additional systems are required to implement a desired virtualized environment, said method comprising:

obtaining a set of virtualization guest candidates and determining aggregate workload requirements based on workload data pertaining to said guest candidates;

obtaining a set of virtualization host candidates and determining aggregate workload capacity based on configuration data pertaining to said host candidates;

comparing said workload requirements against said workload capacity to determine if sufficient capacity exists to satisfy said workload requirements; and

if there is insufficient capacity, adding hypothetical server models to said host candidates to meet said workload requirements.

7. A method for validating an existing virtualized environment comprising a plurality of virtual machines placed on one or more virtual hosts, said method comprising:

obtaining a data set for each of said plurality of virtual machines, each data set comprising information pertaining to technical, business and workload constraints associated with a corresponding virtual machine;

evaluating the placement of said virtual machines in said virtualized environment using said data sets by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to

US 8,209,687 B2

39

said technical, business and workload constraints to determine guest-host placements; and identifying the existence of virtual machines with suboptimal placements to enable alternative placements for said virtual machines.

8. The method according to claim 1, wherein said second compatibility analysis comprises:

performing a third compatibility analysis of said plurality of systems using said data sets and a first rule set pertaining to virtualization specific constraints by evaluating each system against said first rule set to determine an intermediate set of virtualization host candidates that are qualified to be virtual hosts; and

performing a fourth compatibility analysis of said intermediate set of virtual host candidates using a second rule set pertaining to migration specific constraints by evaluating each intermediate candidate against each other to determine which of said intermediate candidates are compatible with each other and form one or more groups of compatible hosts to be used as said set of virtualization hosts.

9. The method according to claim 8 further comprising incorporating one or more hypothetical hosts into said set of virtualization hosts based on workload requirements for said virtualized environment.

10. The method according to claim 1, further comprising: facilitating the deployment of a virtualized environment according to said virtual environment design; and on an ongoing basis:

obtaining data pertaining to systems being used in said virtualized environment;

validating placement of said systems in said virtualized environment by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to technical, business and workload constraints, to determine guest-host placements based on compatibilities of a plurality of virtual design scenarios;

if necessary rebalancing said systems according to one of said plurality of virtual design scenarios; and refining said virtualized environment according to said one of said plurality of virtual design scenarios.

11. The method according to claim 1, further comprising: comparing workload requirements of said virtual guests against said workload capacity of said virtual hosts, to determine if sufficient capacity exists to satisfy said workload requirements; and

if there is insufficient capacity, adding hypothetical server models to virtual host candidates to meet said workload requirements.

12. The method according to claim 1, further comprising: evaluating the placement of virtual machines in a virtualized environment based on said virtual environment design by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to said technical, business and workload constraints to determine guest-host placements; and

identifying the existence of virtual machines with suboptimal placements to enable alternative placements for said virtual machines.

13. A computer readable medium comprising computer executable instructions for designing a virtualized environment based on an existing physical environment comprising a plurality of systems, said computer executable instructions comprising instructions for:

40

obtaining a data set for each of said plurality of systems, each data set comprising information pertaining to parameters associated with a corresponding system; performing a first compatibility analysis on said systems to determine candidate virtual guests;

performing a second compatibility analysis on said systems to determine candidate virtual hosts; and

performing a third compatibility analysis using said candidate virtual hosts, said candidate virtual guests and one or more rule sets pertaining to technical, business and workload constraints to generate a virtual environment design for virtualizing said plurality of systems by evaluating each candidate virtual guest against each candidate virtual host and other candidate virtual guests using said one or more rule sets to determine guest-host placements based on compatibilities of a plurality of virtual design scenarios.

14. The computer readable medium according to claim 13, wherein said second compatibility analysis comprises instructions for:

performing a third compatibility analysis of said plurality of systems using said data sets and a first rule set pertaining to virtualization specific constraints by evaluating each system against said first rule set to determine an intermediate set of virtualization host candidates that are qualified to be virtual hosts; and

performing a fourth compatibility analysis of said intermediate set of virtual host candidates using a second rule set pertaining to migration specific constraints by evaluating each intermediate candidate against each other to determine which of said intermediate candidates are compatible with each other and form one or more groups of compatible hosts to be used as said set of virtualization hosts.

15. The computer readable medium according to claim 14 further comprising instructions for incorporating one or more hypothetical hosts into said set of virtualization hosts based on workload requirements for said virtualized environment.

16. The computer readable medium according to claim 13, further comprising instructions for:

facilitating the deployment of a virtualized environment according to said virtual environment design; and on an ongoing basis:

obtaining data pertaining to systems being used in said virtualized environment;

validating placement of said systems in said virtualized environment by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to technical, business and workload constraints, to determine guest-host placements based on compatibilities of a plurality of virtual design scenarios;

if necessary rebalancing said systems according to one of said plurality of virtual design scenarios; and refining said virtualized environment according to said one of said plurality of virtual design scenarios.

17. The computer readable medium according to claim 13, further comprising instructions for:

comparing workload requirements of said virtual guests against said workload capacity of said virtual hosts, to determine if sufficient capacity exists to satisfy said workload requirements; and

if there is insufficient capacity, adding hypothetical server models to virtual host candidates to meet said workload requirements.

US 8,209,687 B2

41

18. The computer readable medium according to claim 13, further comprising instructions for:
evaluating the placement of virtual machines in a virtualized environment based on said virtual environment design by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to said technical, business and work-

42

load constraints to determine guest-host placements; and
identifying the existence of virtual machines with suboptimal placements to enable alternative placements for said virtual machines.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 8,209,687 B2
APPLICATION NO. : 12/201323
DATED : June 26, 2012
INVENTOR(S) : Tom Silangan Yuyitung et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In the Claims

Column 38, lines 2 and 4 delete “virtual quests” and replace with “virtual guests”.

Signed and Sealed this
Twenty-fourth Day of March, 2015



Michelle K. Lee
Director of the United States Patent and Trademark Office

Exhibit 13

VMWARE TECHNICAL JOURNAL

Editors: Steve Muir, Rita Tavilla and Ben Verghese

TABLE OF CONTENTS

- 1 Introduction**
Steve Herrod, CTO
- 2 VisorFS: A Special-purpose File System for Efficient Handling of System Images**
Olivier Cremel
- 3 A Software-based Approach to Testing VMware® vSphere® VMkernel Public APIs**
Lan Xue, Sreevathsa Sathyanarayana, Thorbjørn Donbaek, Ramesh Pallapotu, James Truong, Sriram Sankaran, Eric Lorimer
- 4 Providing Efficient and Seamless Desktop Services in Ubiquitous Computing Environments**
Lizhu Zhang, Wenlong Shao, Jim Grandy
- 5 Comprehensive User Experience Monitoring**
Lawrence Spracklen, Banit Agrawal, Rishi Bidarkar, Hari Sivaraman
- 6 StatsFeeder: An Extensible Statistics Collection Framework for Virtualized Environments**
Vijayaraghavan Soundararajan, Balaji Parimi, Jon Cook
- 7 VMware Distributed Resource Management: Design, Implementation, and Lessons Learned**
Ajay Gulati, Anne Holler, Minwen Ji, Ganesha Shanmuganathan, Carl Waldspurger, Xiaoyun Zhu
- 8 Identity, Access Control, and VMware Horizon**
Will Pugh, Kyle Austin
- 9 VMworld 2011 Hands-On Labs: Implementation and Workflow**
Adam Zimman, Clair Roberts, Mornay Van Der Walt

VMware Distributed Resource Management: Design, Implementation, and Lessons Learned

Ajay Gulati

VMware, Inc.
agulati@vmware.com

Ganesha Shanmuganathan

VMware, Inc.
sganesh@vmware.com

Anne Holler

VMware, Inc.
anne@vmware.com

Carl Waldspurger

carl@waldspurger.org

Minwen Ji

Facebook, Inc.
mji@fb.com

Xiaoyun Zhu

VMware, Inc.
xzhu@vmware.com

Abstract

Automated management of physical resources is critical for reducing the operational costs of virtualized environments. An effective resource-management solution must provide performance isolation among virtual machines (VMs), handle resource fragmentation across physical hosts and optimize scheduling for multiple resources. It must also utilize the underlying hardware infrastructure efficiently. In this paper, we present the design and implementation of two such management solutions: DRS and DPM. We also highlight some key lessons learned from production customer deployments over a period of more than five years.

VMware's Distributed Resource Scheduler (DRS) manages the allocation of physical resources to a set of virtual machines deployed in a cluster of hosts, each running the VMware ESX hypervisor. DRS maps VMs to hosts and performs intelligent load balancing in order to improve performance and to enforce both user-specified policies and system-level constraints. Using a variety of experiments, augmented with simulation results, we show that DRS significantly improves the overall performance of VMs running in a cluster. DRS also supports a "what-if" mode, making it possible to evaluate the impact of changes in workloads or cluster configuration.

VMware's Distributed Power Management (DPM) extends DRS with the ability to reduce power consumption by consolidating VMs onto fewer hosts. DPM recommends evacuating and powering off hosts when CPU and memory resources are lightly utilized. It recommends powering on hosts appropriately as demand increases, or as required to satisfy resource-management policies and constraints. Our extensive evaluation shows that in clusters with non-trivial periods of lowered demand, DPM reduces server power consumption significantly.

Categories and Subject Descriptors

- C.4 [Performance of Systems]: Modeling techniques;
- C.4 [Performance of Systems]: Measurement techniques;
- C.4 [Performance of Systems]: Performance attributes;
- D.4.8 [Operating Systems]: Performance—*Modeling and prediction*;
- D.4.8 [Operating Systems]: Performance—*Measurements*;
- D.4.8 [Operating Systems]: Performance—*Operational analysis*

General Terms

Algorithms, Design, Experimentation, Management, Measurement, Performance

Keywords

VM, Virtualization, Resource Management, Scheduling, Cluster, Hosts, Load Balancing, Power Management

1 Introduction

Initially, the rapid adoption of virtualization was fueled by significant cost savings resulting from server consolidation. Running several virtual machines (VMs) on a single physical host improved hardware utilization, allowing administrators to "do more with less" and reduce capital expenses. Later, more advanced VM capabilities such as cloning, template-based deployment, checkpointing, and live migration [43] of running VMs led to more agile IT infrastructures. As a result, it became much easier to create and manage virtual machines.

The ease of deploying workloads in VMs is leading to increasingly large VM installations. Moreover, hardware technology trends continue to produce more powerful servers with higher core counts and increased memory density, causing consolidation ratios to rise. However, the operational expense of managing VMs now represents a significant fraction of overall costs for datacenters using virtualization. Ideally, the complexity of managing a virtualized environment should also benefit from consolidation, scaling with the number of hosts, rather than the number of VMs. Otherwise, managing a virtual infrastructure would be as hard — or arguably harder, due to sharing and contention — as managing a physical environment, where each application runs on its own dedicated hardware.

In practice, we observed that a large fraction of the operational costs in a virtualized environment were related to the inherent complexity of determining good VM-to-host mappings, and deciding when to use vMotion [8], VMware's live migration technology, to rebalance load by changing those mappings. The difficulty of this problem is exacerbated by the fragmentation of resources across many physical hosts and the need to balance the utilization of multiple resources (including CPU and memory) simultaneously.

DPM only: In this case, the HPM policy was set to “High Performance”, effectively disabling HPM. Since DPM had been enabled, six of the eight hosts remained in the powered-off state during the initial 30 minutes of the idle period. This led to a total cluster power consumption of about 566W, a 60% reduction as compared with the 1500W in the HPM-only case. Notice the power saving is not exactly 75% here because in the DPM case, the two hosts that remained powered on had a much higher CPU utilization, resulting in higher per-host power consumption. After the VM load increase, DPM first powered on four of the standby hosts in the 33rd minute, and then powered on two additional hosts in the 38th minute, bringing the cluster back to full capacity. The total cluster power consumption increased to approximately 2100W after all the hosts were powered on. When the VMs became idle again, DPM (by design) kept all the hosts powered on for more than 30 minutes, resulting in a total cluster power consumption of approximately 1900W.

DPM+HPM combined: This case is similar to the DPM-only case, except for the following three observations. First, after the VM load increase, DPM initially powered on two of the standby hosts in the 33rd minute, and then powered on the four remaining hosts in the 38th minute. Second, the total cluster power after all the hosts were powered on was roughly 2000W, 100W lower compared to the DPM-only case. This was because the newly powered-on hosts had fewer VMs running on them, resulting in lower host utilization and providing HPM with an opportunity to reduce the power consumption on these hosts. Third, in the last 30 minutes when the cluster was idle, the DPM+HPM combined policy provided an additional 400W of power reduction (1500W vs. 1900W) compared to the DPM-only case. Overall, this combined policy provides the maximum power savings among the three power management policies we tested.




Figure 15: Per-host power consumption for the 8 hosts in the cluster.

Figure 15 shows the per-host power consumption as a function of time for the DPM+HPM combined case. We can clearly see how each host’s power consumption varied as the host power state or the VM load level changed. In the last 30 minutes, periodic spikes are visible in the host power consumption, due to vMotion-induced higher CPU utilization on these hosts. These vMotions were recommended by DRS every five minutes to balance the load in the cluster.

6. DRS and DPM in the Field

When DRS was introduced in early 2006, vMotion was just beginning to gain widespread adoption, but customers were wary of automated migration of virtual machines. In the early days, one request from customers was support for manual-move recommendations. A human administrator would inspect the recommendations and apply the moves only if they made sense. Some administrators would run DRS in manual mode and if the same move was recommended over a substantial number of DRS invocations, then the administrator would apply the move. The use of DRS manual mode diminished over time as the DRS algorithm became more mature and as administrators became more comfortable with automated VM migration; as of vSphere 5.0, the use of DRS manual mode is very low.

The first version of DRS did not have cost-benefit analysis turned on, as the code was considered experimental. This led to the problem that DRS could make recommendations that moved VMs back and forth in response to frequently changing demand. In the very next release, cost-benefit feature was enabled by default, leading to higher-quality moves and fewer migrations.

The DRS algorithm tries to get the biggest bang for its vMotion buck, *i.e.*, to minimize the total number of moves needed for load-balancing. Moving the largest, most-active VMs can have the highest impact on correcting imbalance, and hence DRS would favor such moves. While choosing such VMs for vMotion in order to issue fewer moves seemed good in theory, some customers did not like this selection since their largest, most active VMs were also their most important and performance-sensitive VMs, and vMotioning those VMs could adversely impact their performance during the migration.

To address this issue, DRS cost-benefit analysis was changed to take into account the impact of vMotion on the workload which it had not done previously. As vSphere’s vMotion continued to be improved, the cost modeling of that impact required updating as well. Over time we learned that the modeling aspects of the algorithm should be separated from the parts of the algorithm that use the model, to ease the maintenance of the algorithm code as the technology changes. For example, we moved to having the algorithm consider the vMotion time, with the details of the parameters relevant to that generation of vMotion technology handled in modeling-specific code.

Earlier versions of DRS did not support affinity between VMs and hosts and it was thought that affinities between VMs should be sufficient. We also wanted administrators to think less about individual hosts and more about aggregate clusters. While VM-to-VM affinity was sufficient for most technical use-cases, there were other requirements such as software licensing that made administrators want to isolate VMs onto a set of hosts. Administrators started rolling out their own solutions to pinning VMs to a set of hosts, such as adding dummy networks to the VMs and adding the networks only to a subset of hosts, making the other hosts incompatible.

Exhibit 14

SELF-DRIVING OPERATIONS BY VMWARE vREALIZE OPERATIONS

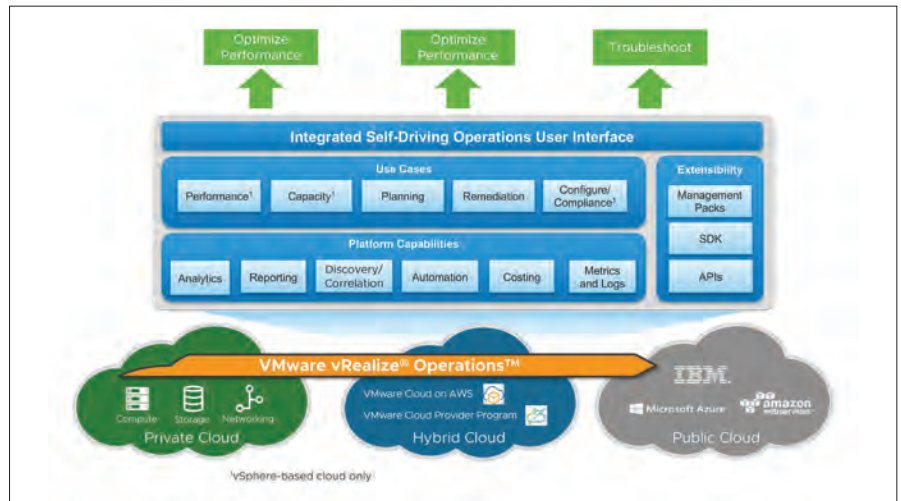
AT A GLANCE

Self-driving operations by VMware vRealize® Operations™ delivers continuous performance optimization at minimal cost driven by business and operational intent, efficient capacity management and planning and intelligent remediation. It automates and simplifies IT operations management and provides unified visibility from applications to infrastructure across physical, virtual, and cloud environments.

KEY BENEFITS

- Centralized management for software-defined data center (SDDC) and multi-cloud environments
- Intent-driven automated workload balancing and predictive DRS to resolve resource contention and enable continuous performance optimization
- Real-time, predictive capacity analytics to proactively alert on capacity risk, forecast future demand, and deliver actionable recommendations
- Integration of costs with capacity analytics to optimize utilization and reduce costs
- Cost transparency across private and public clouds to optimize planning
- Integration with vRealize Log Insight™¹ for 360-degree troubleshooting with metrics and logs in context
- Complete vSAN operations management that starts in VMware vCenter® and delivers full stack troubleshooting and capacity management
- Fully open and extensible platform

¹ Sold separately as standalone and included in vCloud Suite and vRealize Suite.



Self-Driving Operations for Physical, Virtual, and Cloud Infrastructures

Self-Driving Operations Overview

VMware delivers self-driving operations from applications to infrastructure to optimize, plan and scale SDDC and multi-cloud deployments. This highly scalable, extensible, and intuitive operations platform automates and centralizes management for SDDC and cloud, delivering continuous performance optimization based on intent, efficient capacity management, proactive planning, and intelligent remediation.

Use Cases

Continuous performance optimization – Assure performance at minimal cost, driven by operational and business intent with predictive analytics driving actions to automatically balance workloads and proactively avoid contention.

Automate workload balancing to reduce software license costs, optimize based on performance tiers, densify clusters, or enforce compliance.

Efficient capacity management and planning – Reduce cost and capacity risk with real-time, predictive capacity analytics delivering optimal densification and proactive planning. Predict future demand, get actionable recommendations, and automate reclamation and right-sizing. Integrate costs and capacity analytics to optimize utilization and reduce costs. Advanced what-if scenarios help plan capacity and model best-fit for new workloads across private cloud and multiple public clouds.

Intelligent remediation – Predict, prevent, and troubleshoot faster with actionable insights correlating metrics and logs and unified visibility from applications to infrastructure. Centralize IT operations management with native SDDC integrations, federated views, and a highly scalable and extensible platform. Manage SDDC technologies, like VMware vSAN™, at scale with operational visibility that starts in vCenter and delivers full environment overview, troubleshooting, and capacity management.

