

US00RE48680E

(19) **United States**
(12) **Reissued Patent**
Yemini et al.

(10) **Patent Number:** **US RE48,680 E**
(45) **Date of Reissued Patent:** ***Aug. 10, 2021**

(54) **MANAGING RESOURCES IN CONTAINER SYSTEMS**

(71) Applicant: **Turbonomic, Inc.**, Boston, MA (US)

(72) Inventors: **Yechiam Yemini**, Alpine, NJ (US);
Mor Cohen, New York, NY (US);
Enlin Xu, New York, NY (US); **Andre Sara**, Briarcliff Manor, NY (US);
Shmuel Kliger, Chappaqua, NY (US)

(73) Assignee: **Turbonomic, Inc.**, Boston, MA (US)

(*) Notice: This patent is subject to a terminal disclaimer.

(21) Appl. No.: **16/940,585**

(22) Filed: **Jul. 28, 2020**

Related U.S. Patent Documents

Reissue of:

(64) Patent No.: **9,888,067**
Issued: **Feb. 6, 2018**
Appl. No.: **14/810,284**
Filed: **Jul. 27, 2015**

U.S. Applications:

(63) Continuation-in-part of application No. 14/539,973,
filed on Nov. 12, 2014, now Pat. No. 9,852,011,
(Continued)

(51) **Int. Cl.**
H04L 29/08 (2006.01)
G06F 9/50 (2006.01)
(Continued)

(52) **U.S. Cl.**
CPC **H04L 67/1023** (2013.01); **G06F 9/455**
(2013.01); **G06F 9/45558** (2013.01);
(Continued)

(58) **Field of Classification Search**

CPC ... H04L 9/00; H04L 67/1023; H04L 67/1008;
G06Q 20/085; G06Q 30/04; G06F
9/5055; G06F 9/45558; G06F 2009/4557
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,055,571 A 4/2000 Fulp et al.
6,480,861 B1 * 11/2002 Kanevsky G06F 9/50
(Continued)

FOREIGN PATENT DOCUMENTS

WO 2009091580 7/2009
WO WO 2009/091580 A1 7/2009
(Continued)

OTHER PUBLICATIONS

“Apache Hadoop NextGen MapReduce (YARN)”, The Apache
Software Foundation, Nov. 2014, 2 pages. <http://hadoop.apache.org/docs/current/hadoop-yam/hadoop-yam-site/YARN.html>.

(Continued)