Exhibit 15

Workload Optimization

<p>msbu-east</p> <p>0 Days Remaining</p> <p>US\$8 Cost Savings</p> <p>Not Optimized</p>	<p>lab-dc</p> <p>> 1 Year Remaining</p> <p>US\$8 Cost Savings</p> <p>Not Optimized</p>	<p>msbu-demo</p> <p>> 1 Year Remaining</p> <p>US\$47 Cost Savings</p> <p>Not Optimized</p>	<p>msbu-mgmt</p> <p>> 1 Year Remaining</p> <p>US\$40 Cost Savings</p> <p>Not Optimized</p>	<p>msbu-west</p> <p>> 1 Year Remaining</p> <p>US\$1 Cost Savings</p> <p>Optimized</p>
---	---	---	---	--

ALL DATACENTERS

msbu-east

msbu-vc-east

Optimization Recommendation

Status: **Not Optimized**

You can optimize your datacenter by moving workloads to avoid performance issues.


OPTIMIZE NOW SCHEDULE

Placement Settings 0 Tag Applied

Utilization Objective: Balance

- ✓ Avoid Performance Issues
- ✓ Balance Workloads

EDIT SETTINGS



Are your clusters meeting your utilization objective?

[VIEW DRS SUMMARY](#) [SET DRS AUTOMATION](#)

Exhibit 16

(12) **United States Patent**
Hillier

(10) **Patent No.:** **US 9,654,367 B2**
 (45) **Date of Patent:** **May 16, 2017**

(54) **SYSTEM AND METHOD FOR DETERMINING AND VISUALIZING EFFICIENCIES AND RISKS IN COMPUTING ENVIRONMENTS**

(71) Applicant: **Cirba Inc.**, Richmond Hill (CA)

(72) Inventor: **Andrew Derek Hillier**, Toronto (CA)

(73) Assignee: **Cirba IP Inc.**, Richmond Hill (CA)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 100 days.

(21) Appl. No.: **14/180,438**

(22) Filed: **Feb. 14, 2014**

(65) **Prior Publication Data**

US 2014/0164612 A1 Jun. 12, 2014

Related U.S. Application Data

(63) Continuation of application No. PCT/CA2012/050561, filed on Aug. 16, 2012.
 (Continued)

(51) **Int. Cl.**
H04L 12/26 (2006.01)
H04L 12/24 (2006.01)
 (Continued)

(52) **U.S. Cl.**
 CPC **H04L 43/0876** (2013.01); **G06F 11/328** (2013.01); **G06F 11/3442** (2013.01);
 (Continued)

(58) **Field of Classification Search**
 CPC G06F 11/3433; G06F 11/3447; G06F 9/5011; G06F 5/161; H04L 43/0876; H04L 41/147; H04L 47/20
 (Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,148,335 A 11/2000 Haggard et al.
 8,087,025 B1 * 12/2011 Graupner G06F 9/505 718/104

(Continued)

FOREIGN PATENT DOCUMENTS

CA 2583582 A1 5/2006
 CA 2649714 A1 11/2007

(Continued)

OTHER PUBLICATIONS

Goiri, I. et al. "Energy-aware Scheduling in Virtualized Datacenters"; 2010 IEEE International Conference on Cluster Computing; Sep. 20, 2010; pp. 58 to 67; ISBN: 978-1-4244-8373-0.

(Continued)

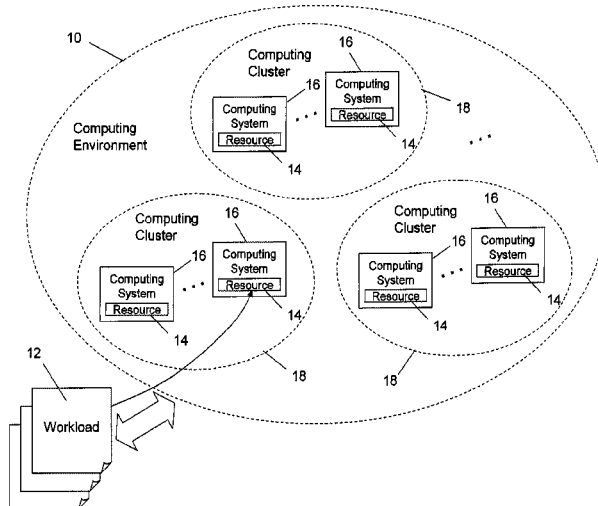
Primary Examiner — Tesfay Yohannes

(74) *Attorney, Agent, or Firm* — Blake, Cassels & Graydon LLP; Brett J. Slaney

(57) **ABSTRACT**

A system and method are provided for determining and visualizing efficiencies and risks in computing environments. The method comprises obtaining resource utilization data and resource capacity data for a plurality of entities in a computing environment; obtaining at least one operational policy defining at appropriate level of at least one resource used by the computing environment according to at least one factor; and computing at least one score quantifying efficiencies and risks associated with the computing environment based on the resource utilization data, resource capacity data, and at least one operational policy.

19 Claims, 12 Drawing Sheets



US 9,654,367 B2

Page 2

Related U.S. Application Data

		2012/0159476 A1*	6/2012	Ramteke	G06F 9/5033 718/1
(60)	Provisional application No. 61/523,912, filed on Aug. 16, 2011.	2012/0284408 A1*	11/2012	Dutta	G06F 9/5066 709/226
		2013/0041768 A1*	2/2013	Llach	G06Q 20/28 705/17

(51) **Int. Cl.**

G06F 11/34 (2006.01)
G06F 11/32 (2006.01)

(52) **U.S. Cl.**

CPC **H04L 41/0896** (2013.01); **G06F 11/3409** (2013.01); **G06F 2201/815** (2013.01)

(58) **Field of Classification Search**

USPC 709/204, 226; 370/223
See application file for complete search history.

(56)

References Cited

U.S. PATENT DOCUMENTS

2003/0046396 A1*	3/2003	Richter	G06F 9/505 709/226
2007/0165019 A1*	7/2007	Hale	G06Q 10/00 345/418
2007/0250829 A1	10/2007	Hillier et al.	
2008/0028409 A1*	1/2008	Cherkasova	G06F 9/5061 718/104
2009/0037367 A1*	2/2009	Wein	G06F 9/5061
2009/0292715 A1	11/2009	Chiaromonte et al.	
2010/0082505 A1	4/2010	Hollingsworth et al.	
2012/0011077 A1*	1/2012	Bhagat	G06F 21/554 705/317
2012/0140620 A1*	6/2012	Hogan	H04L 41/0893 370/230

FOREIGN PATENT DOCUMENTS

WO	2007/147258 A1	12/2007
WO	2011/011408 A1	1/2011

OTHER PUBLICATIONS

Lanches, P.; Supplementary European search report from corresponding European Application No. 12824582.6; search completed May 6, 2015.

"How to Measure & Manage Efficiency in Virtual Environments Using CiRBA's Efficiency & Risk Dashboard"; Online Mar. 23, 2011; CiRBA.

Hillier, Andrew; "Server Capacity Defrag: Maximizing Infrastructure Efficiency through Strategic Workload Placements (Defrag Isn't Just for Storage Anymore!"; Jan. 24, 2011; CiRBA Inc.; <<http://www.cirba.com/forms/Server-Capacity-Defrag.htm>>.

Hillier, Andrew; "Transformational Analytics: Virtualizing IT Environments"; Apr. 2008; CiRBA Inc.; <<http://www.cirba.com/forms/Transformational-Analytics-Virtualizing-IT-Environments.htm>>.

Hillier, Andrew; "Advanced Workload Analysis Techniques"; Apr. 2008; CiRBA Inc.; <<http://www.cirba.com/forms/Advanced-Workload-Analysis-Techniques.htm>>.

"Thanks for attending yesterday's live demo of the Efficiency & Risk dashboard"; Twitter Status Update from CiRBA @CiRBA; Mar. 24, 2011; <<https://twitter.com/CiRBA/status/51018299923894272>>.

Matheson, L.; Search Report from corresponding PCT/CA2012/050561; search completed Nov. 15, 2012.

* cited by examiner




FIG. 1




FIG. 2

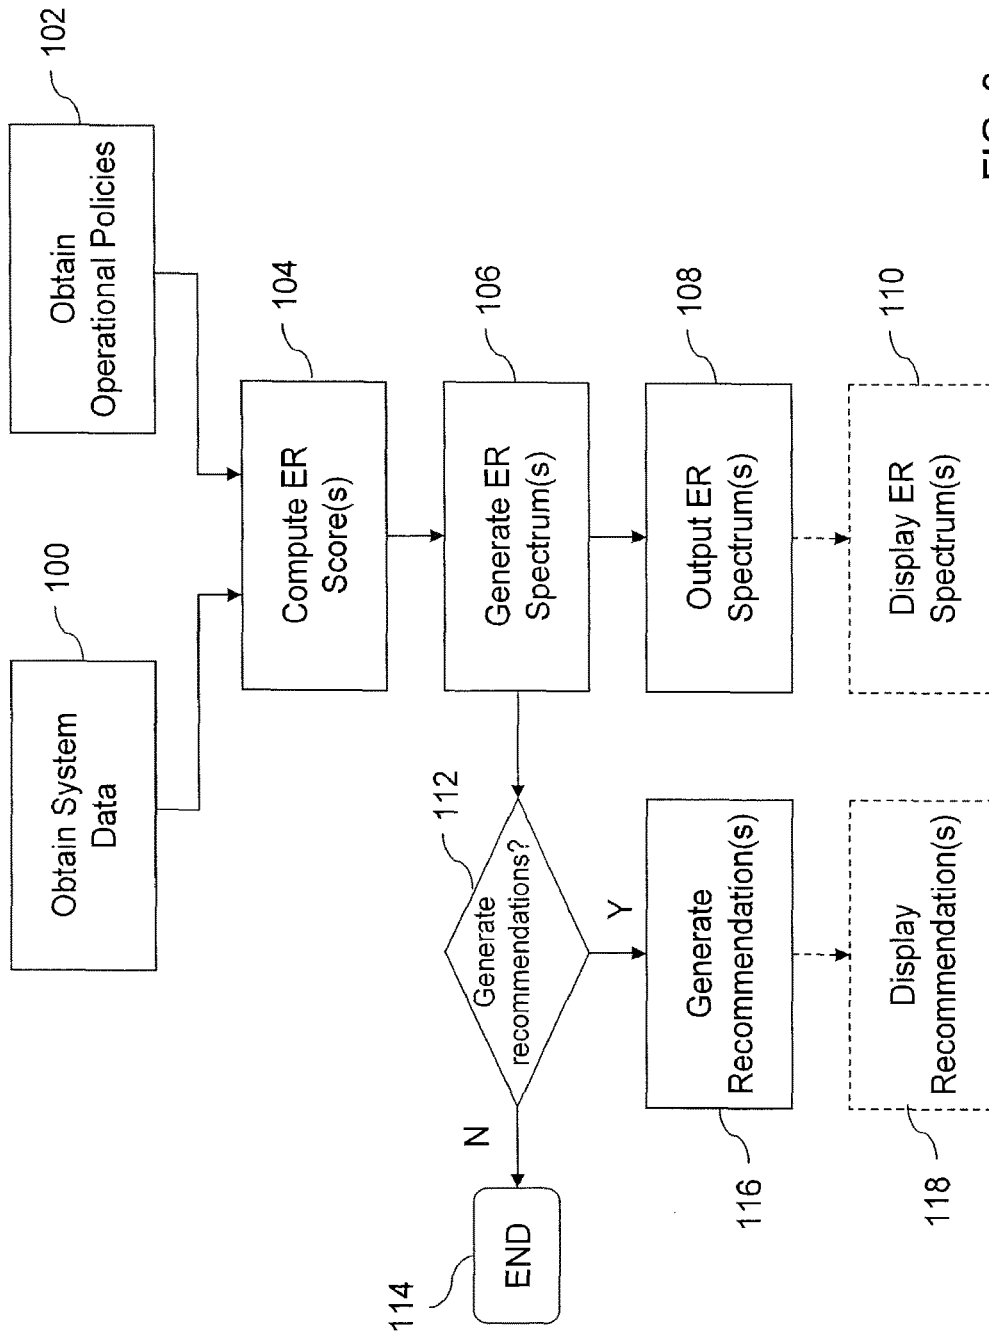


FIG. 3

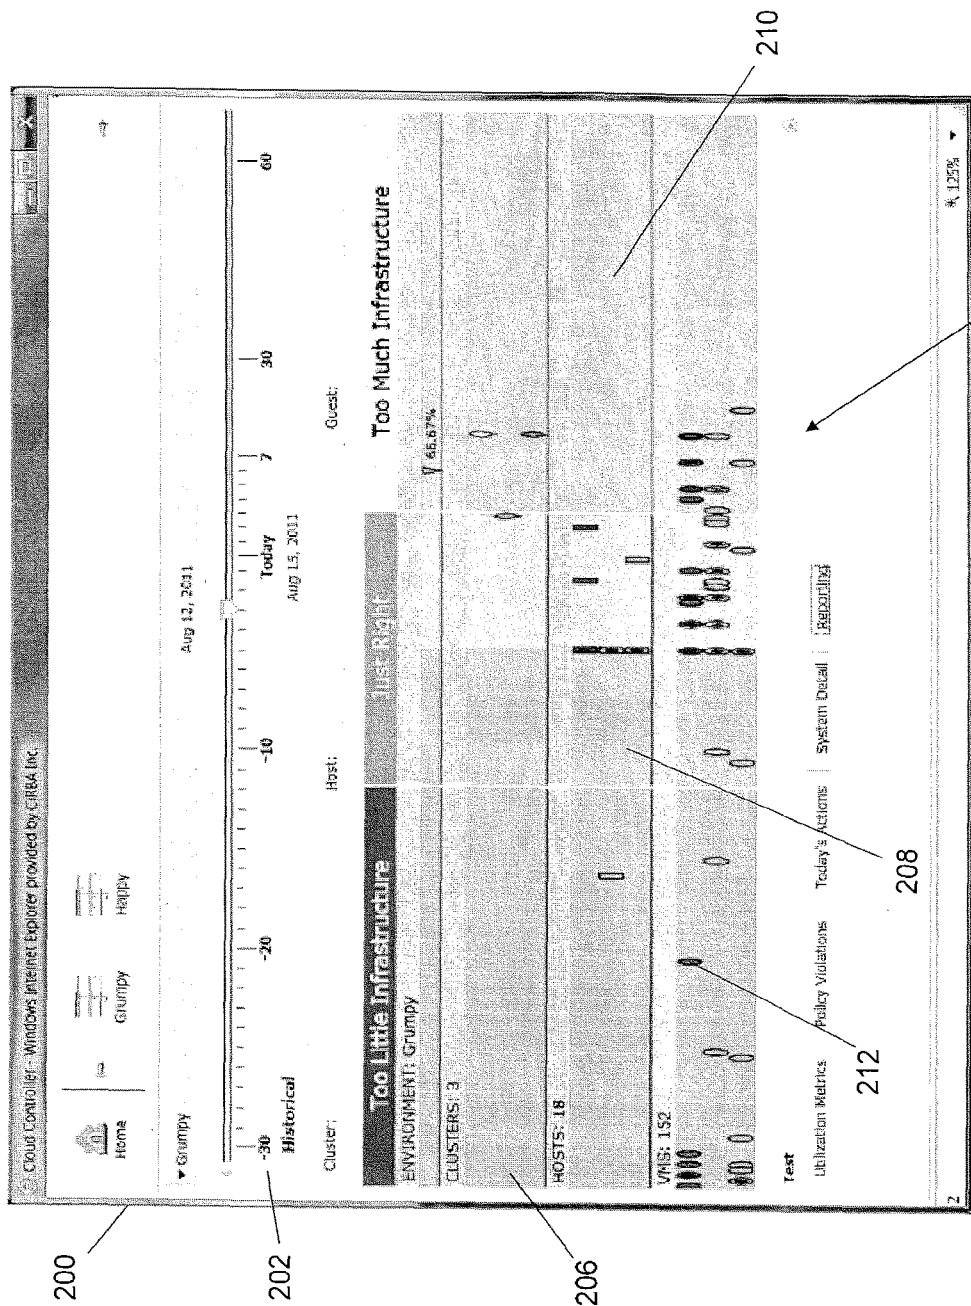


FIG. 4

200'

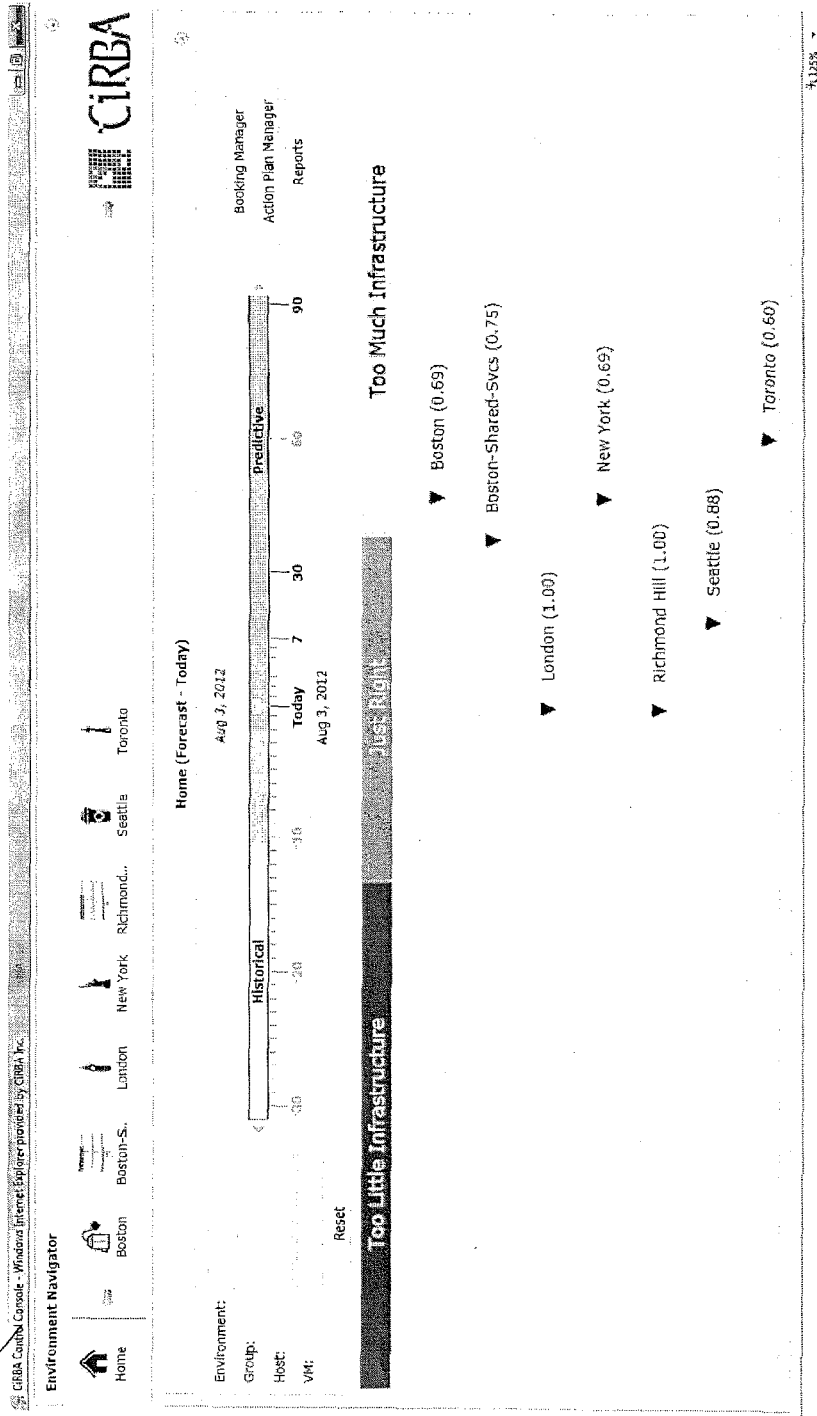


FIG. 5

200"

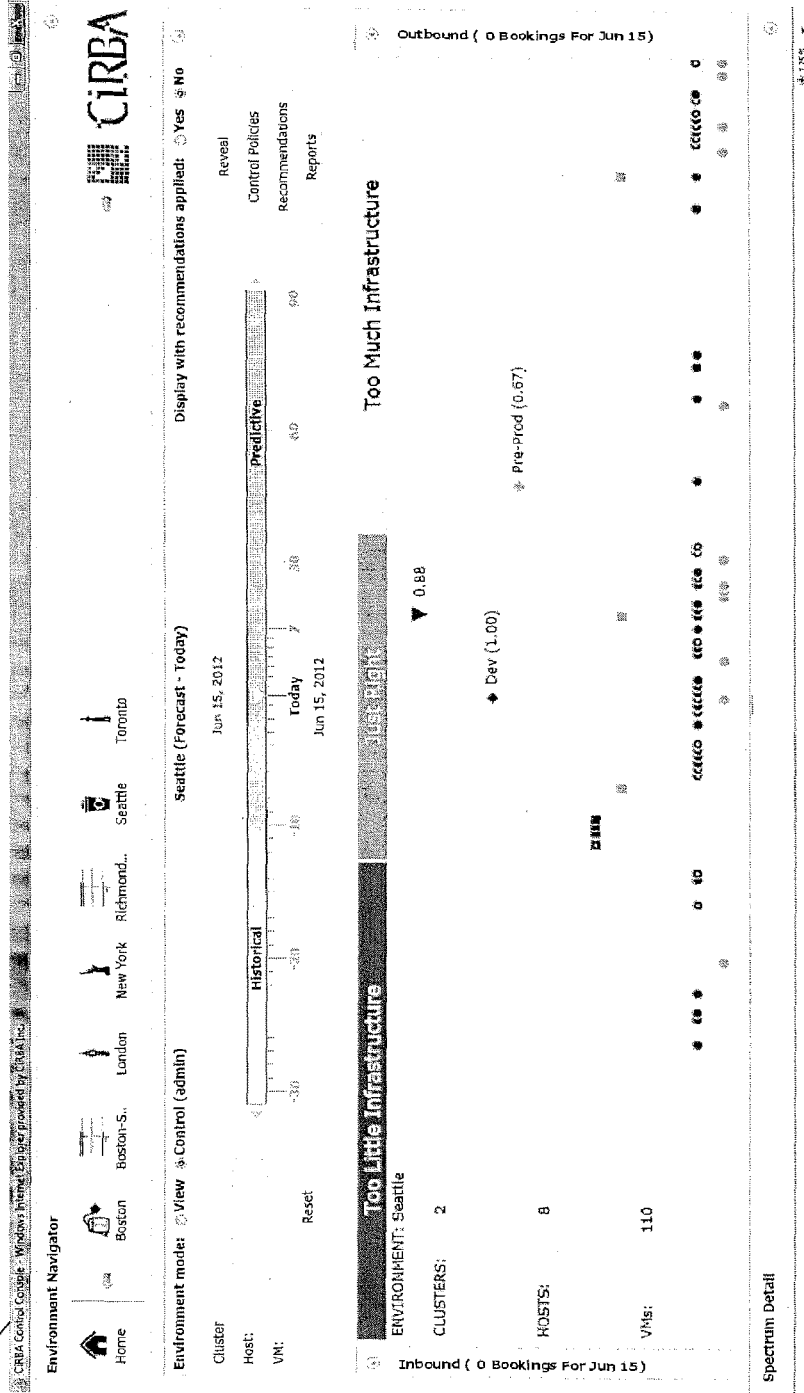


FIG. 6

204"




FIG. 7

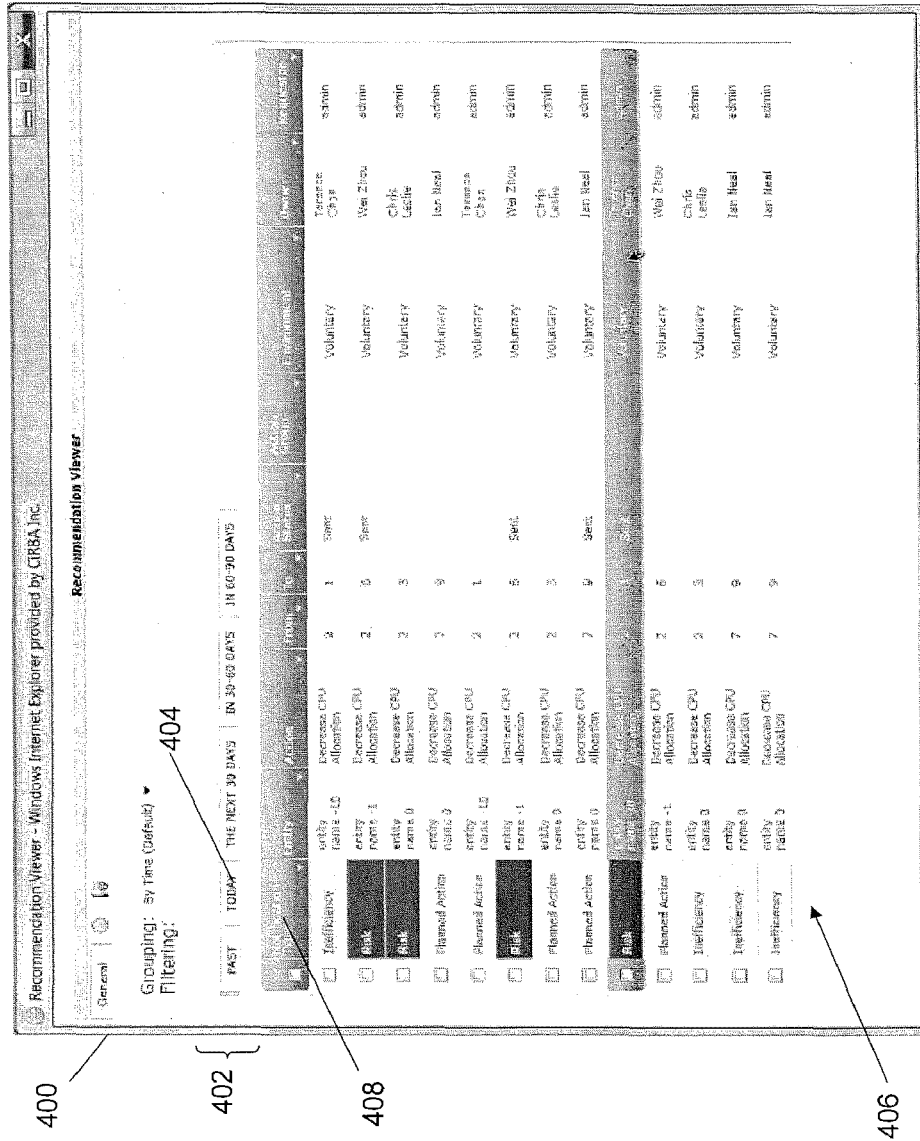



FIG. 8

400'



406'

FIG. 9

200''

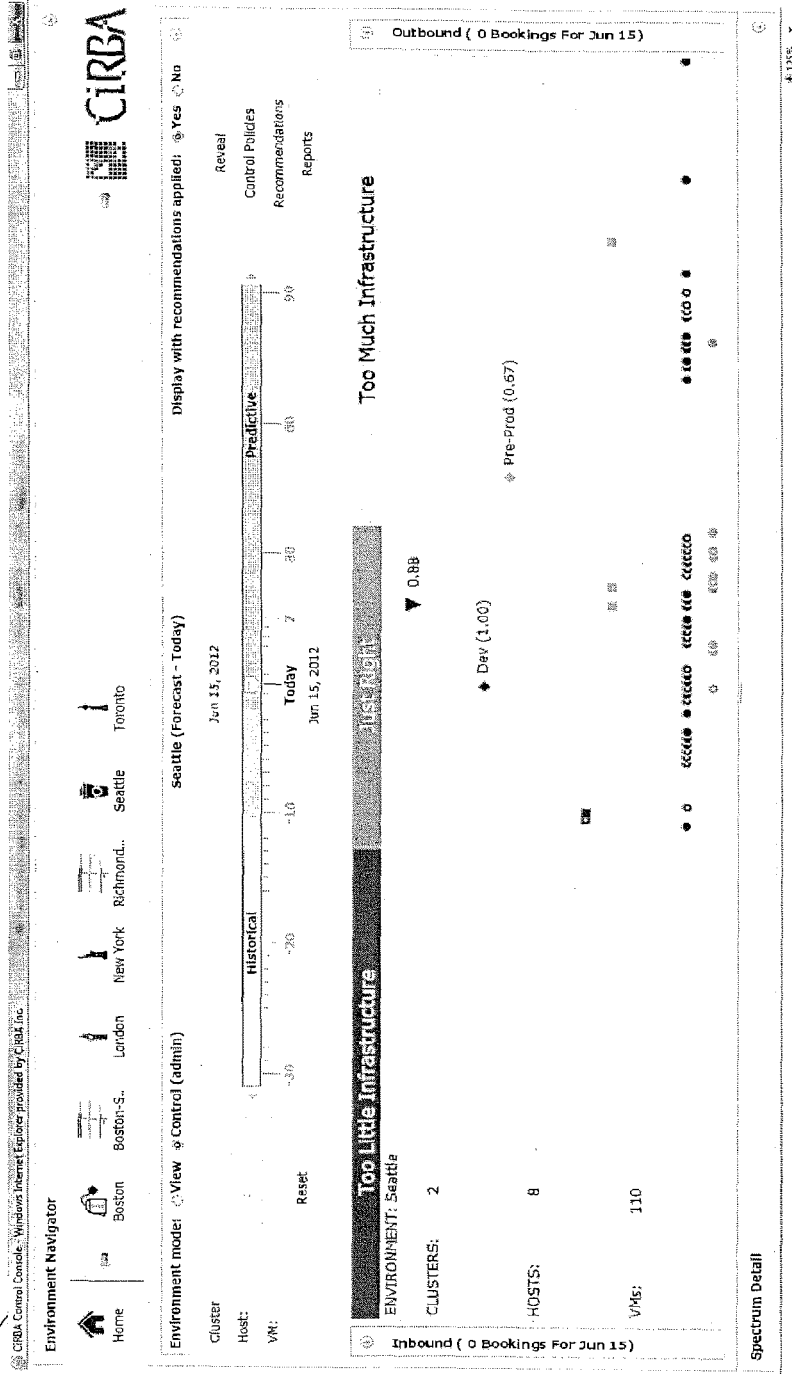
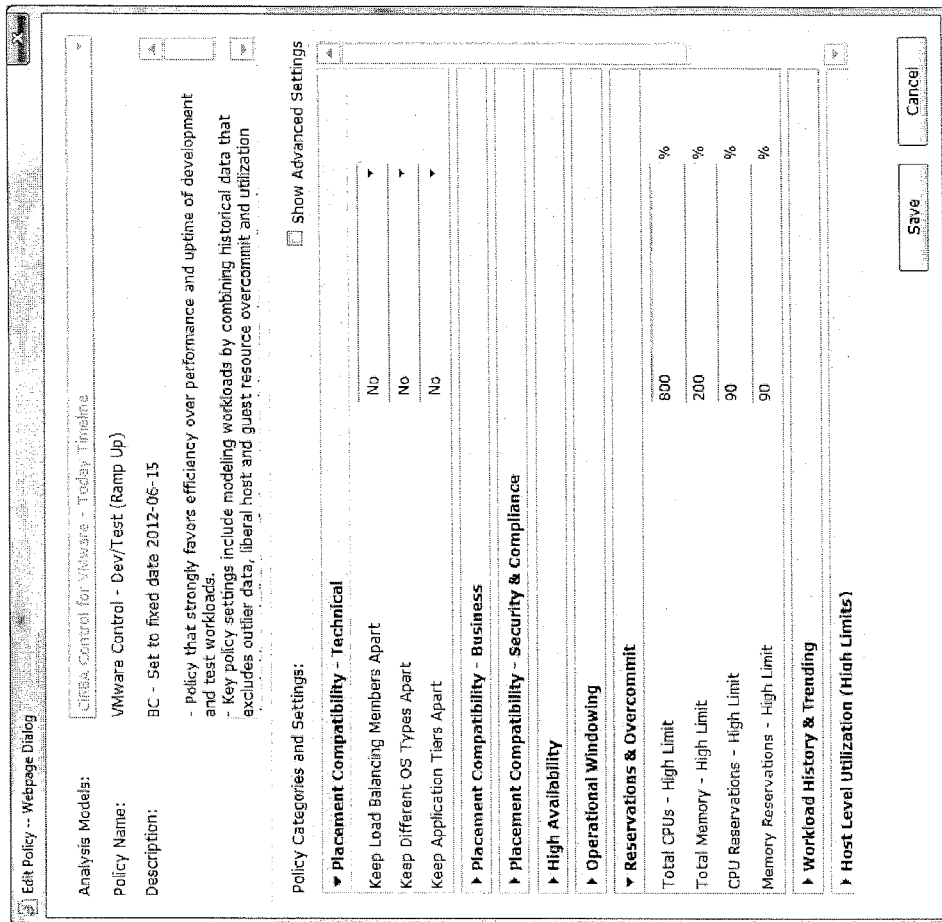


FIG. 10



500

FIG. 11

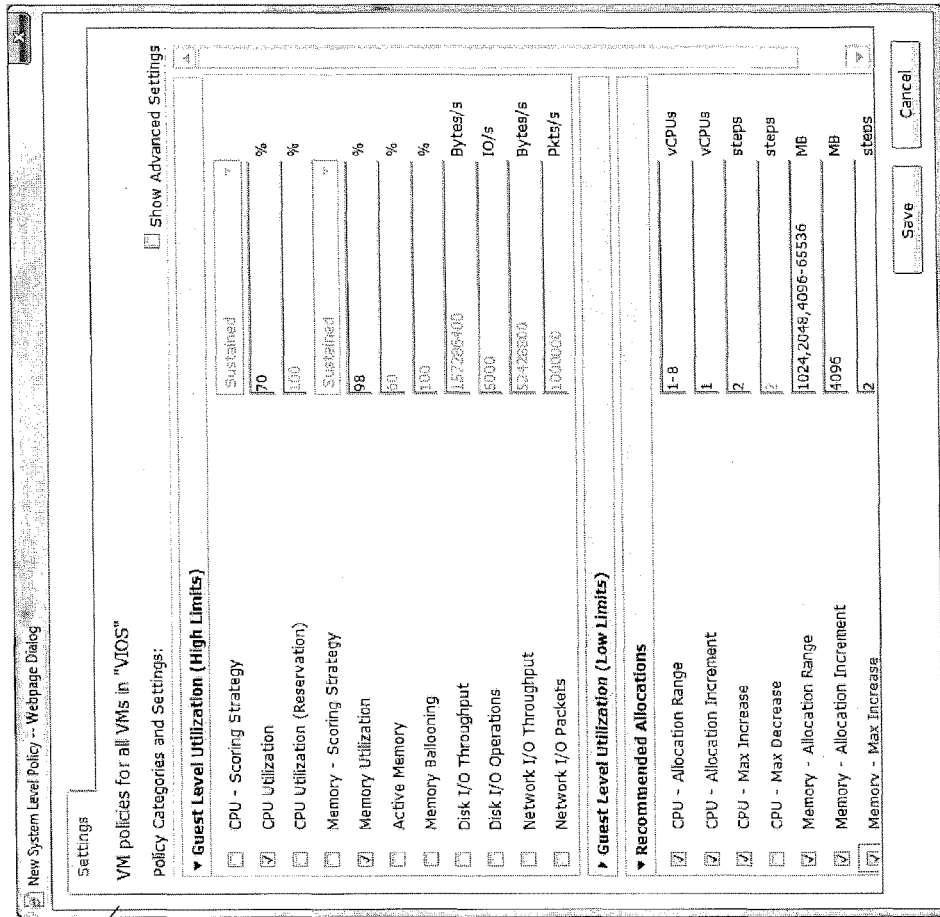


FIG. 12

US 9,654,367 B2

1

**SYSTEM AND METHOD FOR
DETERMINING AND VISUALIZING
EFFICIENCIES AND RISKS IN COMPUTING
ENVIRONMENTS**

CROSS-REFERENCE TO RELATED
APPLICATIONS

This application is a continuation of International PCT Application No. PCT/CA2012/050561 filed on Aug. 16, 2012 which claims priority from U.S. Provisional Patent Application No. 61/523,912 filed on Aug. 16, 2011, both incorporated herein by reference.

TECHNICAL FIELD

The following relates to systems and methods for determining and visualizing efficiencies and risks in computing environments.

DESCRIPTION OF THE RELATED ART

Modern data centers typically comprise hundreds if not thousands of servers. Each server supplies a finite amount of resource capacity, typically in the form of, but not limited to: central processing unit (CPU) capacity, memory or storage capacity, disk input/output (I/O) throughput, and network I/O bandwidth. Workloads running on these servers consume varying amounts of these resources. With the advent of virtualization and cloud technologies, individual servers are able to host multiple workloads.

Percent CPU utilization, which corresponds to the ratio of CPU usage relative to CPU capacity, is a common measure of how effectively servers are being utilized. Various other metrics may be used to determine resource utilization for computing systems. Organizations may wish to measure and evaluate efficiencies and risks in computing environments but often do not have convenient ways to perform such measurements and evaluations.

SUMMARY

In one aspect, there is provided a method comprising: obtaining resource utilization data and resource capacity data for a plurality of entities in a computing environment; obtaining at least one operational policy defining at appropriate level of at least one resource used by the computing environment according to at least one factor; and computing at least one score quantifying efficiencies and risks associated with the computing environment based on the resource utilization data, resource capacity data, and at least one operational policy.

In another aspect, there is provided a computer readable storage medium comprising computer executable instructions for performing the method.

In yet another aspect, there is provided a system for analyzing efficiencies and risks in a computing environment, the system comprising a processor and at least one memory, the memory comprising computer executable instructions for performing the method.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments will now be described by way of example only with reference to the appended drawings wherein:

FIG. 1 is a schematic diagram of a computing environment;

2

FIG. 2 is a block diagram of functional components configured to perform an efficiency and risk analysis using resource utilization and capacity data and operational policies.