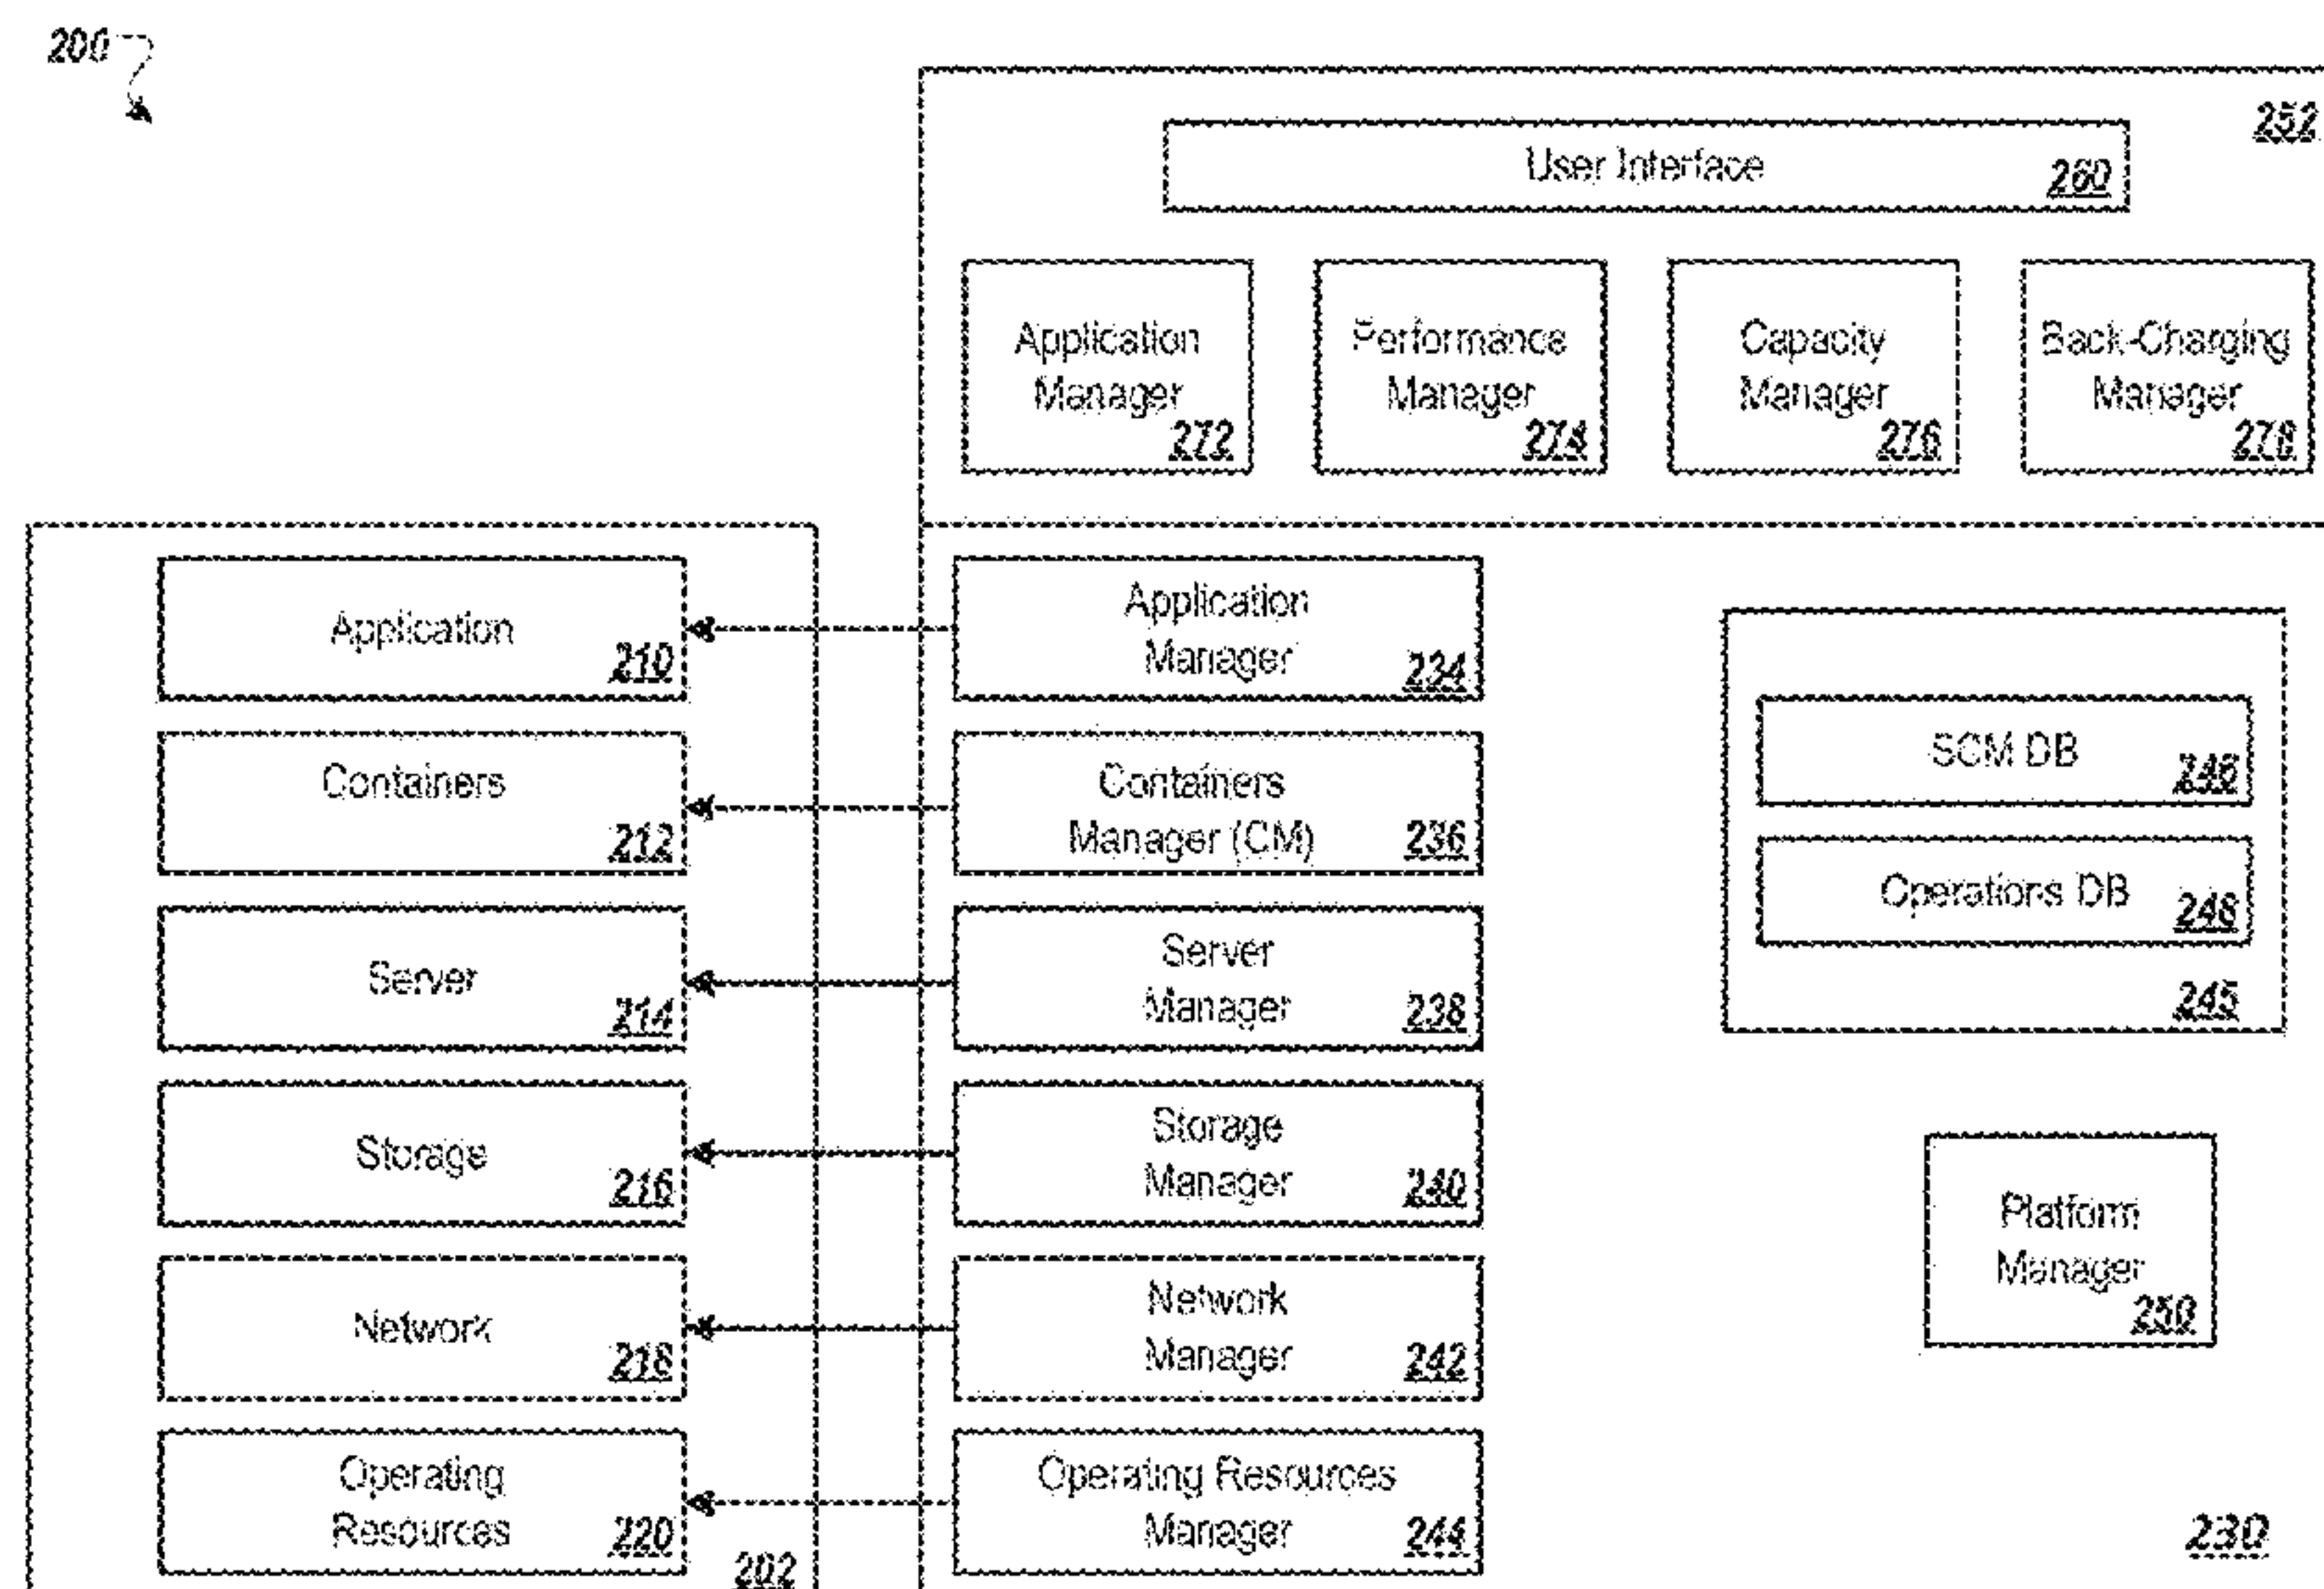
Primary Examiner — Minh Dieu Nguyen

(74) *Attorney, Agent, or Firm* — Morgan, Lewis &
Bockius LLP

(57) **ABSTRACT**

Methods, systems, and apparatus, including computer program products, for managing resources in container systems, including multi-cloud systems. The use of supply chain economics alone and in combination with other techniques offers a unified platform to integrate, optimize or improve, and automate resource management in a container system. These techniques may be used to monitor and control the delivery of service level agreements and software licenses. They may also be used to monitor and control contention of computing resources in a container system, and to suspend or terminate computing resources.

23 Claims, 18 Drawing Sheets



Related U.S. Application Data					
which is a continuation-in-part of application No. 12/492,890, filed on Jun. 26, 2009, now Pat. No. 8,914,511.			2005/0132363	A1	6/2005 Tewari et al.
			2005/0256683	A1	11/2005 Hillermeier
			2006/0020939	A1	1/2006 Fellenstein et al.
			2006/0045100	A1	3/2006 Klausberger et al.
			2006/0090163	A1 *	4/2006 Karisson G06F 11/3409 718/105
			2006/0143617	A1	6/2006 Knauerhase et al.
			2006/0167984	A1	7/2006 Fellenstein et al.
			2006/0168584	A1	7/2006 Dawson et al.
			2006/0184937	A1	8/2006 Abels et al.
			2006/0188011	A1	8/2006 Goldszmidt et al.
			2006/0190482	A1	8/2006 Kishan et al.
			2006/0212332	A1	9/2006 Jackson
			2006/0244607	A1	11/2006 Liu et al.
			2006/0265470	A1 *	11/2006 Rolia G06F 9/5027 709/217
			2007/0105630	A1 *	5/2007 Van Luchene G06Q 30/04 463/42
			2008/0052387	A1	2/2008 Heinz et al.
			2008/0080552	A1	4/2008 Gates et al.
			2008/0083031	A1	4/2008 Meijer et al.
			2008/0104608	A1	5/2008 Hyser et al.
			2008/0109241	A1	5/2008 D'Alto et al.
			2008/0127348	A1	5/2008 Largman et al.
			2008/0154837	A1	6/2008 Morimura et al.
			2008/0155169	A1	6/2008 Hiltgen et al.
			2008/0163194	A1	7/2008 Dias et al.
			2008/0201409	A1	8/2008 Vul et al.
			2008/0244607	A1 *	10/2008 Rysin G06Q 30/06 718/104
			2008/0250143	A1	10/2008 Garg et al.
			2008/0301027	A1	12/2008 Boss et al.
			2008/0319926	A1	12/2008 Alam et al.
			2008/0320474	A1 *	12/2008 Jelinek G06F 9/455 718/1
			2009/0007126	A1 *	1/2009 Jelinek G06F 9/5016 718/104
			2009/0052333	A1	2/2009 Vankov et al.
			2009/0089406	A1 *	4/2009 Roush G06F 9/5077 709/220
			2009/0164356	A1	6/2009 Bakman
			2009/0172666	A1 *	7/2009 Yahalom G06F 3/0605 718/1
			2009/0182565	A1	7/2009 Erickson et al.
			2009/0240628	A1	9/2009 Blair et al.
			2009/0249488	A1 *	10/2009 Robinson G06F 21/10 726/26
			2009/0276271	A1	11/2009 Munson et al.
			2009/0276771	A1	11/2009 Nickolov et al.
			2010/0005173	A1 *	1/2010 Baskaran G06F 15/16 709/226
			2010/0057625	A1	3/2010 Boss et al.
			2010/0077449	A1	3/2010 Kwok et al.
			2010/0088205	A1	4/2010 Robertson
			2010/0106820	A1	4/2010 Gulati et al.
			2010/0142522	A1	6/2010 Gardner
			2010/0153945	A1 *	6/2010 Bansal G06F 9/4881 718/1
			2010/0180275	A1 *	7/2010 Neogi G06F 1/3203 718/1
			2010/0318454	A1	12/2010 Warncke et al.
			2010/0319004	A1	12/2010 Hudson et al.
			2011/0131335	A1	6/2011 Spaltro et al.
			2011/0225277	A1	9/2011 Freimuth et al.
			2011/0276951	A1	11/2011 Jain
			2012/0072762	A1 *	3/2012 Atchison G06F 9/5072 714/2
			2012/0131591	A1	5/2012 Moorthi et al.
			2012/0226796	A1	9/2012 Morgan
			2012/0317578	A1 *	12/2012 Kansal G06F 9/5077 718/104
			2014/0081766	A1	3/2014 Maass
(60) Provisional application No. 62/077,898, filed on Nov. 10, 2014, provisional application No. 62/077,852, filed on Nov. 10, 2014, provisional application No. 62/077,860, filed on Nov. 10, 2014, provisional application No. 62/159,883, filed on May 11, 2015.					
(51) Int. Cl.					
G06F 9/455 (2018.01)					
G06Q 10/06 (2012.01)					
G06Q 30/06 (2012.01)					
G06Q 40/02 (2012.01)					
(52) U.S. Cl.					
CPC G06F 9/5055 (2013.01); G06F 9/5077 (2013.01); G06F 9/5083 (2013.01); G06Q 10/06 (2013.01); G06Q 30/06 (2013.01); H04L 67/1008 (2013.01); G06F 2209/5022 (2013.01); G06Q 40/02 (2013.01)					
(56) References Cited					
U.S. PATENT DOCUMENTS					
6,735,553	B1	5/2004	Frogner et al.		
7,013,296	B1	3/2006	Yemini et al.		
7,140,039	B1	11/2006	Yemini et al.		
7,272,855	B1	9/2007	Yemini et al.		
7,441,023	B2	10/2008	Benjamin et al.		
7,464,132	B1	12/2008	Florissi et al.		
7,496,655	B2	2/2009	Gopalan et al.		
7,590,601	B2	9/2009	Shea et al.		
7,685,270	B1 *	3/2010	Vermeulen G06Q 30/0603 370/252		
7,711,789	B1	5/2010	Jnagal et al.		
7,774,457	B1 *	8/2010	Talwar G06F 9/5072 709/203		
7,840,517	B2	11/2010	Morimura et al.		
7,877,754	B2	1/2011	Birkestrand et al.		
8,051,017	B2	11/2011	Munson et al.		
8,060,875	B1 *	11/2011	Lambeth G06F 9/5077 709/223		
8,181,175	B1 *	5/2012	McKee G06F 11/3419 718/1		
8,271,345	B1	9/2012	Milgrom et al.		
8,286,174	B1	10/2012	Schmidt et al.		
8,347,302	B1	1/2013	Vincent et al.		
8,370,898	B1	2/2013	Jackson		
8,396,807	B1	3/2013	Yemini et al.		
8,429,097	B1 *	4/2013	Sivasubramanian G06F 17/30483 706/12		
8,433,801	B1	4/2013	Yemini et al.		
8,478,878	B2	7/2013	Freimuth et al.		
8,650,296	B1 *	2/2014	Herington G06F 9/5077 709/226		
8,661,131	B1	2/2014	Yemini et al.		
8,762,531	B1	6/2014	Yemini et al.		
8,914,511	B1	12/2014	Yemini et al.		
8,959,249	B1 *	2/2015	Love G06F 3/0611 710/6		
9,589,495	B2	3/2017	Tsuei et al.		
9,589,501	B2	3/2017	Hirama		
2002/0147611	A1 *	10/2002	Greene G06Q 10/063112 705/1.1		
2003/0154123	A1	8/2003	Subbloie et al.		
2005/0005271	A1	1/2005	Clymer et al.		