FIG. 3 is a flow chart illustrating example computer executable operations that may be performed in conducting an efficiency and risk analysis;

FIG. 4 is an example screen shot including an efficiency risk spectrum for a single computing environment;

FIG. 5 is an example screen shot including an efficiency risk spectrum for a single computing environment;

FIG. 6 is an example screen shot including an efficiency risk spectrum for a single computing cluster;

FIG. 7 is an example screen shot including an efficiency risk spectrum for multiple computing environments;

FIG. 8 is an example screen shot including a recommended actions output;

FIG. 9 is an example screen shot including a recommended actions output;

FIG. 10 is an example screen shot including an efficiency risk spectrum for a cluster with recommendations applied;

FIG. 11 is an example screen shot for an operational policy user interface; and

FIG. 12 is an example screen shot for a system policy user interface.

DETAILED DESCRIPTION

It will be appreciated that for simplicity and clarity of illustration, where considered appropriate, reference numerals may be repeated among the figures to indicate corresponding or analogous elements. In addition, numerous specific details are set forth in order to provide a thorough understanding of the examples described herein. However, it will be understood by those of ordinary skill in the art that the examples described herein may be practiced without these specific details. In other instances, well-known methods, procedures and components have not been described in detail so as not to obscure the examples described herein. Also, the description is not to be considered as limiting the scope of the examples described herein.

It will be appreciated that the examples and corresponding diagrams used herein are for illustrative purposes only. Different configurations and terminology can be used without departing from the principles expressed herein. For instance, components and modules can be added, deleted, modified, or arranged with differing connections without departing from these principles.

A system and method are provided for quantifying and visualizing the efficiency and risks related to resource utilization levels, relative to the provisioned capacity of computing environments, with consideration of operational policies. In addition, the system may be configured to determine and presents recommended actions that mitigate the inefficiencies and risks detected for the computing environments being analyzed. The capabilities of the system herein described enable organizations to accurately measure efficiency and risks in physical, virtual and cloud computing environments. It has been recognized that through the recommended actions, organizations can increase efficiency and reduce risks in their computing environments.

An example of a computing environment 10 is shown in FIG. 1. Computing environments 10 can be virtual or physical. Virtual computing environments 10 may be based on various virtualization platforms such as VMware vSphere, IBM PowerVM, Microsoft Hyper-V, Oracle/Sun Logical Domains, etc. Physical computing environments 10

may be based on various server platforms such as IBM Power, Oracle/Sun SPARC and x86-based servers, etc. As illustrated in FIG. 1, a computing environment 10 is designed and/or provided to run at least one workload 12 that performs business functions and consume compute resources 14, e.g., resources 14 related to CPU, memory, disk, network, etc. The workloads 12 run on computing systems 16 such as servers that supply the computing resources 14. Each computing system 16 has a finite capacity of resources 14. Multiple computing systems 16 can form a computing cluster 18 reflecting, for example, an administrative group or a management domain. Such groups or domains may support advanced capabilities such as live migration of workloads 12 between computing systems 16, load balancing and/or high availability. Multiple computing clusters 18 can be logically grouped (e.g., by location, line of business, etc.) to form a computing environment 10 as is illustrated in FIG. 1.

It can be appreciated that the principles discussed herein apply to any one or more workloads 12 consuming any one or more resources 14 provided by any one or more computing systems 16, in any one or more computing clusters 18, in one or more computing environments 10. As such, the example shown in FIG. 1 is for illustrative purposes only.

Many computing environments 10 may be modeled through the entity types shown in FIG. 1, and may include associated parent-child relationships. Workloads 12 are considered resource consumers, typically with configurable resource allocations. Computing systems 16 such as servers are considered resource suppliers containing one or more workloads 12. Computing clusters 18 are considered collections of computing systems 16 (e.g. a server farm) with mobility of workloads 12 between computing systems 16 in a computing cluster 18 being possible. The computing environments 10 are typically defined by a collection of one or more computing clusters 18.

For example, VMware vSphere computing environments 10 can be modeled with the following entity types. A guest is considered a virtual machine running on a host for performing actual workloads 12. A host is a physical computing system 16 running the ESX hypervisor capable of running one or more virtual machines. A computing cluster 18 therefore enables hosts to be managed as a group capable of supporting capabilities such as live migration of workloads 12 between hosts, automated workload balancing and high availability of guest workloads 12. A datacenter in this example is considered a computing environment 10 including one or more computing clusters 18.

In another example, IBM PowerVM computing environments 10 can be modeled with the following entity types. Logical Partitions (LPARs) are considered virtual machines running on managed computing systems 16 for performing actual workloads 12. Managed systems are considered physical computing systems 16 (e.g. servers) capable of running one or more LPARs. A domain is considered a group of managed systems administered by a common hardware management controller (HMC). An environment in this example is a computing environment 10 including one or more management domains.

It can be appreciated that depending on the computing environment 10 and technology being modeled, additional entity types and parent-child relationships are possible. For example, workloads 12 can often be divided into multiple applications. In addition, some virtualization technologies support the creation of resource pools to divide processor and/or memory resources that are allocated by servers to their workloads 12.

Turning now to FIG. 2 an example of an analysis system 20 is shown. It can be appreciated that the analysis system 20 may be configured using software, hardware or any combination of software and hardware. For example, the analysis system 20 may reside on a personal computer, embedded computer, mobile computing device, etc. It can also be appreciated that the configuration and functional delineations shown in FIG. 2 are for illustrative purposes only. The system 20 includes an analysis engine 22 that comprises an efficiency and risk (ER) processor 24. The ER processor 24 utilizes system data 26 related to the computing systems 16 in a particular cluster 18 and/or computing environment 10 to quantify and visualize the efficiency and risks for a computing environment 10. The system data 26 includes, without limitation, resource utilization data 28 and resource capacity data 30 for conducting the analyses (as shown), and well as, for example, system configuration data and business related data (e.g., guest and host operating systems, guest workload uptime requirements, guest workload security level requirements, guest workload and host maintenance windows, guest workload balancing groups, guest workload high availability groups, etc.) The ER processor 24 also obtains operational policies 32 to be considered when analyzing such efficiencies and risks. In evaluating the efficiencies and the risks, the analysis engine 22 may output at least one ER spectrum 34 related to the computing environment 10 which, as described below, depicts efficiencies and risks in the computing environment 10 based on ER scores. The analysis engine 22 may also output recommended actions 36 based on the ER scores. The outputs 34, 36 shown in FIG. 2 may be displayed graphically as illustrated below.

As discussed above, computing resources 14 are consumed by workloads 12 and supplied by computing systems 16 such as servers. Typically, the resources 14 fall into four main areas: a) CPU—processing capacity, b) Memory—physical and virtual memory, c) Disk—disk storage and disk I/O bandwidth, and d) Network I/O—network interfaces and network I/O bandwidth.

The operational policies 32 help define the appropriate levels of resources 14 required by a computing environment 10 by considering factors such as, without limitation: performance/service level requirements, workload growth assumptions (planned and trended), uptime-related requirements (hardware failures, disaster recovery, maintenance windows, etc.), and workload placement affinity and anti-affinity (data security, load balancing, failover, etc.). It has been recognized that by combining the operational policies 32 with the actual resource utilization levels indicated in the resource utilization data 28, resource capacities indicated in the resource capacity data 30, system configuration data, and business attributes, the efficiencies and risks of a computing environment 10 can be assessed.

The efficiency and risks of a computing environment can be quantified through an efficiency/risk (ER) score for each entity. The ER score for an entity is based on its utilization levels, allocated or available resources (e.g., determined from system data 26) and operational policies 32. At a high level, the ER score reflects whether the resources for the entity are appropriately provisioned, under-provisioned, or over-provisioned.

An example range for ER scores is from 0 to 200 and the significance of the score is summarized below in Table 1 for illustrative purposes only.

TABLE 1

Example ER score ranges and descriptions	
ER Score	Description
0 to 74	Entity is under-provisioned. Lower scores indicate more severe levels of under-provisioning.
75 to 125	Entity is appropriately provisioned. Scores closer to 100 indicate more optimal provisioning levels.
126 to 200	Entity is over-provisioned. Higher scores indicate greater levels of over-provisioning

The ER score may be used to generate ER spectrums and, optionally, recommended actions and/or other recommendations for addressing efficiency and/or risk issues identified via the computed ER scores. FIG. 3 illustrates computer executable operations that may be performed by the ER processor 24 in conducting an analysis of system data 26 and operational policies 32. At 100 the system data 26 is obtained in order to analyze the resource utilization data 28 and the resource capacity data 30. At 102 the operational policy 32 (or policies 32) are obtained. The system data 26 and the operational policies 32 are used at 104 to compute one or more ER scores according to the nature of the computing environment 10 being evaluated. As will be explained in greater detail in the examples below, the ER score(s) is/are used at 106 to generate one or more ER spectrums. The ER spectrums are output at 108 and, if applicable, are displayed at 110.

As noted above, the ER scores may also be used to generate recommended actions and/or other recommendations. At 112, the ER processor 24 determines whether or not such recommendations are to be generated, e.g., by determining whether an option or input has been selected. If not, the process ends at 114. If a recommendation is to be generated, the recommendation(s) is/are generated at 116 and, if applicable, displayed at 118.

An example for computing ER scores for various entity types will now be described below.

Workload-Level ER Score

The ER score for a workload entity (e.g., vSphere guest, LPAR) is based on the following:

- Resource utilization levels of the entity (e.g. CPU utilization, memory utilization);
- Resource allocations (e.g. CPU allocation, memory allocation); and

Operational policies 32 that define the criteria to determine whether sufficient CPU and memory resources have been allocated for the entity. It may be noted that some operational policies 32 can be specified on a per-workload level. For example, different % CPU utilization high limits can be specified for different workloads 12, depending on the business or operational requirements of the different workloads 12 (e.g. production workloads 12 may have lower limits than non-production workloads 12).

The ER scores for the workload entities can be based on the results of two types of analyses:

1) Under-Provisioned Analysis—evaluates each workload entity by checking whether the entity’s resource utilization levels exceed high limits defined by operational policies 32. The check generates an under-provisioned score (UPS), in this example, ranging between 0 and 100 that reflects whether the entity is under-provisioned. For example, scores less than 75 may indicate that the entity is under-provi-

sioned, whereas scores greater than 75 indicate that the entity is appropriately provisioned.

2) Over-Provisioned Analysis—evaluates each workload entity by checking whether the entity’s resource utilization levels are below low limits defined by operational policies. The check generates an over-provisioned score (OPS) ranging, in this example, between 0 and 100 that reflects whether the entity is over-provisioned. For example, scores less than 75 may indicate that the entity is over-provisioned whereas scores greater than 75 indicate that the entity is appropriately provisioned.

Based on the under-provisioned and over-provisioned scores, the ER score for a workload entity can be determined as follows:

- If $UPS < 100$, $ER\ score = UPS$
- If $UPS = 100$, $ER\ score = 200 - OPS$

As such, the UPS may be made to take precedence over the OPS when computing the ER score of a workload entity which reflects its overall provisioning level. For example, an entity may be under-provisioned with respect to CPU utilization but over-provisioned with respect to memory. Based on the overall ER score, the entity is designated to be under-provisioned. This is appropriate since the shortage of resources typically result in more severe consequences than having excess resources (i.e. risks vs. inefficiency).

Server-Level ER Score

The ER score for computing system 16 such as a server entity (e.g. vSphere host, managed system) is based on the following:

- Resource utilization levels of the server (CPU, memory, disk I/O, network I/O utilization);
- Resource capacity of the server (CPU capacity, memory capacity, maximum disk and network I/O throughput); and
- Operational policies (criteria to determine whether server has sufficient resources).

The ER scores for server entities may be determined in the same way as those for workload entities, discussed above.

Cluster-Level ER Score

The ER score for a computing cluster 18 may be based on the results of a “defrag” analysis of workloads 12 and computing devices 16 included in the cluster 18.

A defrag analysis as herein described attempts to determine the maximum number of workloads 12 that can be placed on the minimum number of computing devices 16 (e.g. servers) subject to constraints defined by the operational policies 32.

The defrag analysis results may include the following metrics, assuming the computing devices 16 being analyzed are servers:

- 1) Fully loaded utilization (U_{FL})—minimum number of servers required to accommodate all the workloads as a percentage of the total number of servers.
- 2) Number of unused servers (SU)—number of servers with no workloads. A number of additional servers required (SR) may also be determined, which indicates the additional servers required in case there are insufficient existing servers.
- 3) Number of unplaced workloads (WU)—number of workloads that were not placed on a server.
- 4) Number of placed workloads (WP)—number of workloads that were placed on a server.
- 5) Normalized lowest placement score among all servers with at least one workload (LPS)—the value of this score ranges from 100 to the minimum target score limit (default=75). If the minimum target score limit modified so that

it is not equal to 75, this score value is normalized to ensure that it always ranges between 75 and 100.

The ER score is derived from these defrag results as follows:

Case 1: All workloads are placed and the fully loaded utilization is less than 100%

The ER score is equal to the 200 minus the fully loaded utilization.

In general, a server group is considered to be over-provisioned if the U_{FL} is less than 75% (which translates to an ER score that is greater than 125).

If the U_{FL} is between 75% and 99%, the cluster is considered to be appropriately provisioned.

Case 2: All workloads are placed and the fully loaded utilization is equal to 100%

The ER score is equal to the normalized lowest placement score which is defined to range between 75 and 100. This score indicates that the server group is provisioned appropriately.

ER scores approaching 100 indicate that cluster is optimally provisioned whereas scores nearing 75 indicate that the cluster is on the verge of being deemed as under-provisioned.

Case 3: One or more workloads are not placed and there are no unused servers

The ER score is equated to the number of placed workloads divided by the total number of workloads multiplied by 75.

In this case, the ER score will range between 0 and 75 with lower scores indicating higher ratios of unplaced workloads.

Case 4: One or more workloads are not placed but there are also unused servers

This indicates that the unplaced workloads are not suitable for the server group.

The ER score is equal to 200 minus the fully loaded utilization—but is also marked as a special case due to the presence of unsuitable workloads.

In summary, the ER score is computed as follows:

Case 1: ($WU=0$ AND $U_{FL}<100$)

$$ER\ score=200-U_{FL}$$

Case 2: ($WU=0$ AND $U_{FL}=100$)

$$ER\ score=LPS$$

Case 3: ($WU>0$ AND $SU=0$)

$$ER\ score=75*WP/(WP+WU)$$

Case 4: ($WU>0$ AND $SU>0$)

$$ER\ score=200-U_{FL}$$

Environment-Level ER Score

The ER score for a computing environment **10** reflects the efficiency and risks associated with the clusters that comprise the environment **10**.

Typically, it may be assumed that workloads **12** and computing systems **16** have no mobility between clusters **18**. For such environments **18**, the ER score is computed from the weighted average of the ER scores for each group of computing devices, e.g., a server group as exemplified below.

The weights used to combine the ER scores for each server group sum to 1 and reflect the relative resource capacities of each server group. If servers in all the groups have identical resource capacities, the weights can simply be based upon the number of servers. If the servers have

different resource capacities, the weights can be based on a particular resource **14** (e.g. CPU or memory).

Alternatively, weights can be based on the resource **14** that represents the primary constraint for the workloads **12** in the environment **10**. The primary constraint can be estimated by comparing the aggregate resource utilization of all the workloads **12** with the capacity of all the servers.

For environments **10** where there is mobility of workloads **12** and servers between the clusters **18**, the ER score can be computed from the results of a defrag analysis for the entire environment **10**—effectively treating the environment **10** as a single cluster **18**.

Efficiency and Risk Spectrums

Based on the ER scores, the efficiency and risks of computing environments **10** can be depicted in an Efficiency and Risk Spectrum **204** as shown in the screen shot **200** illustrated in FIG. 4. For a given computing environment **10**, ER spectrums display the relevant entities in one or more vertically arranged two-dimensional (x-y) coordinate systems. As shown in FIG. 4, a time scale bar **202** can be provided to allow a user to focus on a particular day or period of time.

The number of coordinate systems corresponds to the number of entity types which the environment **10** comprises. For example, the ER spectrum for a computing environment **10** modeled using 4 entity types (e.g. environment **10**, cluster **18**, host and guest) will also contain 4 coordinate systems.

The coordinate systems share a common horizontal axis representing the ER score. This axis is typically divided into three regions, corresponding to under-provisioned **206** (too little infrastructure), optimally provisioned **208** (just right) and over-provisioned **210** (too much infrastructure) entities, respectively.

Each entity is depicted as a single dot **212** in the spectrum **204**. The entity's type determines the coordinate system in which the entity is depicted. The ER score of the entity defines its x-coordinate. For environments **10** having multiple entity groups based on a parent entity type (e.g., workloads **12** and servers belonging to specific clusters **18**), the entity's group membership effectively defines its y-coordinate.

Types of ER spectrums that may be generated include:

ER Spectrum for a single environment **10**;

ER Spectrum for multiple environments **10**; and

ER Spectrum for multiple timeframes.

ER Spectrum for a Single Environment

Based on the ER scores, efficiency and risks of the entities in a computing environment **10** can be depicted in an Efficiency and Risk Spectrum **204** such as that shown in FIG. 4.

In FIG. 4, the spectrum **204** for a single environment (Houston) is organized into four vertically stacked sections corresponding to the four entity types: environment **10**, cluster **18**, host and guest. Each dot **212** in the spectrum **204** corresponds to an entity. Entities of each type are depicted in the corresponding section. If the environment **10** includes multiple clusters **18**, entities associated with each cluster **18** may be depicted in a different color and arranged vertically into separate rows.

Each entity's ER score determines where to draw the corresponding dot **212** on the horizontal axis. The horizontal axis ranges from 0 to 200 with 0 at the left-most edge, 100 at the center and 200 at the right-most edge, consistent with the above exemplary ER score ranges.

The analyses can be based on a variety of historical or projected timeframes selectable from the timeline bar **202**,

which define the scope of entities to be assessed and their respective resource utilization levels, allocations and capacities.

FIG. 5 illustrates a screen shot **200'** of another environment level ER spectrum **204'** and FIG. 6 illustrates a screen shot **200''** of a cluster level ER spectrum **204''** for the "Seattle" cluster shown in FIG. 5.

ER Spectrum for Multiple Environments

The screen shot **300** shown in FIG. 7 illustrates a pair of ER spectrums **304a**, **304b**, one for each of multiple environments **10**. Providing multiple spectrums **304a**, **304b** together as shown in FIG. 7 allows multiple environments **10** to be compared by depicting the key metrics (e.g., fully loaded utilization) of each environment **10**. It can be appreciated that users can interact with these spectrums **304a**, **304b** by selecting a specific environment **10** to access the ER spectrum **304a**, **304b** for the selected environment **10** (showing all the entities comprising the environment **10** as shown in FIGS. 4 and 5).

ER Spectrum for Multiple Timeframes

Another variant of the ER spectrum **204**, **304** can depict the key metrics of an environment **10** over time. For example, the fully loaded utilization of an environment can be charted for a given time range (e.g., each day for the last 30 days). For example, for a given environment **10** for which the fully loaded utilization has been computed over the last 30 days, a spectrum charting the historical values over the given time period can be generated. The spectrum can be oriented with ER score with on the x-axis and the time line on the y-axis. The desired spectrum snapshot may then be selected using the timeline bar **202**, **302**. Alternatively, the ER score and timelines can be transposed so that the spectrum **204**, **304** shows the ER-score on the y-axis and the time line of the x-axis.

Recommended Actions

Based on the analyses performed for each environment **10**, recommendations to mitigate the inefficiencies and risks depicted in the ER spectrum **204**, **304** can be generated.

Examples of recommended actions include:

- Per-guest resource allocation adjustments (e.g. CPU allocations and memory allocations of guests that match their actual resource utilization patterns);

- Workload rebalancing by changing guest-host placements within a given cluster **18** (e.g. move guests from busier to less busy hosts to better balance workloads within a cluster **18**);

- Adjustment of host server capacity for a given cluster **18** (e.g. addition or removal of server capacity to match requirements of actual guest workloads); and

- Number of additional host servers required to host the existing guest workloads.

These recommendations can correspond to different timeframes for a given environment. An example set of recommended actions are provided in the screen shot **400** shown in FIG. 8. As shown in FIG. 8, the screen shot **400** may include a series of tabs **402** including different time periods. The "Today" tab **404** is shown in FIG. 8 and includes a modification type **408** at the beginning of each row of the recommendation chart **406** to identify, e.g., whether the recommendation relates to an inefficiency or risk. FIG. 9 illustrates another example screen shot **400'** showing a recommendation chart **406'**

In general, it can be appreciated that the implementation of the recommended actions should reduce inefficiencies and risks, resulting in entities moving towards the optimal (just right) region of the spectrum **204**, **304**.

Another example of a recommended action applies to virtual environments managed as a cloud computing environment. Specifically, many cloud computing environments are managed with cloud instance sizes that are configured with pre-defined resource allocations (e.g. small=1 virtual CPU and 2 GB of memory, medium=2 virtual CPUs and 4 GB of memory, large=4 virtual CPUs and 8 GB of memory, etc.). In such environments, the recommended action may be to propose an alternate cloud instance size based on the workload's actual utilization levels and applicable policies.

An additional mode for the ER spectrum **204**, **304** can be generated where the recommended actions are assumed to have been performed, e.g., as shown in screen shot **200'''** illustrated in FIG. 10. In this scenario, the ER scores for each level (workloads **12**, hosts, clusters **18**, environment **10**) are recomputed based on the application of the recommended actions such as allocation changes and workload placements. In general, the position of the entities in the resulting spectrum **204**, **304** where the recommended actions are performed will tend to move towards the center of the spectrum. Another possible mode involves recomputing the ER spectrum **204**, **304** based on a subset of the recommended actions. Another possible mode involves computing the ER spectrum based on a set of actions specified by the user to model a desired scenario—e.g. add workloads **12**, remove workloads **12**, add hosts, upgrade resource capacity of hosts, etc.

FIG. 11 illustrates an example screen shot **500** for an operational policy user interface. As shown in FIG. 11, the policy name and description and various settings can be edited and/or set by the user. Such settings include those related to high limits for CPUs, memory, CPU reservations, and memory reservations.

FIG. 12 illustrates an example screen shot **600** for a system-level policy user interface. As shown in FIG. 12, various policy categories and settings can be edited through the illustrated user interface, e.g., those for guest level utilization (high limits), guest level utilization (low limits), recommended allocations, etc.

It will be appreciated that any module or component exemplified herein that executes instructions may include or otherwise have access to computer readable media such as storage media, computer storage media, or data storage devices (removable and/or non-removable) such as, for example, magnetic disks, optical disks, or tape. Computer storage media may include volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information, such as computer readable instructions, data structures, program modules, or other data. Examples of computer storage media include RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by an application, module, or both. Any such computer storage media may be part of the analysis engine **22**, ER processor **24**, any component of or related to the system **20**, etc., or accessible or connectable thereto. Any application or module herein described may be implemented using computer readable/executable instructions that may be stored or otherwise held by such computer readable media.

The steps or operations in the flow charts and diagrams described herein are just for example. There may be many variations to these steps or operations without departing

US 9,654,367 B2

11

from the principles discussed above. For instance, the steps may be performed in a differing order, or steps may be added, deleted, or modified.

Although the above principles have been described with reference to certain specific examples, various modifications thereof will be apparent to those skilled in the art as outlined in the appended claims.

The invention claimed is:

1. A method performed by a processor in a computing system, the method comprising:

obtaining resource utilization or performance data pertaining to a plurality of computing entities in a computing environment, and capacity data specifying resource capacities for the plurality of computing entities in the computing environment;

obtaining at least one operational policy defining criteria to determine whether the utilization or performance of an entity is in an acceptable range relative to its capacity or performance limits;

computing at least one score quantifying efficiencies and/or risks associated with corresponding ones of the entities in the computing environment, based on the resource utilization or performance data, the capacity data, and the at least one operational policy; and

displaying an indicator for at least one of the plurality of computing entities in a graphical representation based on the corresponding score;

wherein each indicator is positioned in the graphical representation according to the corresponding score such that the positioned indicator shows in a spatial manner, relative efficiencies and/or risks for the corresponding entity by positioning the indicator in one of a first portion indicative of risk associated with having infrastructure in the computing environment that cannot service workload demands and meet criteria specified in the at least one operational policy, a second portion indicative of an amount of infrastructure in the computing environment that can service workload demands based on the at least one operational policy, or a third portion indicative of inefficiencies associated with having more than the required amount of infrastructure in the computing environment to service workload demands based on the at least one operational policy.

2. The method of claim 1, wherein the at least one score is based on an under provisioned score and an over provisioned score.

3. The method of claim 2, wherein the under provisioned score is given precedence over the over provisioned score.

4. The method of claim 1, wherein the at least one score is generated using factors specific to a corresponding one of a plurality of entity types.

5. The method of claim 4, the plurality of entity types comprising any one or more of a workload, a computing system, a cluster of computing systems, and the computing environment.

6. The method of claim 4, wherein for an entity representing a cluster of computing systems, the at least one score being determined by considering a fully loaded utilization value (U_{FL}), an unused servers value (SU), an unplaced workloads value (WU), policies business separation and a placed workloads value (WP).

7. The method of claim 6, wherein:

for $WU=0$ and $UFL < 100$, the at least one score is computed as $200 - U_{FL}$;

for $WU=0$ and $UFL=100$, the at least one score equals a lowest placement score;

12

for $WU > 0$ and $SU=0$, the at least one score is computed as $75 * WP / (WP + WU)$; and

for $WU > 0$ and $SU > 0$, the at least one score is computed as $200 - UFL$.

8. The method of claim 1, wherein the graphical representation displays the scores spatially relative to each other, for each computing entity in the computing environment by placing each indicator in a corresponding portion of the graphical representation.

9. The method of claim 1, wherein the portions of the graphical representation are arranged to create a spectrum.

10. The method of claim 1, further comprising displaying a plurality of graphical representations.

11. The method of claim 1, further comprising using the at least one score to generate at least one recommendation for addressing inefficiencies and/or risks in the computing environment.

12. The method of claim 11, further comprising displaying the at least one recommendation in the graphical representation, the graphical representation comprising one or more recommendations for each entity in the computing environment that has been evaluated.

13. A non-transitory computer readable storage medium comprising computer executable instructions for:

obtaining resource utilization or performance data pertaining to a plurality of computing entities in a computing environment, and capacity data specifying resource capacities for the plurality of computing entities in the computing environment;

obtaining at least one operational policy defining criteria to determine whether the utilization or performance of an entity is in an acceptable range relative to its capacity or performance limits;

computing at least one score quantifying efficiencies and/or risks associated with corresponding ones of the entities in the computing environment, based on the resource utilization or performance data, the capacity data, and the at least one operational policy; and

displaying an indicator for at least one of the plurality of computing entities in a graphical representation based on the corresponding score;

wherein each indicator is positioned in the graphical representation according to the corresponding score such that the positioned indicator shows in a spatial manner, relative efficiencies and/or risks for the corresponding entity by positioning the indicator in one of a first portion indicative of risk associated with having infrastructure in the computing environment that cannot service workload demands and meet criteria specified in the at least one operational policy, a second portion indicative of an amount of infrastructure in the computing environment that can service workload demands based on the at least one operational policy, or a third portion indicative of inefficiencies associated with having more than the required amount of infrastructure in the computing environment to service workload demands based on the at least one operational policy.

14. A system for analyzing efficiencies and risks in a computing environment, the system comprising a processor and at least one memory, the memory comprising computer executable instructions for:

obtaining resource utilization or performance data pertaining to a plurality of computing entities in a computing environment, and capacity data specifying resource capacities for the plurality of computing entities in the computing environment;

US 9,654,367 B2

13

obtaining at least one operational policy defining criteria to determine whether the utilization or performance of an entity is in an acceptable range relative to its capacity or performance limits;
computing at least one score quantifying efficiencies and/or risks associated with corresponding ones of the entities in the computing environment, based on the resource utilization or performance data, the capacity data, and the at least one operational policy; and
displaying an indicator for at least one of the plurality of computing entities in a graphical representation based on the corresponding score;
wherein each indicator is positioned in the graphical representation according to the corresponding score such that the positioned indicator shows in a spatial manner, relative efficiencies and/or risks for the corresponding entity by positioning the indicator in one of a first portion indicative of risk associated with having infrastructure in the computing environment that cannot service workload demands and meet criteria specified in the at least one operational policy, a second portion indicative of an amount of infrastructure in the computing environment that can service workload demands based on the at least one operational policy, or a third portion indicative of inefficiencies associated with having more than the required amount of infrastructure in the computing environment to service workload demands based on the at least one operational policy.

14

15. The method of claim **5**, wherein scoring a cluster of computing systems comprises considering at least one of: workload affinity or anti-affinity, data security, load balancing, failover capacity, or other business related data that impacts quantification of efficiencies and risks.
16. The non-transitory computer readable medium of claim **13**, wherein the graphical representation displays the scores spatially relative to each other, for each computing entity in the computing environment by placing each indicator in a corresponding portion of the graphical representation.
17. The non-transitory computer readable medium of claim **13**, wherein the portions of the graphical representation are arranged to create a spectrum.
18. The non-transitory computer readable medium of claim **13**, further comprising instructions for displaying a plurality of graphical representations.
19. The non-transitory computer readable medium of claim **13**, further comprising instructions for:
using the at least one score to generate at least one recommendation for addressing inefficiencies and/or risks in the computing environment; and
displaying the at least one recommendation in the graphical representation, the graphical representation comprising one or more recommendations for each entity in the computing environment that has been evaluated.

* * * * *

Exhibit 17

Application No: 14/180,438
Amendment Dated: April 22, 2016
Reply to Office Action of: October 22, 2015

Docket No: 59612/00064

REMARKS

Applicant wishes to thank the Examiner for reviewing the present application.

Claim Amendments

In the claims, claim 1 has been amended to make it explicit that the method is performed by a processor in a computing system (see Fig. 2 and paragraph [0030]); to clarify the technical meaning of the terms “resource utilization or performance data” and “capacity data”, based on e.g. paragraphs [0024], [0030], [0032] and [0039]); and to amend the definition of the term “operational policy” for clarity and consistency with the description e.g. at paragraphs [0041] to [0044].

Claim 1 has also been amended to add a step of “displaying an indicator for at least one of the plurality of computing entities in a graphical representation to indicate, in a spatial manner, relative efficiencies and risks for the at least one of the plurality of computing entities according to a corresponding score”; based on subject matter from original claims 8-9 and the disclosure at paragraph [0030], last sentence “The outputs ... may be displayed graphically ...”; see also Figures 4-6 and associated passages of the description.

Claims 8-11, and 13 have been amended to be consistent with claim 1 as amended.

Claims 14 and 15 have been amended in a manner consistent with claim 1 as amended.

Claim 14 has also been amended inserting “non-transitory” into the preamble.

Claim 16 is new and is dependent on claim 1. Support for new claim 16 can be found in at least paragraphs [0030], [0032], and [0041] of the application as filed.

Claims 17-21 are new, are dependent on claim 14, and correspond to the subject matter recited in claims 8-13 as amended.

Applicant respectfully submits that no new subject matter has been added by way of these amendments.

Claim Rejections – 35 U.S.C. 101

Claims 1-14 have been rejected under 35 U.S.C. 101 for being directed to non-statutory subject matter.

Claim 1 has been amended to specify that the method is performed by a processor in a computing system, consistent with claim 15, which has not been rejected under 35 U.S.C. 101. Furthermore, claim 1 has been amended to specify a step of “displaying”. Claim 14 has been

Application No: 14/180,438
Amendment Dated: April 22, 2016
Reply to Office Action of: October 22, 2015

Docket No: 59612/00064

amended to insert “non-transitory” as suggested by the Examiner.

Applicant notes that giving visual indications about technical resource capacity/utilization conditions in a complex computing system (e.g. a data centre), as is done in the claimed invention, constitutes a technical problem.

This technical problem is solved by the present invention by computing efficiency scores for computing entities based on resource utilization data, resource capacity data and operational policies relating to resource allocation. These scores thus reflect conditions prevailing in an apparatus or system – in particular conditions as to the utilization and availability of resources (such as CPU/memory capacity) – and these conditions are visually indicated by generating a graphical display with an indicator for at least one computing entity.

In this way, the claimed invention thus provides visual indications about conditions prevailing in the system (see e.g. Figure 4, allowing the user to recognize which parts of the system are low on resources or have surplus resources from the placement of the indicators). Thus, Applicant respectfully submits that the claims, as herein amended, comply with 35 U.S.C. 101.

Claim Rejections – 35 U.S.C. 103

Claims 1-5 and 8-15 have been rejected under 35 U.S.C. 103(a) as being unpatentable over Hogan (US 2012/0140620) in view of Bhagat (US 2012/0011077). Claims 6 and 7 have been rejected under 35 U.S.C. 103(a) as being unpatentable over Hogan in view of Cherkasova (US 2008/0028409). Applicant respectfully traverses the rejections as follows.

The present application describes and claims a system and method for quantifying and visualizing the efficiency and risks related to resource utilization levels, relative to the provisioned capacity of computing environments, with consideration of operational policies. The capabilities of the system herein described enable organizations to accurately measure efficiency and risks in physical, virtual and cloud computing environments.

As described in paragraph [0030] and shown in Figure 2, an ER processor (24) obtains operational policies (32) to be considered when analyzing such efficiencies and risks. In evaluating the efficiencies and the risks, the analysis engine (22) may output at least one ER spectrum (34) related to the computing environment 10 which, as described below, depicts efficiencies and risks in the computing environment (10) based on ER scores. An example of such a graphical representation provided by the ER spectrum (34) is shown in Figure 4. The screen shot (200) shown in Figure 4 indicates, in a spatial manner, relative efficiencies and risks for the at least one of the plurality of computing entities that has been analyzed, according to a corresponding score. In

Application No: 14/180,438
Amendment Dated: April 22, 2016
Reply to Office Action of: October 22, 2015

Docket No: 59612/00064

this way, the efficiencies and risks can be visualized and compared to each other or to themselves over time, etc. (see also paragraphs [0090] – [0097]).

Claim 1 has been amended to clarify the protection being sought and recites in part:

“obtaining at least one operational policy defining criteria to determine whether the utilization or performance of an entity is in an acceptable range relative to its capacity or performance limits;
computing at least one score quantifying efficiencies and risks associated with the computing environment based on the resource utilization or performance data, the capacity data, and the at least one operational policy; and
displaying an indicator for at least one of the plurality of computing entities in a graphical representation to indicate, in a spatial manner, relative efficiencies and risks for the at least one of the plurality of computing entities according to a corresponding score;” [emphasis added]

Applicant respectfully submits that the combinations of cited references fail to teach or suggest what is recited in claim 1 and in particular, none of the cited references teaches or suggests: “displaying an indicator for at least one of the plurality of computing entities in a graphical representation to indicate, in a spatial manner, relative efficiencies and risks for the at least one of the plurality of computing entities according to a corresponding score”.

Hogan teaches a policy control system for communications networks. A predictive indicator is used to forecast network resource utilization. A policy-based decision is then made based on the predictive indicator. While Hogan generally considers utilization and considers operational policies, as acknowledged by the Examiner on page 5 of the Office Action, Hogan fails to teach or suggest computing a score quantifying efficiencies and risks associated with a computing environment. Applicant also notes that because Hogan is silent regarding such a score, let alone utilizing such a score, it follows that Hogan also fails to teach or suggest displaying an indicator in a graphical representation as recited in claim 1.

Bhagat has been cited as teaching what is missing from Hogan, in particular the aforementioned scores quantifying efficiencies and risks associated with the computing environment. Applicant, however, is unclear regarding the relevance of Bhagat and requests further clarification should the Examiner re-apply Bhagat. For instance, in paragraph [0027] of Bhagat, there is no

Application No: 14/180,438
Amendment Dated: April 22, 2016
Reply to Office Action of: October 22, 2015

Docket No: 59612/00064

discussion regarding computing a score, let alone one that considers policies or quantifies efficiencies and risks. In fact, paragraph [0027] is completely silent regarding quantifying any of the “risks” enumerated in that paragraph. Similarly, while claim 1 of Bhagat mentions tracking particular risks, there is nothing to suggest a computation that quantifies such risks.

Nevertheless, Bhagat also fails to teach or suggest: “displaying an indicator for at least one of the plurality of computing entities in a graphical representation to indicate, in a spatial manner, relative efficiencies and risks for the at least one of the plurality of computing entities according to a corresponding score”; and thus also fails to teach this additional aspect missing from Hogan.

Regarding Cherkasova, while also failing to teach what is missing from Hogan per the above, Applicant also respectfully disagrees that Cherkasova teaches or suggests what is recited in claims 6 and 7. It is unclear how paragraphs [0003], [0114], and [0115] even suggest the type of computation recited in claims 6 and 7. Applicant therefore kindly requests that the rejections that include the Cherkasova also be withdrawn.

In view of the foregoing, Applicant respectfully submits that the claims presented herein are patentable over the references cited by the Examiner.

* * *

Applicant requests early reconsideration and allowance of the present application.

Respectfully submitted,

/Brett J. Slaney/

Brett J. Slaney
Agent for Applicant
Registration No. 58,772

Date: April 22, 2016

BLAKE, CASSELS & GRAYDON LLP
199 Bay Street
Suite 4000, Commerce Court West
Toronto ON M5L 1A9
Canada

Tel: 416-863-2518

BS/lxi

Exhibit 18

vRealize Operations 6.2: Intelligent Workload Placement with DRS

Today I want to discuss [VMware vRealize Operations Manager v6.2](#) and specifically its Intelligent Workload Placement feature. This feature works in conjunction with, and complements DRS to help VMs get the required resources they need, ensuring better performance of the environment and applications.

[Distributed Resource Scheduler](#) also known as DRS is a well-known and proven vSphere feature that moves VMs *within* a cluster to ensure VMs are always running on a host with adequate resources to support it. vRealize Operations Manager's moves VMs *between* clusters to ensure the clusters are balanced in the environment, which in the end helps DRS.