(56)

References Cited

U.S. PATENT DOCUMENTS

2014/0376368 A1 12/2014 Xiao et al.
 2014/0380324 A1 12/2014 Xiao et al.

FOREIGN PATENT DOCUMENTS

WO 2010036731 4/2010
 WO WO 2010/036731 A2 4/2010

OTHER PUBLICATIONS

Kubemetes Design Overview:, Github, GoogleCloudPlatform, Oct. 2014, 4 pages. <https://github.com/GoogleCloudPlatform/kubernetes/blob/master/DESIGN.md>.

Dailianas et al., "Use of Currency for Access Control in Large-Scale Information Systems", Ph.D. Thesis Proposal, Columbia University, Department of Computer Science, Sep. 1998, 33 pages.

Erickson et al., "Using Network Knowledge to Improve Performance in Virtualized Data Centers", PhD thesis, Stanford University, May 2013, 10 pages.

Ferguson et al., "Economic Models for Allocating Resources in Computer Systems", Market Based Control of Distributed Systems. World Scientific, 2006, 26 pages.

Sutherland et al., "A Futures Market in Computer Time", Aiken Computation Lab, vol. 11, No. 6, Jun. 1968, 3 pages.

Turner et al., "A Lightweight Current Paradigm for the P2P Resource Market", Electronic Commerce Research—ECR, 2003, 15 pages.

U.S. Appl. No. 09/589,501, filed Jun. 7, 2000, Yemini et al.

U.S. Appl. No. 09/589,495, filed Jun. 7, 2000, Yemini et al.

"Big Ideas: Demystifying Hadoop", EMC, YouTube Video, Published Jul. 2012. <https://www.youtube.com/watch?v=xJHv5t8jcM8>.

"MapReduce Tutorial", The Apache Software Foundation, Nov. 2014, 24 pages. <http://hadoop.apache.org/docs/current/hadoop-mapreduce-client/hadoop-mapreduce-client-core/MapReduceTutorial.html>.

Byde et al., "Market-Based Resource Allocation for Utility Data Centers", HP Laboratories Bristol, Sep. 2003, 16 pages.

Dailianas et al., "MarketNet: Market-based Protection of Network Systems and Services—an Application to SNMP Protection", INFOCOM 2000, Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies, IEEE, vol. 3, Mar. 2000, 10 pages.

Dias et al., "Online Traffic-aware Virtual Machine Placement in Data Center Networks", Global Information Infrastructure and Networking Symposium (GHS), 2012, 8 pages.

Ferguson et al., "Economic Models for Allocating Resources in Computer Systems", Market Based Control of Distributed Systems. World Scientific, 2006, 26 pages.

Ferguson et al., "Microeconomic Algorithms for Load Balancing in Distributed Computer Systems", IBM T.J. Watson Research Center, Nov. 2008, 9 pages.

Gao et al., Economics-Inspired Decentralized Control Approach for Adaptive Grid Services and Applications, International Journal of Intelligent Systems, vol. 21, 2006, 20 pages.

Huberman et al., "Distributed Computation as an Economic System", Journal of Economic Perspectives, vol. 9, No. 1, 1995, 12 pages.

Meng et al., "Improving the Scalability of Data Center Networks with Traffic-aware Virtual Machine Placement", INFOCOM, 2010 Proceedings IEEE, Mar. 2010, 9 pages.

Sandholm, "Managing Service Levels in Grid Computing Systems: Quota Policy and Computation; Computational Market Approaches", KTH Computer Science and Communication, Licentiate Thesis, Stockholm, Sweden, 2007, 105 pages.

Sarangan et al., "Comparative Study of Protocols for Dynamic Service Negotiation in the Next-Generation Internet", IEEE Communications Magazine, vol. 44, No. 3, Mar. 2006, 6 pages.

Stonebraker et al., "An Economic Paradigm for Query Processing and Data Migration in Mariposa", IEEE Xplore, 1994, 10 pages.

Turner et al., "A Lightweight Currency Paradigm for the P2P Resource Market", Electronic Commerce Research—ECR, 2003, 15 pages.

Waldspurger et al., "Spawn: A Distributed Computational Economy", IEEE Transactions on Software Engineering, vol. 18, No. 2, Feb. 1992, 15 pages.

Wang et al., "Pricing Network Resources for Adaptive Applications", IEEE/ACM Transactions on Networking, vol. 14, No. 3, Jun. 2006, 14 pages.

Yemini et al., "MarketNet: Market-Based Protection of Information Systems", ICE '98 Proceedings of the first international conference on Information and computation economics, 1998, 12 pages.

Yemini et al., "MARKETNET: A Market-Based Architecture for Survivable Large-Scale Information Systems", Columbia University, Department of Computer Science, Aug. 2002, 6 pages.

Yemini et al., "MarketNet: A Survivable, Market-Based Architecture for Large-Scale Information Systems", Final Technical Report, Air Force Research Laboratory, Aug. 2002, 27 pages.


Yemini et al., "MarketNet: Protecting Access to Information Systems Through Financial Market Controls", Decision Support Systems—Special Issue on Information and Computational Economics, vol. 28, No. 1-2, Mar. 2000, 20 pages.

Yemini et al., "MarketNet: Using Virtual Currency to Protect Information Systems", Research and Advanced Technology for Digital Libraries, Second European Conference, ECDL '98, 1998, 12 pages.

Yuksel et al., "Effect of Pricing Intervals on Congestion-Sensitivity of Network Prices", Telecommunication Systems, vol. 28, No. 1, 2005, 21 pages.

* cited by examiner

FIG. 1



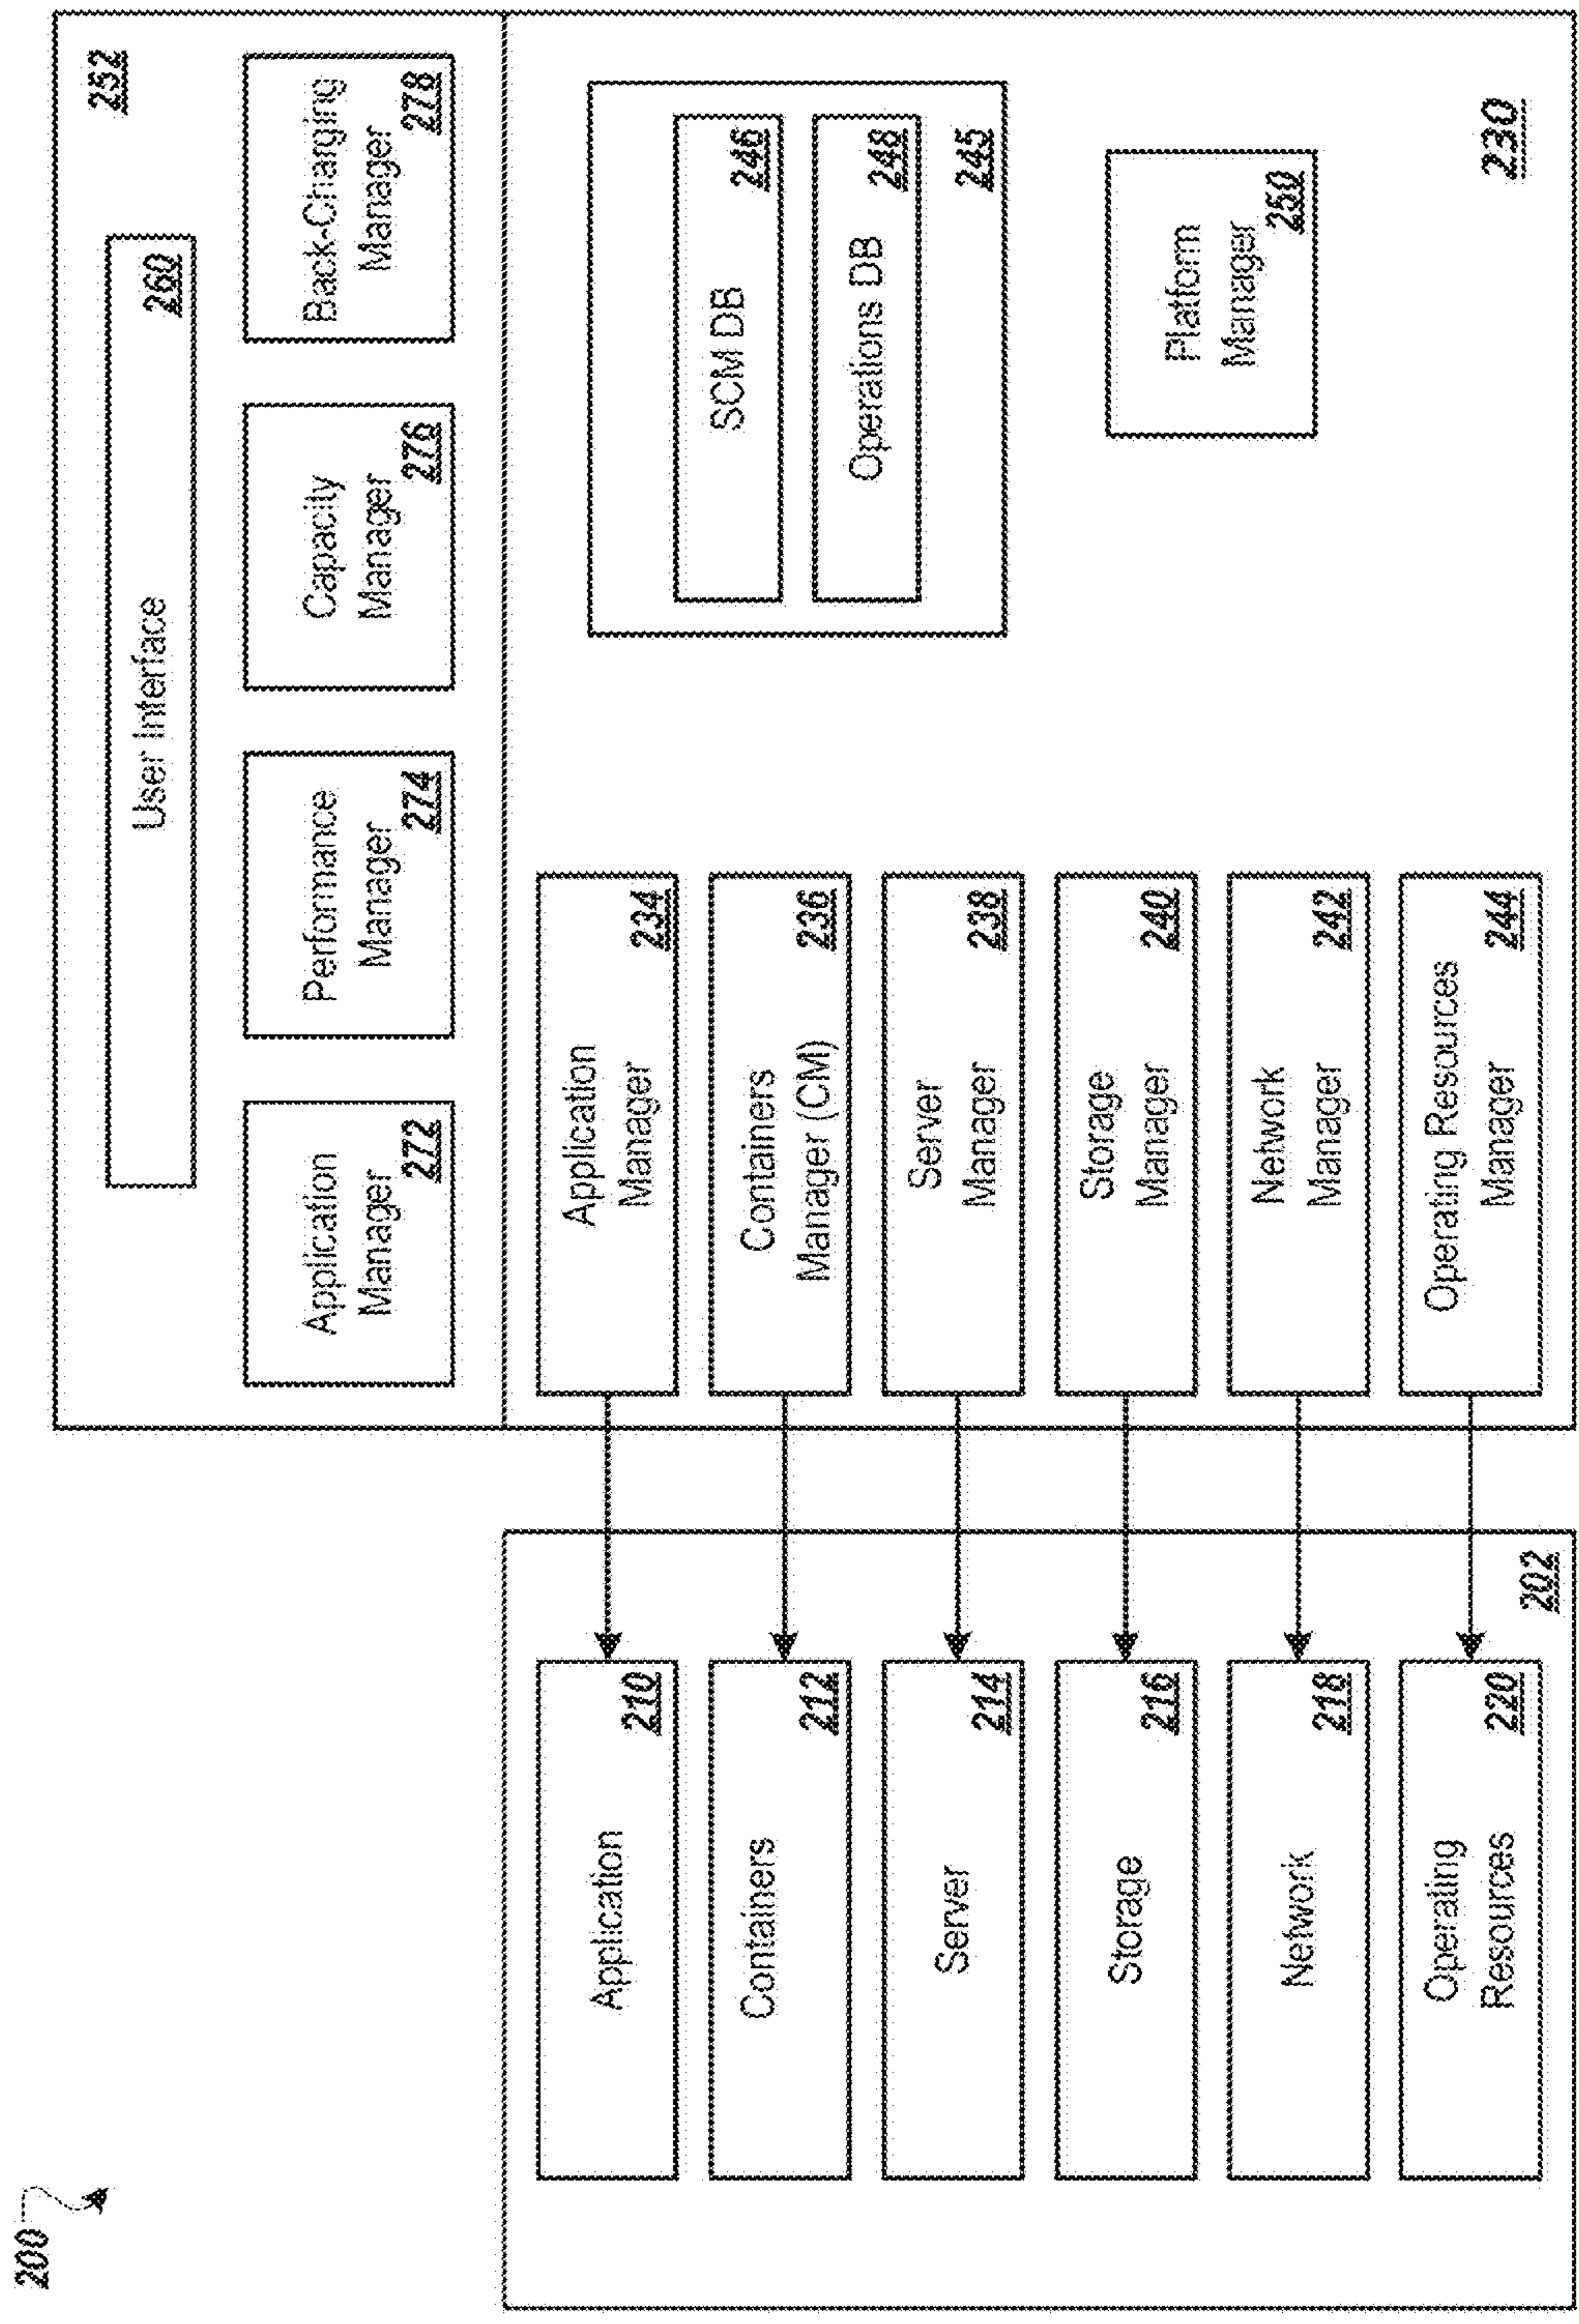



FIG. 2






FIG. 4

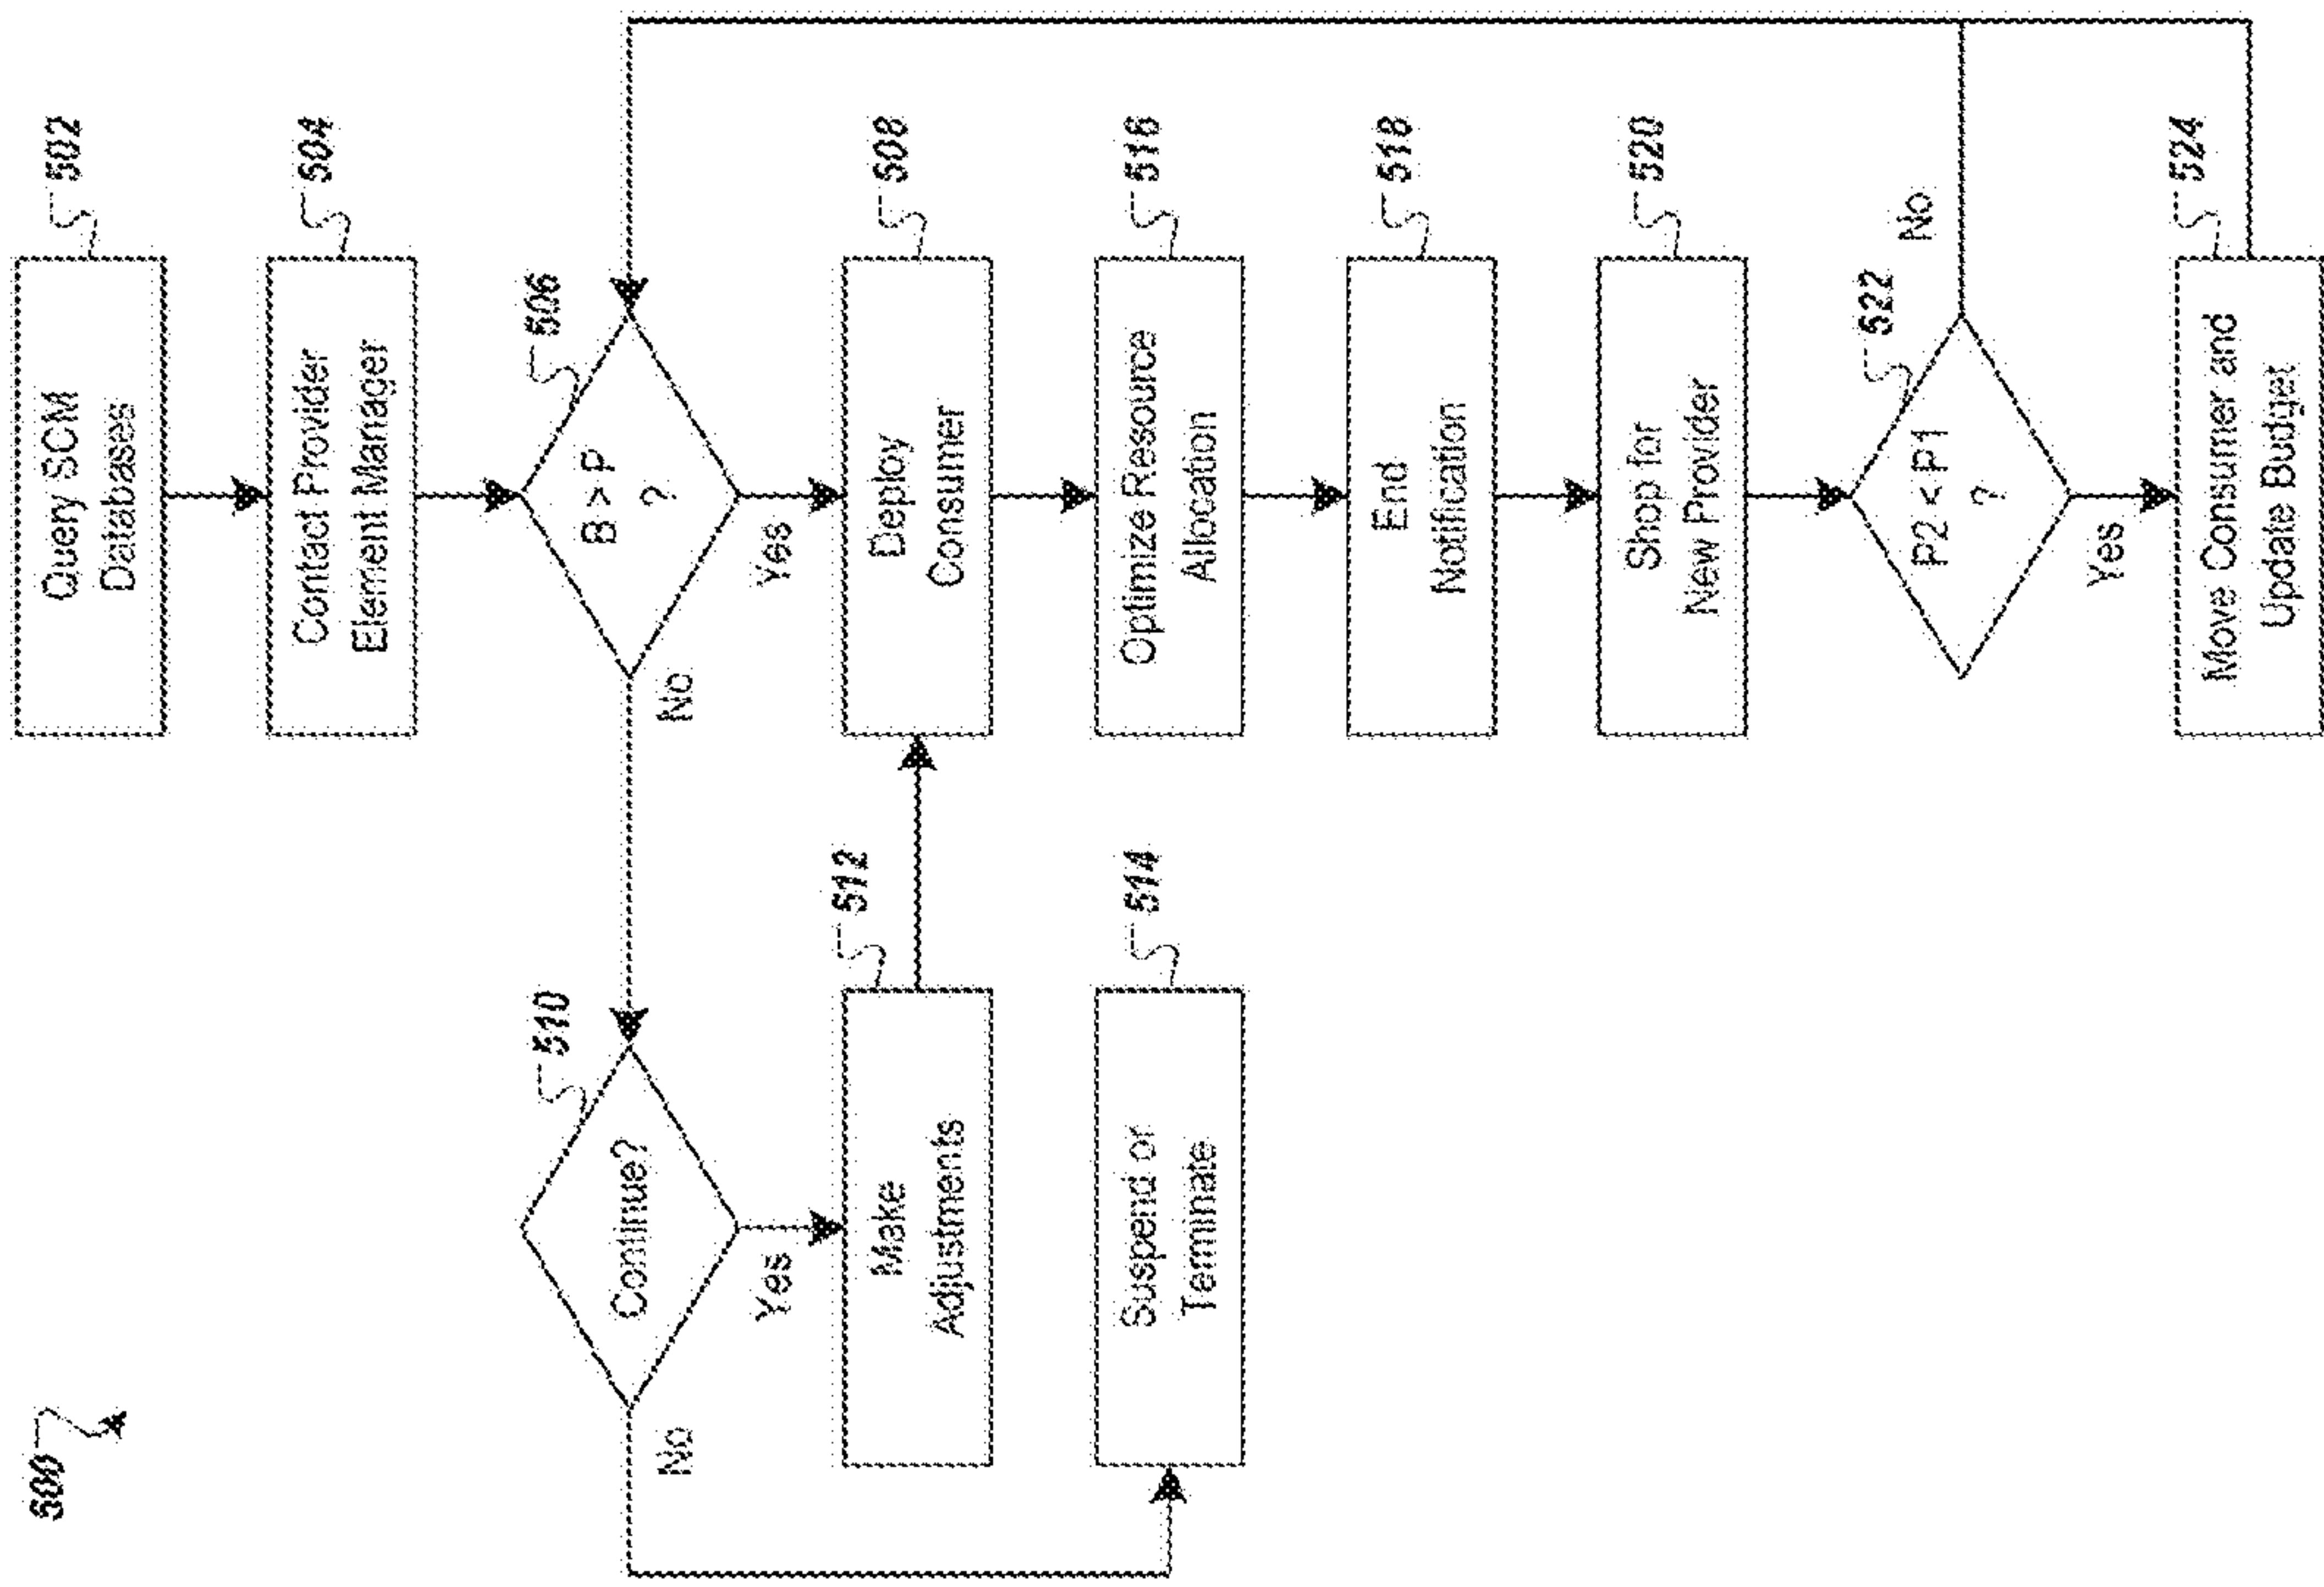



FIG. 5



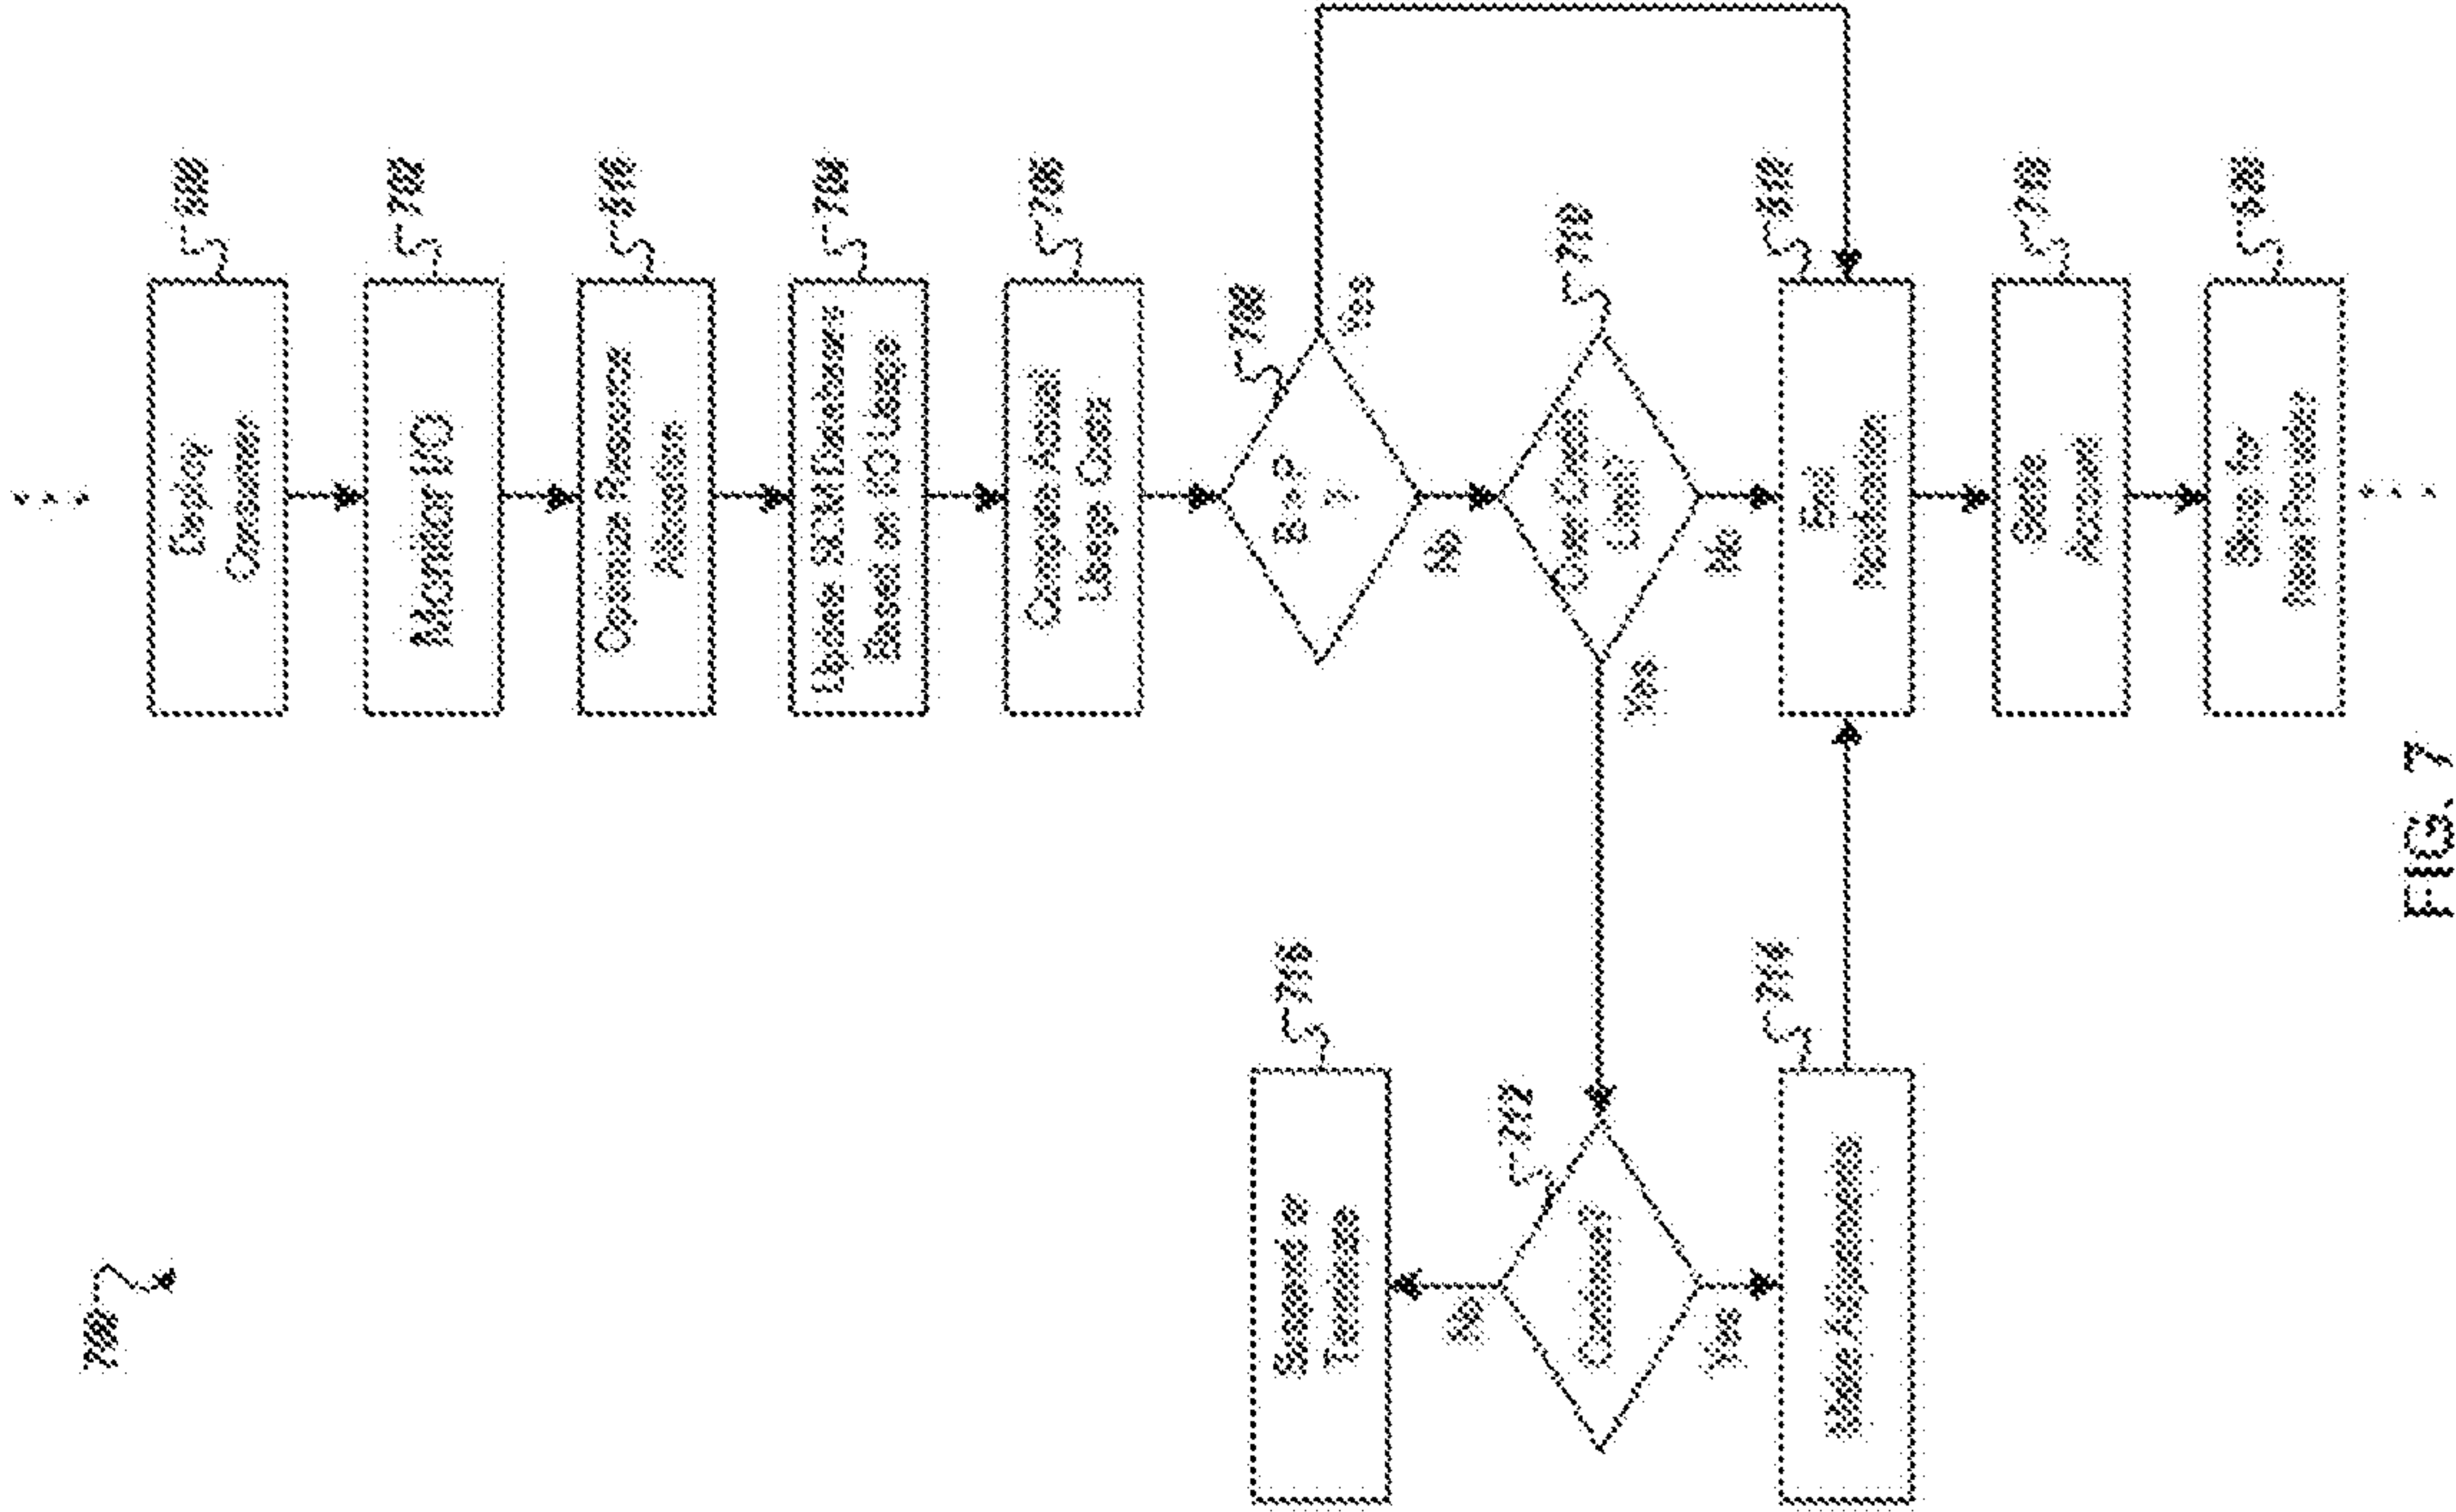


FIG. 7

800

Source-ID 802	Provider-ID 804	Transaction-ID 806	Service 808	Payment 810	Authentication 812
------------------	--------------------	-----------------------	----------------	----------------	-----------------------

FIG. 8A

850

Source-ID 852	Provider-ID 854	Transaction-ID 856	Service Confirmation 858	Authentication 860
------------------	--------------------	-----------------------	--------------------------------	-----------------------