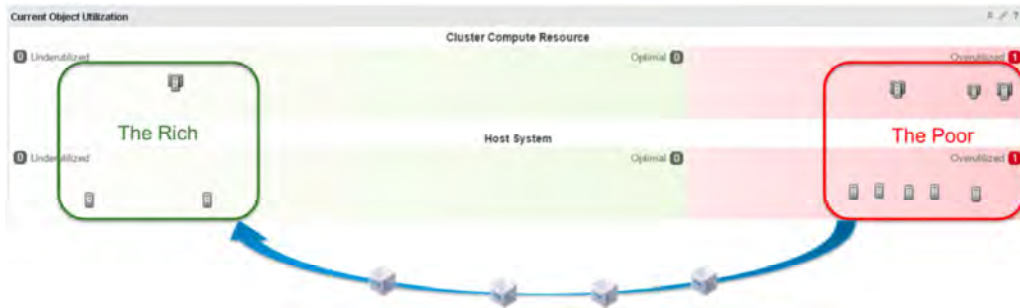
[vRealize Operations Manager's](#) new Rebalance Container action drives the placement and allows you to balance workloads between the clusters in your Datacenter or Custom Datacenters by providing you move recommendations. These move recommendations come in the form of a rebalance action plan. The plan lists move recommendations, includes the Source Cluster -> Destination Cluster mapping and provides a Reason like CPU or Memory imbalance. Once you review the recommendations, you simply need to click the "Begin Action" button to start the moves. Simple!




Next lets look at the new the Current Object Utilization widget, where we can see how well balanced our environment truly is and determine if a rebalance action is called for. This widget shows the workloads spread across the clusters and hosts in the environment. The chart is partitioned into three ZONES based on the workload levels: Underutilized, Optimal and Overutilized. The objects are placed onto the graph based on the amount of workload demand it is experiencing.



Let's look at a very common situation that I like to call the ***Robin Hood Scenario***. In this example, we have a datacenter with a number of clusters in them. As you can see from the chart below, some clusters are Overutilized and the VMs might be starving for resources, while some others are Underutilized and have plenty of room for available. The resolution is to run the Rebalance Container action to move some VMs to the open cluster and “rob from the rich to give to the poor” to better balance out the infrastructure and alleviate potential resource bottlenecks.

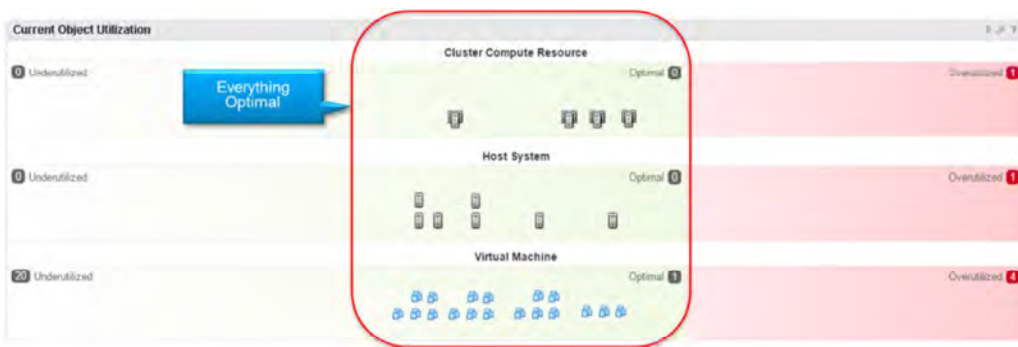


In this next example, we have a datacenter where it looks like everything is Overutilized, but when we run Rebalance Container action and look at the recommendations we find that some of the clusters are struggling with CPU and others with memory. By simply moving a few VMs between them, we can better balance the load and relieve this stressful situation.



In either scenario, it will take only a few minutes for the rebalance action to complete. Once its done, we can verify the results in the Current Object Utilization widget. [vRealize Operations Manager](#) has successfully rebalanced these clusters and they now have very similar workload levels.

In addition, looking at the host systems below, since we have better distributed the workloads across the clusters, DRS has more resources available to leverage. As a result, DRS has been able to move some VMs within the clusters, due to which, the hosts are also more balanced across the entire datacenter.



If you would like to see this in action, you can watch the video:

https://youtu.be/itZ_Mz8Kn5A

Thank you!

Exhibit 19

Self-Driving all the way to the Host? Oh yeah Host Based Placement!

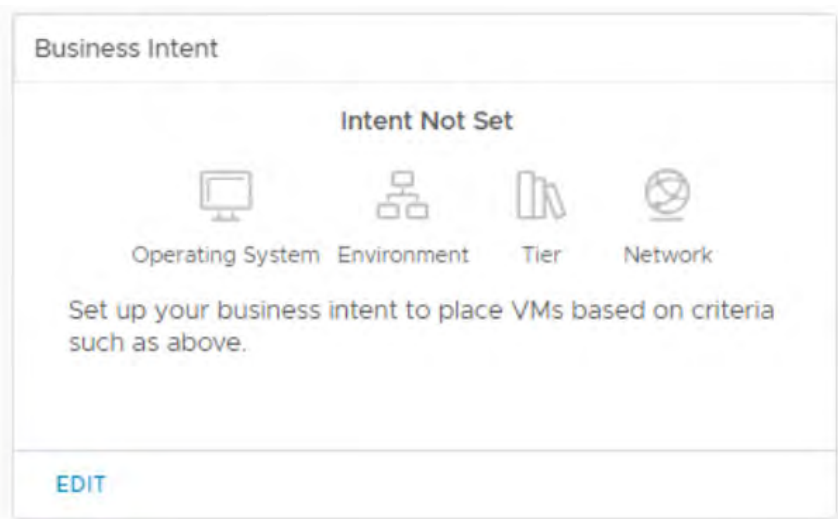
As you hopefully saw in my first blog, [vRealize Operations 7.0](#) is the key to your new self-driving driving datacenter! If you missed it you can read it [HERE](#). Enjoy! But that's not the end of the self-driving story. You can now drive your business intent (license enforcement, SLAs, compliance, etc) all the way down to your hosts with Host Based Placement! Let me explain...

Many customers have LARGE clusters containing hosts serving different business purposes or datacenters with single massive clusters in them. Driving business intent at the cluster level just doesn't make sense in these situations. Instead they need the ability to create logical boundaries around hosts within the cluster and honor business intent within these boundaries (these hosts for Oracle, these hosts for MSFT, etc). Until now customers could only do this with DRS, but that is complex and maintaining DRS rules to do this is too difficult. Wouldn't it be GREAT if you could automatically drive your business intent from ONE place? The answer is [vRealize Operations](#)! With 7.0 you can:

- Have your business intent jointly honored by vRealize Operations and DRS
- Greatly simplify the creation of DRS rules for your business needs
- Automatically fix business intent violations at the cluster AND host level

Let me show you how it works...

- First go into the Business Intent widget in the Workload Optimization screen.



- Next select Host Based tag placement. Note: You can use either Host Based or Cluster Based intents, but not both in the same DC/CDC. This capability also uses vSphere tagging to match VMs with hosts. Simply pick the tags you want to enforce/honor.

Edit Business Intent

Move VMs to either clusters or hosts with matching tags



Select the criteria you would like to use for placement of VMs. For host-based placement, only one category can be used at a time. This will ensure VMs are mapped to the appropriate hosts if moved.

[CLEAR ALL TAGS](#)

Operating System

License Enforcement [INCLUDE TAG](#)

[License Enforcement Linux](#) [License Enforcement MSFT](#)

[License Enforcement Oracle](#)

- Host Based Placement works *WITH* vROps will create DRS rules to enforce your business intent between hosts. Basically, the tagging information for VMs and hosts is used to create VM group and host group affinity rules.
- If you have DRS rules for VM-VM and VM-Host affinity/anti-affinity then vROps will check them and let you know if your business intent conflicts with your user-created DRS rules. vROps will also check every 5 mins to make sure now NEW rules conflict with your business intent. You have the option to review them and if you agree, the user-created rules will be disabled. They can be enabled later if you decide to not use Host Based Placement.

To set host level placement constraints, vRealize Operations will automatically create and manage DRS rules. All conflicting user-created DRS rules need to be **DISABLED**. These include the following:

- VM-VM affinity / anti-affinity
- VM-Host affinity / anti-affinity

1 user-created DRS Rule(s) will be disabled across clusters in this datacenter

Cluster 01		
1 Rule(s)		
Name	Type	Conflicts With
VM to Host Rule 1	VMs should/must run...	2 tags

I understand that vRealize Operations will disable all my current and future DRS rules.

[CANCEL](#) [SAVE](#)

- vROps will now create DRS groups and rules in vCenter. The objects created by vROps have naming conventions to make them easy to identify. For instance, each rule will follow a vROps_◦ <Tag Category>_<Tag>_AR format. For example vROps_License_Oracle_AR (the “AR” stands for “Automated Rule”)

VM/Host Groups

+ Add... X Delete

Name	Type
Host Group 1	Host Group
VM Group 1	VM Group
vROps_License Enforcement_Oracle_HG	Host Group
vROps_License Enforcement_Oracle_VG	VM Group
vROps_License Enforcement_Linux_HG	Host Group
vROps_License Enforcement_Linux_VG	VM Group
vROps_License Enforcement_MSFT_HG	Host Group
vROps_License Enforcement_MSFT_VG	VM Group

VM/Host Rules

+ Add... Edit... X Delete

Name
VM to Host Rule 1
vROps_License Enforcement_Oracle_AR
vROps_License Enforcement_Linux_AR
vROps_License Enforcement_MSFT_AR

This short video on Host Based Placement really explains it best.

Exhibit 20

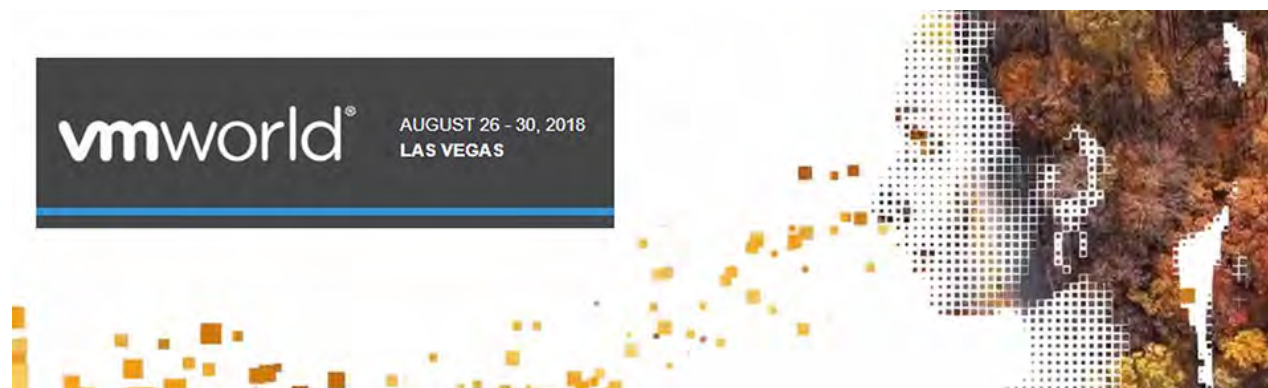
VMworld Las Vegas – Recommended Self-Driving Operations Sessions

Yang Liang (https://blogs.vmware.com/management/author/yang_liang) posted August 16, 2018

0 Comments

Please Enable Cookies Below To View.

VMworld Las Vegas is less than two weeks away, get excited! With the latest release of vRealize Operations (http://www.vmware.com/products/vrealize-operations.html?src=so_5703fb3d92c20&cid=70134000001M5td), there are many sessions that amplify the message of Self-Driving Operations. See below some sessions that we recommend you attend to learn more!



(<https://www.vmworld.com/myvmworld.jspa>)

Recommended Self-Driving Operations Sessions:

- **MGT1201BU** – Self-Driving Operations: What's New with vRealize Operations
- **MGT2934BU** – Optimize Workload Cost and Performance Using the vRealize Suite
- **MGT2552BU** – Troubleshooting Made Easy
- **MGT2566BU** – HCI Management Using vRealize Operations and vCenter
- **MGT1534BU** – vRealize Operations Capacity and Cost Management
- **MGT1440BU** – Operationalize Your World: Practical Steps Toward Proactive Operations
- **MGT1640BU** – Self-Driving Operations with vRealize Operations Performance Optimization
- **PRV1765BU** – Advanced Operations for your VMware Cloud Foundation Based Private Cloud
- **HYP2470BU** – Quickly Getting the Most out of Your VMware Cloud on AWS
- **NET2764BU** – Introduction to vRealize Network Insight
- **NET1879BU** – vRealize Network Insight Deep Dive
- **SAI2555BU** – Accelerate App Security and Availability with vRealize Network Insight

How to Register

Planning to attend the upcoming VMworld 2018 in Las Vegas (<https://www.vmworld.com/en/us/index.html>) or Barcelona (<http://www.vmworld.com/en/europe/index.html>)? This is still a good time to register (<https://www.vmworld.com/en/us/index.html>), and see you soon!

Exhibit 21

Enrollments



Labs



Transcript

Cloud Management Platform

vRealize Automation, Operations and other vRealize Suite products.

Cloud Native Applications

PhotonOS, vSphere Integrated Containers

Emerging Technologies

Blockchain, Machine Learning, OpenStack, NFV, VMware Learning Platform

Hybrid Cloud

Hybrid Cloud and Cloud Services

Hyper-Converged Infrastructure

vSAN, Virtual Volumes and VxRail

Network

Virtualization

NSX, NSX-T, Partner Solutions

Software Defined Data Center

VMware Software Defined Data Center

Solution Labs

Solve Business Challenges with VMware Solutions

Unified Endpoint Management

This catalog contains VMware Workspace ONE UEM labs

Virtualized Workspace Services

Labs Focused on Horizon, Horizon Cloud and Workspace ONE

Language: 日本語

Labs that offer a Lab Manual in Japanese

Language: 繁体中文

CLOUD MANAGEMENT PLATFORM



HOL-1971-01-CMP - What's New in vRealize Operations Manager 7.0



See what's new and improved with vRealize Operations 7. Discover Workload Optimization, new What-If planning scenarios, and much more!

This lab is available in **English**

Lab Modules

#	Title	Time	Level
1	Optimize Workloads Based on Business or Operational Intent with vRealize Operations	60	B
2	Rightsize the Configuration of Oversized or Undersized VMs with vRealize Operations	45	B
3	vSphere Optimization Recommendations with vRealize Operations	15	B
4	Assess vSphere Configuration with vRealize Operations	15	B
5	Assess vSphere Compliance Against Standards with vRealize Operations	15	B

A **Module Time** can be either 15, 30, 45 or 60 Minutes in Length.

A **Module Level** can be either Beginning, Intermediate, or Advanced.

[Close](#)

ENROLL

LAB DETAILS

TIME LIMIT: 1 hour 30 minutes
VALID PERIOD: 2 hours
EXIT ALLOWED: No
END ALLOWED: Yes



Products: [vRealize Operations \(vROps\)](#)
 VMs: Microsoft Windows Server 2012 (64-bit)



HOL-1901-01-CMP - What's New in vRealize Operations Manager 6.7 and vRealize Log Insight 4.6



Explore the new and most popular features and use cases of Self-Driving Operations from VMware vRealize Operations 6.7 and vRealize Log Insight 4.6. This lab is available in English, 日本語 [Read more](#)

ENROLL

Powered by



Exhibit 22



Exhibit 23

Watch Densify's #78—
Ryan Ellis—at Daytona
during NASCAR's
biggest weekend!



Densify
Densify

Exhibit 24



Density

Exhibit 25

Dedicated Cloud & Infrastructure Experts at Your Service

Our Densification Advisors Are Focused on Delivering
Positive Outcomes for You

Exhibit 26

Upgrade to Self-Driving Operations at up to 65% off!

Matt Jones (<https://blogs.vmware.com/management/author/jonesmatt>) posted October 16, 2018

0 Comments

Tweet



For current users of vSphere with Operations Management (vSOM) or vRealize Operations (http://www.vmware.com/products/vrealize-operations.html?src=so_5703fb3d92c20&cid=70134000001M5td) Standard, there has never been a better time to upgrade to the self-driving operations features of vRealize Operations (http://www.vmware.com/products/vrealize-operations.html?src=so_5703fb3d92c20&cid=70134000001M5td) Advanced or vRealize Suite (http://www.vmware.com/products/vrealize-suite.html?src=so_5703fb3d92c20&cid=70134000001M5td) Standard. For a limited time we're offering this upgrade at half off, or 65% off when you combine it with an upgrade to vSphere Platinum (also half off!)

You use vRealize to manage vSphere performance and capacity today, but where is your IT infrastructure heading in the future? Let's think about some common questions you might be considering:

- Are you planning data center modernization, consolidation, or build out?
- Do you need to be able to better forecast and model capacity needs and plan procurement?
- Do you need a single tool to manage vSphere (https://www.vmware.com/products/vsphere.html?int_cid=70134000001CaZ5&src=WWW_us_VMW_0cnHu1FmS2pidRE57WJU) AND other data center components (including vSAN (https://www.vmware.com/products/vsan.html?int_cid=70134000001CaZ5&src=WWW_us_VMW_0cnHu1FmS2pidRE57WJU) storage, VMware Cloud on AWS, and VMware Cloud Foundation) in a unified console?
- Are you planning a cloud migration strategy? Do you need to know the cost of running workloads in private vs multiple public clouds?

If your answer is "yes!" to any of the above, then you might strongly consider upgrading to the Advanced edition of vRealize Operations, or stepping up to the Standard edition of vRealize Suite (which bundles vRealize Operations Advanced with vRealize Log Insight and vRealize Suite Lifecycle Manager).