FIG. 8B

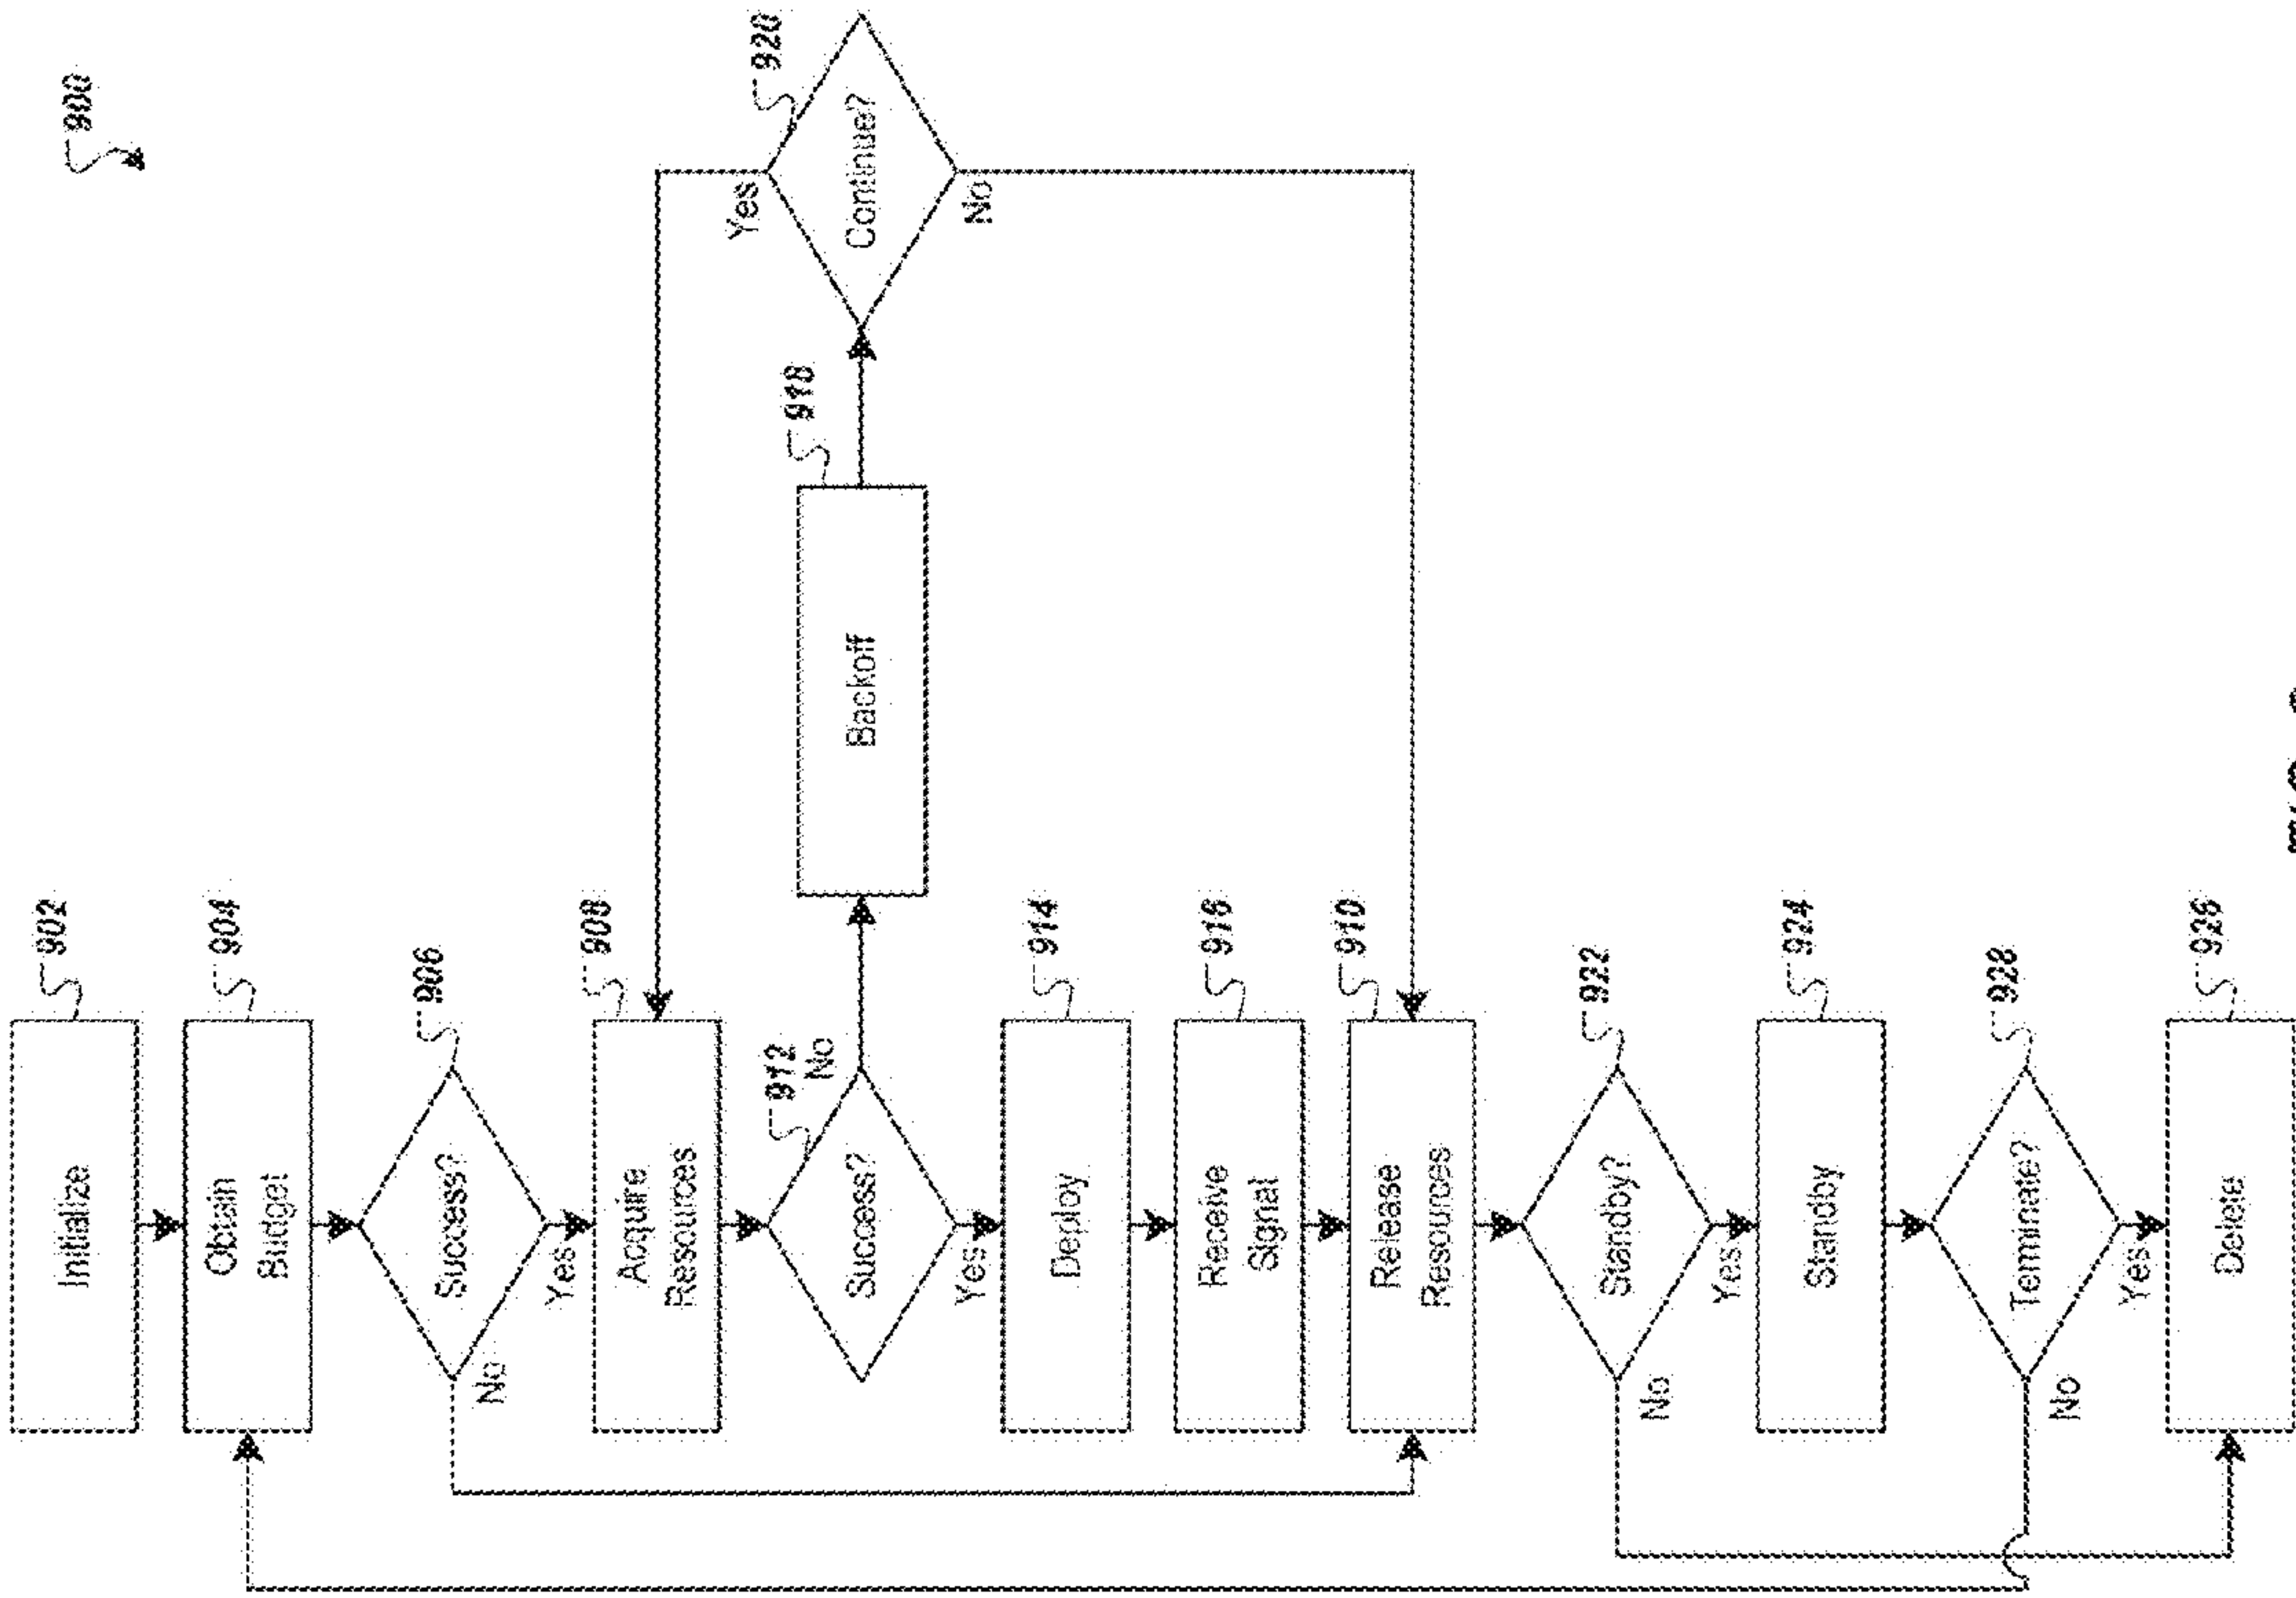