There are three reasons why now is the perfect time to upgrade:

1. vRealize Operations 7.0 is the best vRealize Operations yet – by far

The vision of Self-Driving Operations (<https://www.vmware.com/products/vrealize-operations.html>) is to automate and simplify operations management by delivering per the "Three Tenets" of Self-Driving Operations:

1. Intent-Driven Continuous Performance Optimization

2. Efficient Capacity Management
3. Intelligent Remediation

“Self-driving” summarizes the strategy and vision of providing continuous and automated closed-loop performance and capacity optimization at minimal cost, based on business and operational intent. In other words – as a customer, all you need to do is define your business or operational goals and let the platform take care of the rest to assure performance, densify clusters or enforce software license separation.

vRealize Operations 7.0 augments the capabilities introduced in the last release, particularly focusing on:

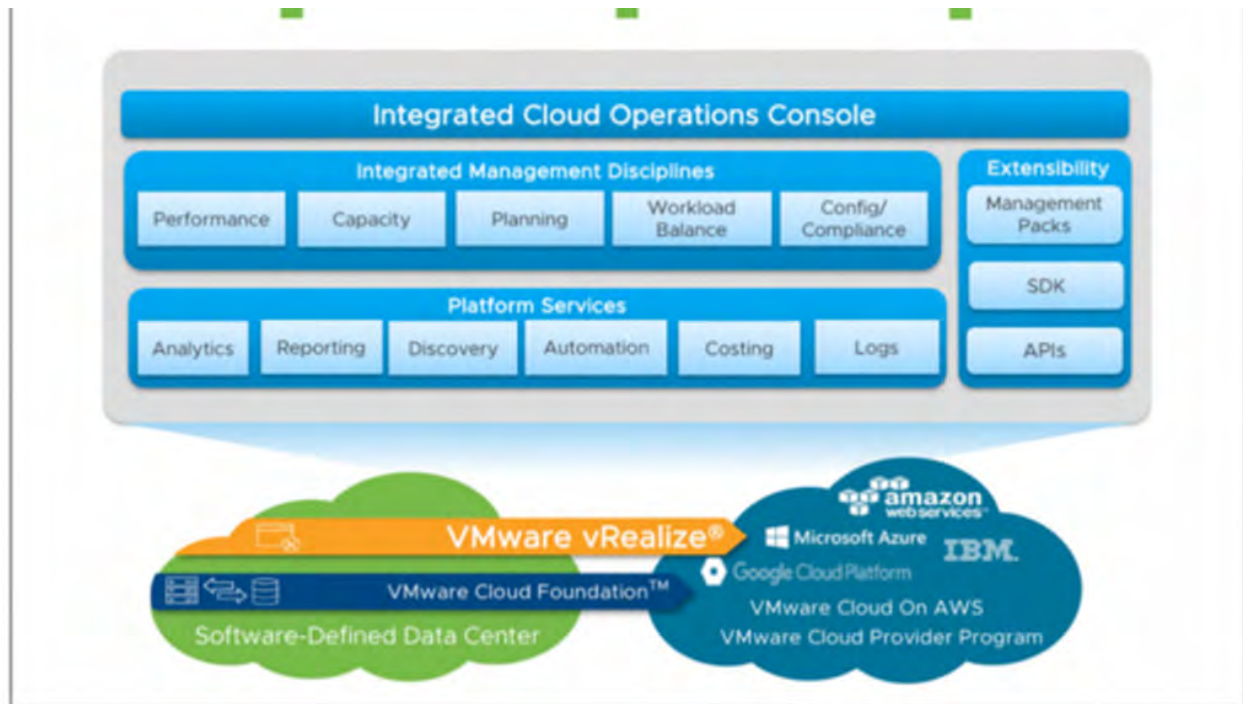
- **Business intent-driven continuous performance optimization** with new automated workload balancing capabilities, enhanced integration with vRealize Automation for both initial placement and ongoing workload balancing, and host-based workload placement
- **Efficient capacity management** by enhancing the real-time, predictive capacity & cost analytics engine and adding multiple what-if planning scenarios
- **Multi-cloud** support with migration planning across VMware Cloud on AWS and native AWS, as well as updated management packs for AWS and Kubernetes
- **Quick time to value** with updates to the UI, simplified custom dashboard creation and sharing, enhanced SDDC integrations and built-in vSphere config & compliance

For more details, be sure to check out John Dias’ technical overview (<https://blogs.vmware.com/management/2018/09/this-baby-is-loaded-whats-new-with-vrealize-operations-7-0-technical-overview.html>) of what’s new in 7.0.

Get the full self-driving capabilities of 7.0 by upgrading to vRealize Operations Advanced or vRealize Suite Standard – obtain visibility and management beyond just vSphere into the entire software-defined data center; gain access to customizable dashboards, capacity and cost optimization, cloud planning, and many of the built in automation features discussed above. Compare vRealize Operations versions here (<https://www.vmware.com/products/vrealize-operations.html#compare>) to see what all comes with an upgrade to the Advanced edition.

With vRealize Suite Standard, gain all the capabilities of vRealize Operations Advanced, in addition to vRealize Log Insight (<https://www.vmware.com/products/vrealize-log-insight.html>). Log Insight’s intelligent log management and analytics allows for speedy troubleshooting across physical, virtual, and cloud environments. See what vRealize Suite can do for you here (<https://www.vmware.com/products/vrealize-suite.html>).





2. All your favorite VMware products are “better together” with self-driving operations

Not only is vRealize Operations better than ever, but it now integrates better than ever with the rest of the VMware ecosystem. Notably, it is tightly integrated with vSAN – you can now manage and troubleshoot both compute and storage operations in the same management console.

vRealize is also fully integrated with Dell EMC VxRail, the standard in hyper-converged infrastructure. Additionally, for telco service providers, vRealize is the best management tool for vCloud NFV. Look for more convergence and integration with the rest of the VMware family moving forward!

3. Promotions!

And finally, for a limited time, **get up to 65% off** when you upgrade from vRealize Operations Standard to vRealize Operations Advanced or vRealize Suite Standard. We are running this promotion in tandem with the vSphere team as follows:

- 50% off upgrades from vSphere ENT or ENT+ to vSphere Platinum
- 50% off upgrades from vRealize Operations Standard to EITHER vRealize Operations Advanced OR vRealize Suite Standard
 - 65% off this upgrade when done together with the vSphere upgrade promotion above

Not only can existing vRealize Operations Standard customers take advantage of this promotion, but

Exhibit 27

Cloud Management Platform (CMP) – Intelligent Provisioning and Optimization

Dave Overbeek (https://blogs.vmware.com/management/author/dave_overbeek) posted November 15, 2018

0 Comments

Tweet



One of the main drivers of a Cloud Management Platform (CMP) is to have the different solutions work together seamlessly. That means the business, automation, compliance, operations, optimization, capacity, etc... All these components need to work to provide you with a unified feel and flow. Without that its really not much of a CMP... its just a bunch of features and functions cobbled together.

One such integration point that is critical to any CMP, and to your business as a whole, is that of your automation and operations tools. For VMware that means vRealize Automation (vRA) and vRealize Operations (vROps). While these two solutions have always delivered better value together, the latest releases of vRA 7.5 and vROps 7.0 have brought major enhancements to the VMware CMP solution.

Initial Workload Placement

The ability to properly provision new workloads into your environment is of utmost importance. Placing it incorrectly could lead to all sorts of issues for compliance, SLAs, application performance, cost etc. It also means you will just need to manually move the VM later! The initial workload placement process between vRA and vROps takes care of all of this for you.

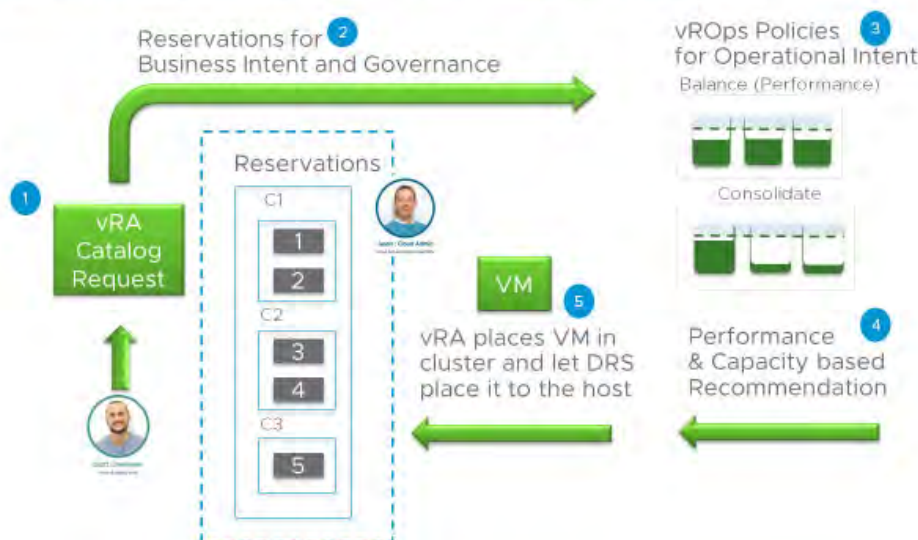
vRA provides the **governance** (infrastructure privileges – what users have access to what part of your infrastructure) and **business intent** (SLAs, license compliance, compliance and cost). vROps on the other hand is responsible for the **operational intent** (densification, balance, headroom and application performance). Together they form a powerful solution for deploying your workloads.

1. It all starts in the vRA customer portal where a user can deploy new workloads and applications from the self-service catalog. vRA's reservation policies dictate where the workload(s) can be placed in the environment based on the blueprint and the user's privileges. Does it need to run on PCI compliant clusters? Should it run on the fastest storage? Should it be placed in the test/dev datacenter?
2. The list of possible places (e.g. datacenter or custom datacenter) to deploy the workloads is sent to vROps which holds your operational intent.
3. If the deployment is for business-critical apps where performance is paramount, then a balanced approach is recommended. But if the workloads to be deployed are for testing, then maybe a more consolidated approach is warranted. vROps also ensures that any headroom settings you may have implemented, to lower the possibility of contention from sudden resource spikes, are honored. Based on your configured operational intent vROps determines the BEST place to deploy these workloads (e.g. cluster).
4. The placement recommendation is then sent back to vRA.
5. vRA deploys the workload to the correct cluster. DRS will then place the workloads on the

appropriate host(s) within the cluster.

DONE! Simple and sweet.

Initial Workload Placement – vRA 7.5 and vROps 7.0



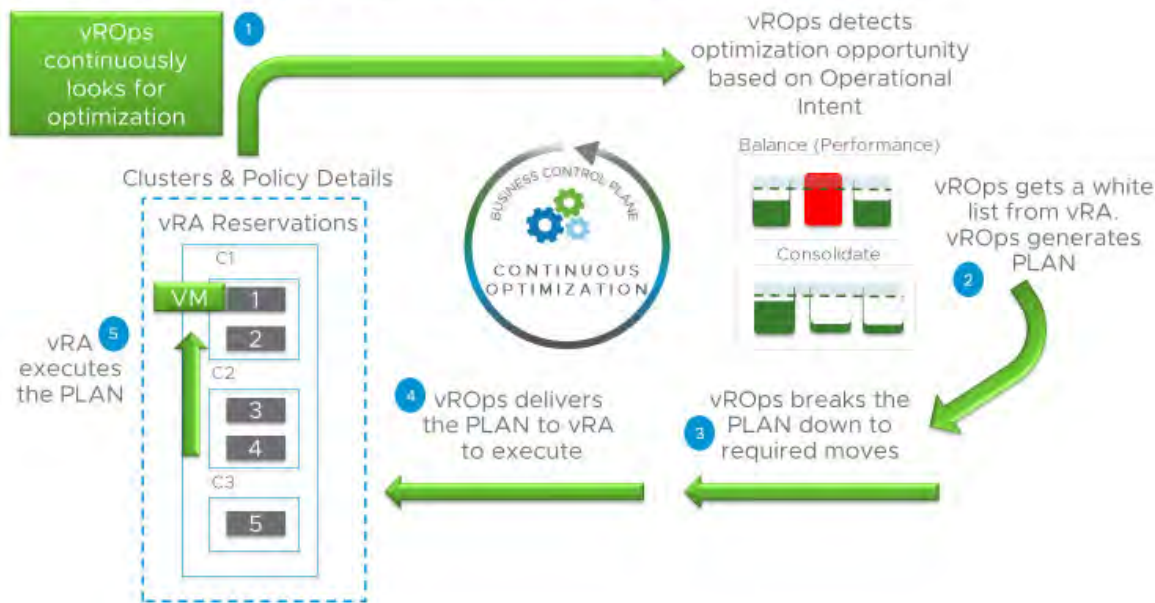
Ongoing Workload Optimization

As we’ve discussed, correctly placing workloads in their initial destination is important. Of course, we hope it stays there throughout its lifetime and never needs to move anywhere, but let’s be honest that’s never the case. Applications change, workloads are added or deleted, HW is refreshed and all of this can affect the health of your workloads/applications.

Together vRA and vROps provide automated workload optimization functionality that will ensure the business and operational intent of your workloads are met throughout the lifecycle of the applications.

1. vROps is continuously looking for optimization opportunities based on your operational intent. Is a cluster experiencing contention? Are there places where we can consolidate? Are any of the clusters breaching my configured headroom setting? If the answer to any of these is “yes” vROps will start a workload optimization process to resolve it.
2. Before recommending any moves vROps needs to understand the business intent of the workloads. To do this it grabs the reservation policies from vRA which tell vROps where workloads CAN be placed. With the knowledge of the current state of the infrastructure and your business needs, vROps formulates a workload optimization placement plan.
3. This plan is broken down into the actual moves that need to be made to bring the environment back to an operational and business “green state”.
4. The list of moves is sent to vRA.
5. vRA executes the plan thereby ensuring BOTH products in the solution are aware of the changes.

On-Going Workload Optimization – vRA 7.5 and vROps 7.0

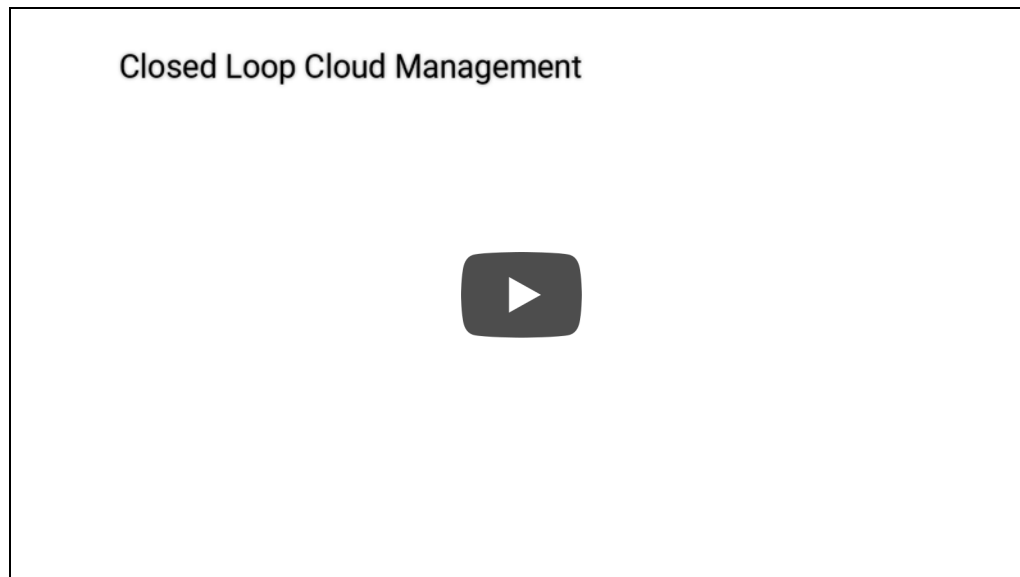


Summary

Together vRA and vROps will help you drive the BEST possible value out of your environment, keep your applications happy and healthy, help lower your overall costs, drive compliance and governance and just make your job easier overall.

Want to see it all in action? Check out this video.

https://youtu.be/SBY_NO_Lvls (https://youtu.be/SBY_NO_Lvls)



Related Posts:

Daniel Zilberman posted April 26, 2018

Ivan Ivanov posted May 3, 2018

Exhibit 28

Start Running a Self-Driving Datacenter – vRealize Operations 7.0

Workload Optimization!

Dave Overbeek (https://blogs.vmware.com/management/author/dave_overbeek) posted September 25, 2018

0 Comments

Tweet



vRealize Operations 7.0 (<https://www.vmware.com/products/vrealize-operations.html>) is the key to your new self-driving datacenter! Going self-driving will save you time and money, reduce your number of fire-drill headaches and make you look like a superstar to your bosses. To turn on self-driving you need to enable workload optimization which lets you automate the business and operational intent in your datacenters. Workload optimization can provide benefits like driving better application performance, improving consolidation ratios, adhering to SLAs, ensuring datacenter compliance or even lowering costs. What's not to like?

Workload optimization is easily accessible from your Quick Start home page. You'll find it in the Optimize Performance pillar.



Configuring workload optimization is as easy as 1, 2, 3 and can be done via the Operational Intent, Business Intent and Optimization Recommendation widgets.



Step 1 – Configuring Your Operational Intent

Select the datacenter or custom datacenter you wish to configure then click on the EDIT button on the Operational Intent widget.

Workload Automation Policy Settings

Affected Policy: BCA and Secondary Apps

Workload Optimization

Balance Consolidate

Balance will spread workloads evenly over the available resources, but may move workloads more often. This is good for more stable populations.

Moderate will minimize workload contention, but will not attempt to move workloads to achieve better balance or consolidation.

Consolidate will place workloads into as few clusters as possible, but allows for less responsive capacity. This is good for populations with steady demand, and may reduce licensing and power costs.

Cluster Headroom

0% 50%

Headroom provides a buffer of space for the cluster, which may cause a rebalance to occur earlier than it would otherwise.



- ✓ Avoid Performance Issues
- ✓ Balance Workloads

Warning: Changes to this policy will affect other containers.

CANCEL

SAVE

The first thing to configure is your target utilization objective (the workload optimization slider). This dictates how vRealize Operation will move workloads between the clusters in the datacenter to best optimize it. The default setting of moderate means workloads will only be moved when a cluster is facing resource contention. Setting it to balance means workloads are evenly spread out across the clusters to drive the best possible performance in the datacenter. On the other hand, with the consolidate setting Workloads are placed to maximize utilization to lower datacenter costs. This is sometimes called “densification”.

The other configuration that needs to be set is cluster headroom which allows you to specify how much risk is acceptable? Workload Optimization will move workloads between clusters in the datacenter and the headroom settings says when a cluster should be considered full. It provides a buffer of space to reduce the risk from bursts or unexpected resource demand in the cluster.

Once both of these are set you simply click save to commit the settings.

Step 2 – Configuring Your Business Intent

That was easy! Now we need to set up your business intent for your datacenter. Click on the EDIT button for Business Intent.

Edit Business Intent


Move VMs to either clusters or hosts with matching tags

Select the criteria you would like to use for placement of VMs. This will ensure VMs are mapped to the appropriate clusters if moved. Only one category can be prioritized at a time. VM with higher priority tags will be moved last.

[CLEAR ALL TAGS](#)

> Operating System

> Environment

▼ Tier  Enable Prioritization

Tag Category ▼ Tag Name (Optional) ▼ [INCLUDE TAG](#)

1 SLA Tiering Gold 2 SLA Tiering Silver

In this screen you specify your business needs so workload optimization can ensure they are being met by placing the workloads on the correct hosts and clusters. It does this by leveraging vSphere tags and placing workloads onto clusters and hosts where the tags match. Depending on your business needs you may wish to enable cluster based or host-based placement. In this blog we will discuss cluster-based placement. I will do a deep dive on host-based placement in an upcoming blog.

Let's start by clicking on cluster-based placement which opens the tag selection area. Here we need to specify what type of business intent we want to drive. These are free text categories and are used to best describe your business needs. You can use one of the out-of-the-box categories or make up one of your own.

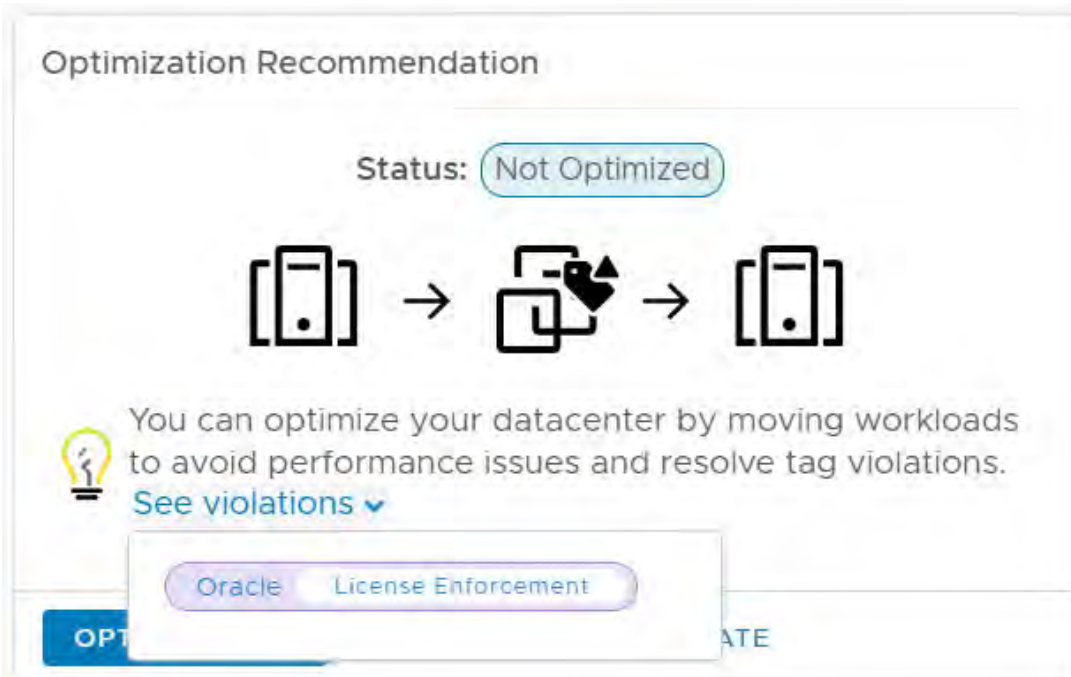
In the drop-down you will be shown any cluster based vCenter tag categories that have been configured. Once you select your category you are shown all the associated vCenter tags. Simply choose the tags you want to use to drive your business intent.

Once both are set you simple click save to commit the settings.

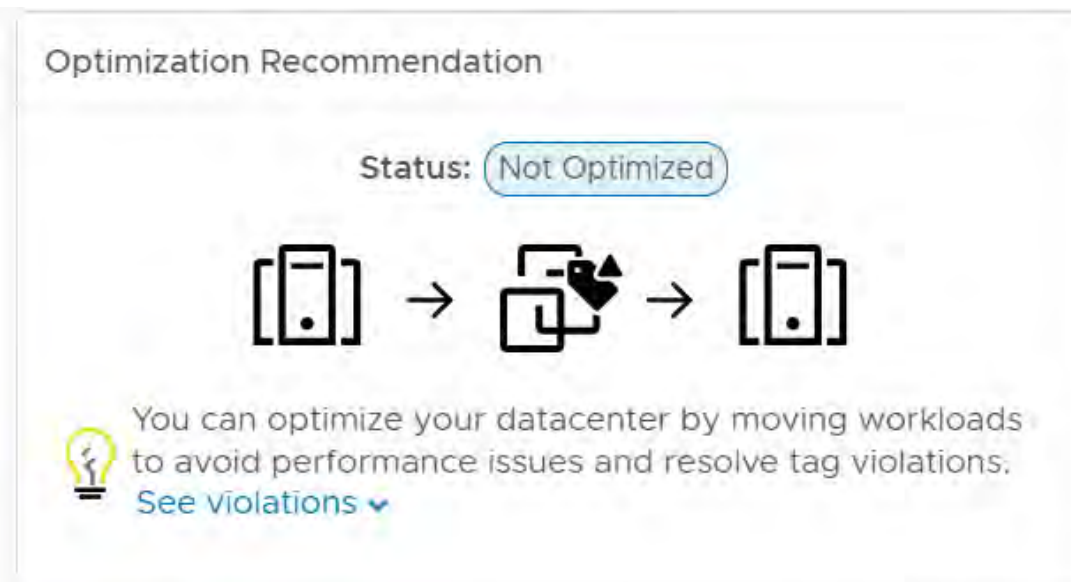
Step 3 – Configuring Automation Level


Now that you have your operational and business intent configured its time to turn workload optimization loose and start making your data center run better! The Optimization Recommendation widget will show you if your datacenter is...um ...optimized. A datacenter is flagged as Not Optimized if its operational intent is not being met. For instance, if you set the datacenter up for balance it will be flagged as Not Optimized if the clusters are out of balance.

New to 7.0 is the idea of Tag Violations for business intent. This means datacenters can be flagged as “Not Optimized” if your business intent is not met. For instance, if you are trying to drive license enforcement and you have Oracle VMs running on Microsoft clusters a Tag Violation will be shown and the datacenter will be labeled as Not Optimized. Even better, these tag violations can be resolved through workload optimization.



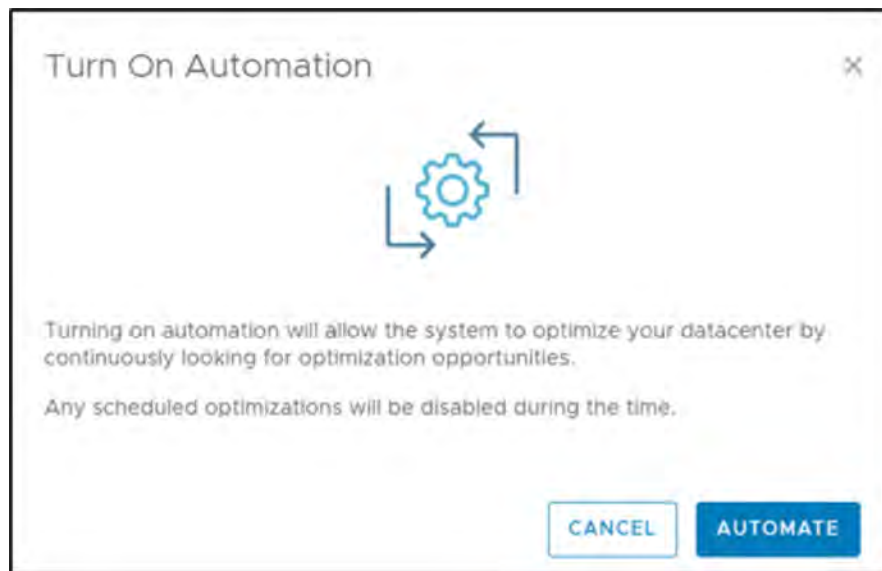
You have 3 options of how to run a workload optimization and deal with a Not Optimized datacenter: Optimize Now, Schedule or Automate.





If you wish to keep your hands on the wheel you can use the Optimize Now button to manually optimize the datacenter when you wish. You can also use the Schedule button to run workload optimizations during your maintenance windows.

However, if you don't feel like logging in a clicking a button or waiting for your maintenance window to fix these issues we have an answer for you: Automate it! vRealize Operations 7.0 enables full automation of workload optimization so you can be sure your workloads are meeting both business and operational intents around the clock. A simple click of the Automate button and vRealize Operations takes over.



If you wish to see this WHOLE THING working together I have created a video of how you can automatically set up cluster-based SLA Tiers in your datacenter in the attached video. It's a bit long, but really worth the time!

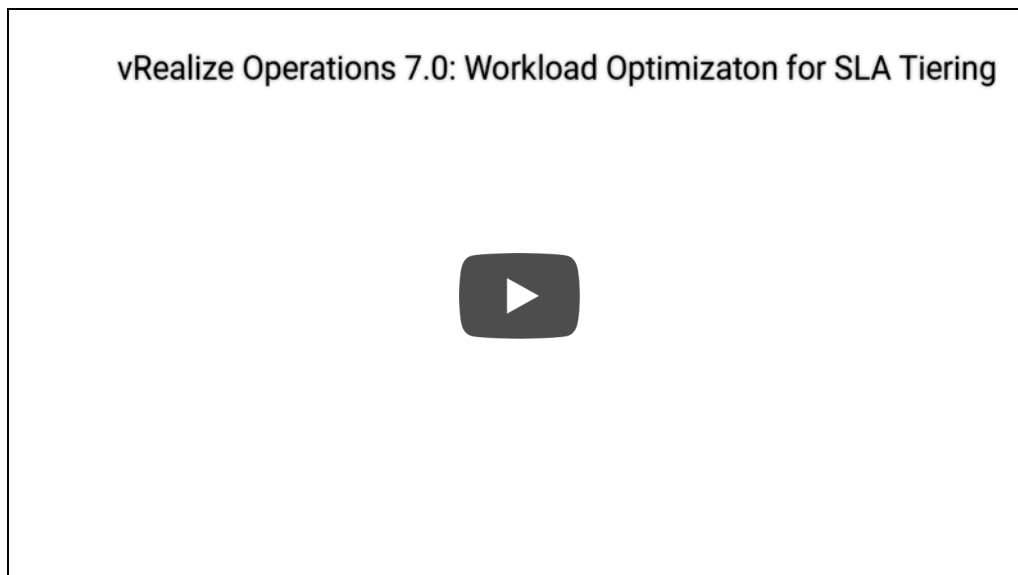


Exhibit 29

The Three Tenets of Self-Driving Operations

Run Production Operations Hands-off and Hassle-free



Continuous Performance Optimization - Assure application performance based on business and operational intent



Efficient Capacity Management - Run infrastructure like AWS - optimal densification, proactive planning and procurement



Intelligent Remediation - Predict, prevent and troubleshoot across SDDC and multiple clouds

2019 03 20 Top Reasons to Upgrade to vROps 7.0 and How to Get the Most Out of It (By VMware EMEA)

Optimized Hybrid Cloud

Cut troubleshooting time in half

40% improvement in operational efficiency

30% Capacity Savings

10% reduction in hardware costs

75% reduction in unplanned downtime

vmware Confidential | ©2018 VMware, Inc.

POWERED BY INXPO

Exhibit 30

VMware vRealize Operations

Workload Optimization -
Densifying to Repurpose Hosts

Exhibit 31

**ELECTRONIC VIDEO
SUBMITTED ON DVD.**

Exhibit 32

**ELECTRONIC VIDEO
SUBMITTED ON DVD.**

Exhibit 33

**ELECTRONIC VIDEO
SUBMITTED ON DVD.**

Exhibit 34

**ELECTRONIC VIDEO
SUBMITTED ON DVD.**

Exhibit 35

**ELECTRONIC VIDEO
SUBMITTED ON DVD.**

Exhibit 36

**ELECTRONIC VIDEO
SUBMITTED ON DVD.**

Exhibit 37

VMWARE vREALIZE AUTOMATION

AT A GLANCE

VMware vRealize® Automation™, part of VMware vRealize Suite, empowers IT to accelerate the provisioning and delivery of IT services, across infrastructure, containers, applications and custom services. Leveraging the extensible framework provided by vRealize Automation, you can streamline and automate the lifecycle management of IT resources from initial service model design, through Day One provisioning and Day Two operations. Whether your IT services are running on private cloud, public cloud or hybrid cloud, the multi-vendor, multi-cloud solution supported by vRealize Automation assures your services will be delivered with speed, control and performance.

KEY BENEFITS

- **Agility** – Automate IT service delivery processes (infrastructure, containers, applications, and any custom IT service) to rapidly respond to business needs
- **Extensibility** – Easily integrate with third-party tools across IT ecosystem to protect investments in existing and future technologies
- **Control** – Embed governance-based policies into IT services across a hybrid cloud environment to ensure compliance, performance and financial outcomes
- **Choice for Developers** – Enable self-service that can rapidly deliver building blocks to developers
- **Complete lifecycle management** – Achieve optimal workload management from initial deployment, on-going rebalance, to retirement and reclamation using the vRealize Suite

Deliver Your IT Services at Cloud Speed

The key to every successful business is agility. In the cloud era, agility gives IT departments a mandate to replace time-consuming, siloed and manual processes with end-to-end automated workflows that enable fast provisioning and delivery of IT services. As more companies start to embrace DevOps practices in order to accelerate application development, building a more automated and collaborative workflow across infrastructure and applications becomes imperative. vRealize Automation enables IT team to remove process inefficiencies through the use of end-to-end automation that helps IT better serve DevOps teams.

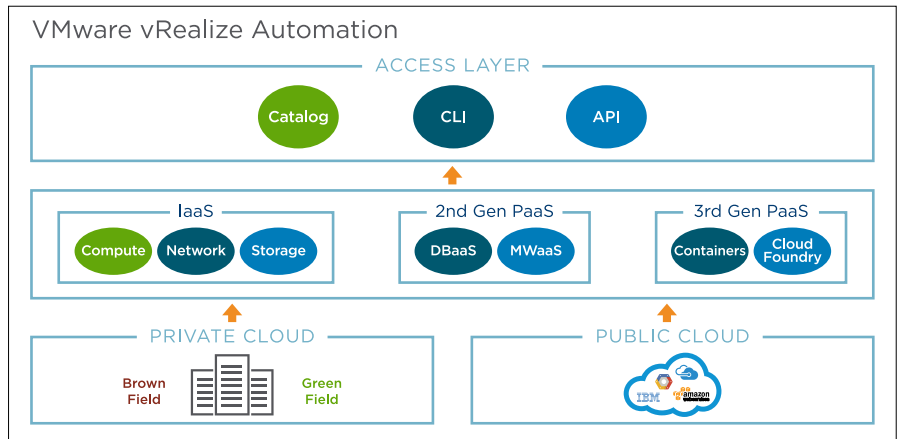


Figure 1: Support Your Cloud Journey by vRealize Automation

Exhibit 38

Foundations and Concepts

7 May 2019

vRealize Automation 7.5



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2008–2018 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Foundations and Concepts

VMware vRealize™ Automation provides a secure portal where authorized administrators, developers, or business users can request new IT services. In addition, they can manage specific cloud and IT resources that enable IT organizations to deliver services that can be configured to their lines of business in a self-service catalog.

This documentation describes the features and capabilities of vRealize Automation. It includes information about the following subjects:

- vRealize Automation components
- Common service catalog
- Infrastructure as a Service
- XaaS
- Software

For information about cost management for VMware vRealize™ Automation, see the documentation for VMware vRealize™ Business™ for Cloud.

Note Not all features and capabilities of vRealize Automation are available in all editions. For a comparison of feature sets in each edition, see <https://www.vmware.com/products/vrealize-automation/>.


Intended Audience

This information is intended for anyone who needs to familiarize themselves with the features and capabilities of vRealize Automation.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

Exhibit 39



VMware® Distributed Power Management Concepts and Use

WHITE PAPER

VMware ESX 4 and VMware vCenter Server 4

Consolidation of physical servers into virtual machines that share host physical resources can result in significant reductions in the costs associated with hardware maintenance and power consumption. VMware® Distributed Power Management (VMware DPM) provides additional power savings beyond this initial benefit by dynamically consolidating workloads even further during periods of low resource utilization. Virtual machines are migrated onto fewer hosts and the un-needed ESX hosts are powered off. When workload demands increase, ESX hosts are powered back on and virtual machines are redistributed to them. VMware DPM is an optional feature of VMware® Distributed Resource Scheduler (VMware DRS).

This information guide provides a technical overview of VMware DPM operation in VMware® ESX™ 4 and VMware vCenter™ Server 4. It is intended for VMware partners, resellers, and customers who want detailed information on VMware DPM functionality in that release.

The guide covers the following topics:

[VMware vSphere™ and Cluster Services](#)

[VMware DPM Usage](#)

[VMware DPM Operation](#)

[VMware DPM Advanced Options](#)

[VMware DPM and Datacenter Monitoring Tools](#)

[VMware DPM Usage Scenario](#)

[Resources](#)

VMware vSphere and Cluster Services

One of the key management constructs in VMware vSphere 4, which comprises VMware ESX 4 and VMware vCenter Server 4, is the cluster. Grouping multiple ESX hosts into a cluster enables you to manage them as a single compute resource. The cluster services that bring about this benefit include VMware® Fault Tolerance (VMware FT), VMware® High Availability (VMware HA), VMware DRS, and VMware DPM.

VMware FT and VMware HA handle host and virtual machine failures in a cluster of ESX hosts. It respects your settings for the desired policies and the associated resources to be set aside for use by virtual machines in the event of a failure. VMware FT and VMware HA implements mechanisms for detecting problems and restarting virtual machines. The comprehensive “VMware vSphere Availability Guide” (see [Resources](#) for a link) presents information on VMware FT and VMware HA operations. VMware FT and VMware HA failover resource constraints are respected by VMware DRS and VMware DPM.

VMware DRS manages the allocation of resources to a set of virtual machines running on a cluster of ESX hosts with the goal of fair and effective use of resources. VMware DRS makes virtual machine placement and migration recommendations that serve to enforce resource-based service level agreements, honor system- and user-specified constraints, and maintain load balance across the cluster even as workloads change. The best practices paper “Resource Management with VMware DRS” (see [Resources](#) for a link) provides material on VMware DRS usage and best practices.

Exhibit 40

vSphere Resource Management

Update 1

11 JAN 2019

VMware vSphere 6.7

VMware ESXi 6.7

vCenter Server 6.7



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2006–2019 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Create a Datastore Cluster

You can manage datastore cluster resources using Storage DRS.

Procedure

- 1 Browse to data centers in the vSphere Client.
- 2 Right-click the data center object and select **New Datastore Cluster**.
- 3 To complete the **New Datastore Cluster** wizard, follow the prompts.
- 4 Click **Finish**.

Enable and Disable Storage DRS

Storage DRS allows you to manage the aggregated resources of a datastore cluster. When Storage DRS is enabled, it provides recommendations for virtual machine disk placement and migration to balance space and I/O resources across the datastores in the datastore cluster.

When you enable Storage DRS, you enable the following functions.

- Space load balancing among datastores within a datastore cluster.
- I/O load balancing among datastores within a datastore cluster.
- Initial placement for virtual disks based on space and I/O workload.

The Enable Storage DRS check box in the Datastore Cluster Settings dialog box enables or disables all of these components at once. If necessary, you can disable I/O-related functions of Storage DRS independently of space balancing functions.

When you disable Storage DRS on a datastore cluster, Storage DRS settings are preserved. When you enable Storage DRS, the settings for the datastore cluster are restored to the point where Storage DRS was disabled.

Procedure

- 1 Browse to the datastore cluster in the vSphere Client.
- 2 Click the **Configure** tab and click **Services**.
- 3 Select **Storage DRS** and click **Edit**.
- 4 Select **Turn ON vSphere DRS** and click **OK**.
- 5 (Optional) To disable only I/O-related functions of Storage DRS, leaving space-related controls enabled, perform the following steps.
 - a Under **Storage DRS** select **Edit**.
 - b Deselect the **Enable I/O metric for Storage DRS** check box and click **OK**.

Exhibit 41

vSphere Availability

Update 2

11 APR 2019

VMware vSphere 6.7

VMware ESXi 6.7

vCenter Server 6.7



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2009–2019 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Creating and Using vSphere HA Clusters

2

vSphere HA clusters enable a collection of ESXi hosts to work together so that, as a group, they provide higher levels of availability for virtual machines than each ESXi host can provide individually. When you plan the creation and usage of a new vSphere HA cluster, the options you select affect the way that cluster responds to failures of hosts or virtual machines.

Before you create a vSphere HA cluster, you should know how vSphere HA identifies host failures and isolation and how it responds to these situations. You also should know how admission control works so that you can choose the policy that fits your failover needs. After you establish a cluster, you can customize its behavior with advanced options and optimize its performance by following recommended best practices.

Note You might get an error message when you try to use vSphere HA. For information about error messages related to vSphere HA, see the VMware knowledge base article at <http://kb.vmware.com/kb/1033634>.

This chapter includes the following topics:

- [How vSphere HA Works](#)
- [vSphere HA Admission Control](#)
- [vSphere HA Interoperability](#)
- [Creating a vSphere HA Cluster](#)
- [Configuring vSphere Availability Settings](#)
- [Best Practices for VMware vSphere® High Availability Clusters](#)

How vSphere HA Works

vSphere HA provides high availability for virtual machines by pooling the virtual machines and the hosts they reside on into a cluster. Hosts in the cluster are monitored and in the event of a failure, the virtual machines on a failed host are restarted on alternate hosts.