FIG. 9


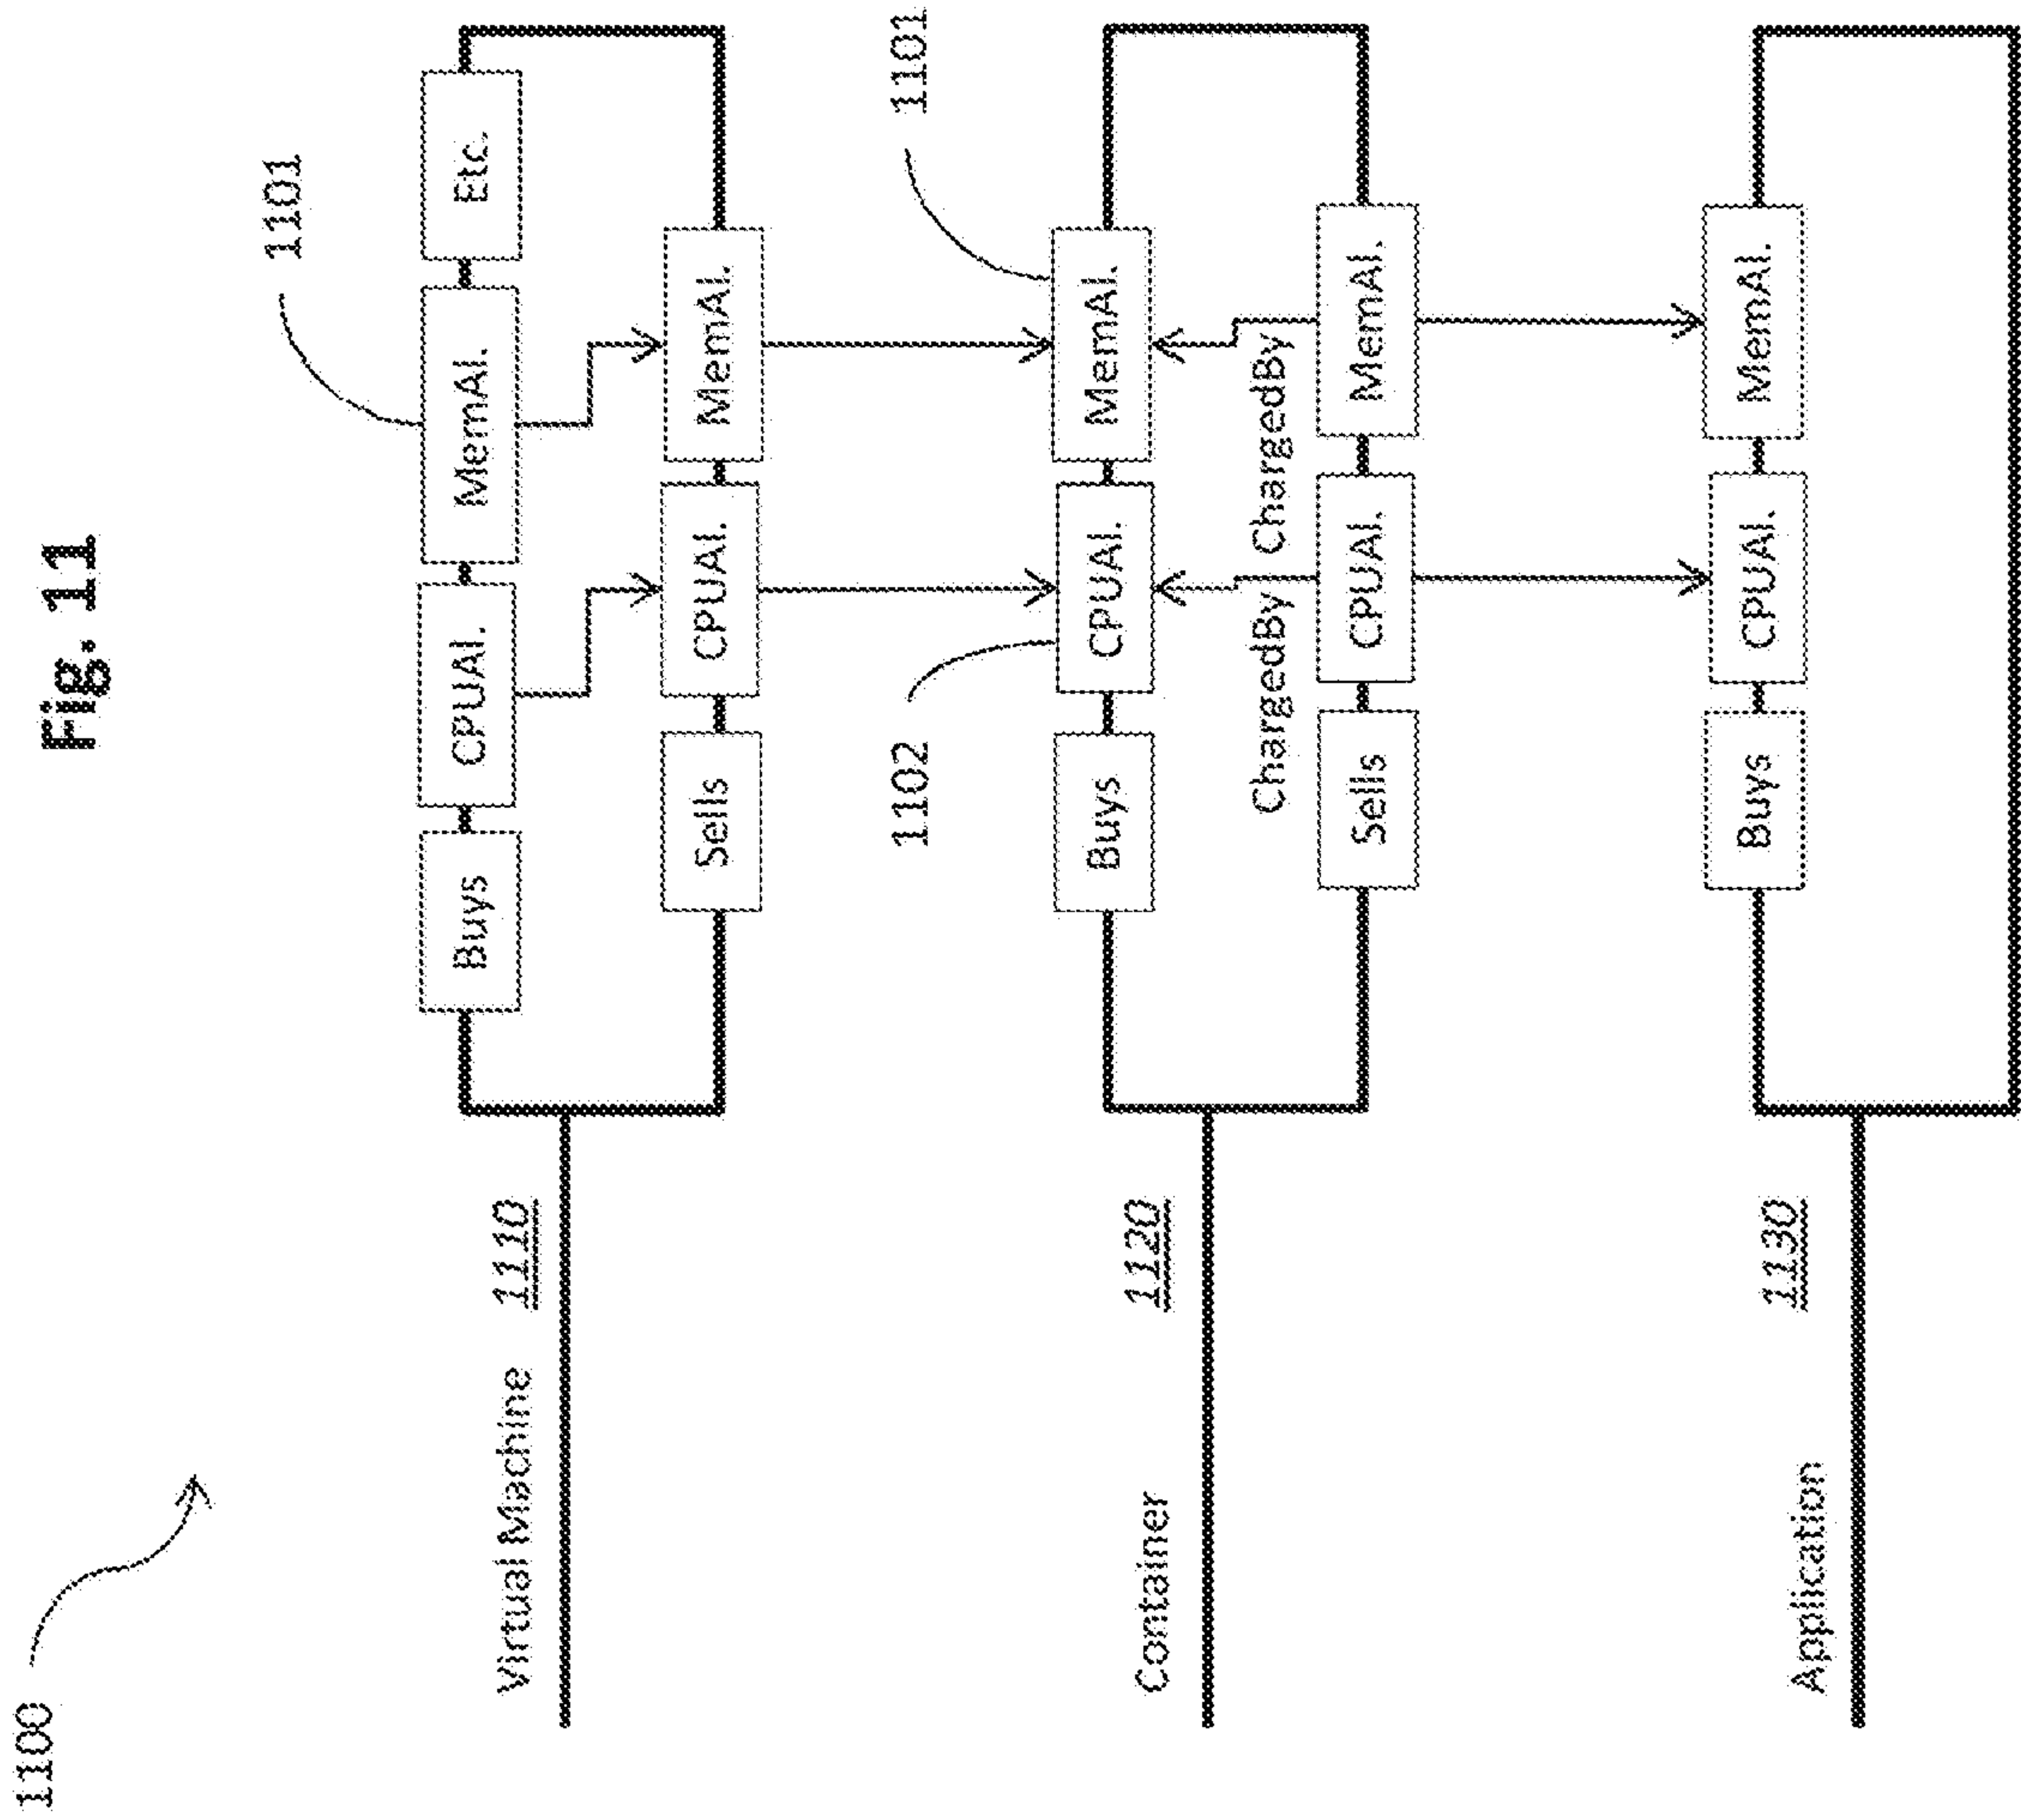
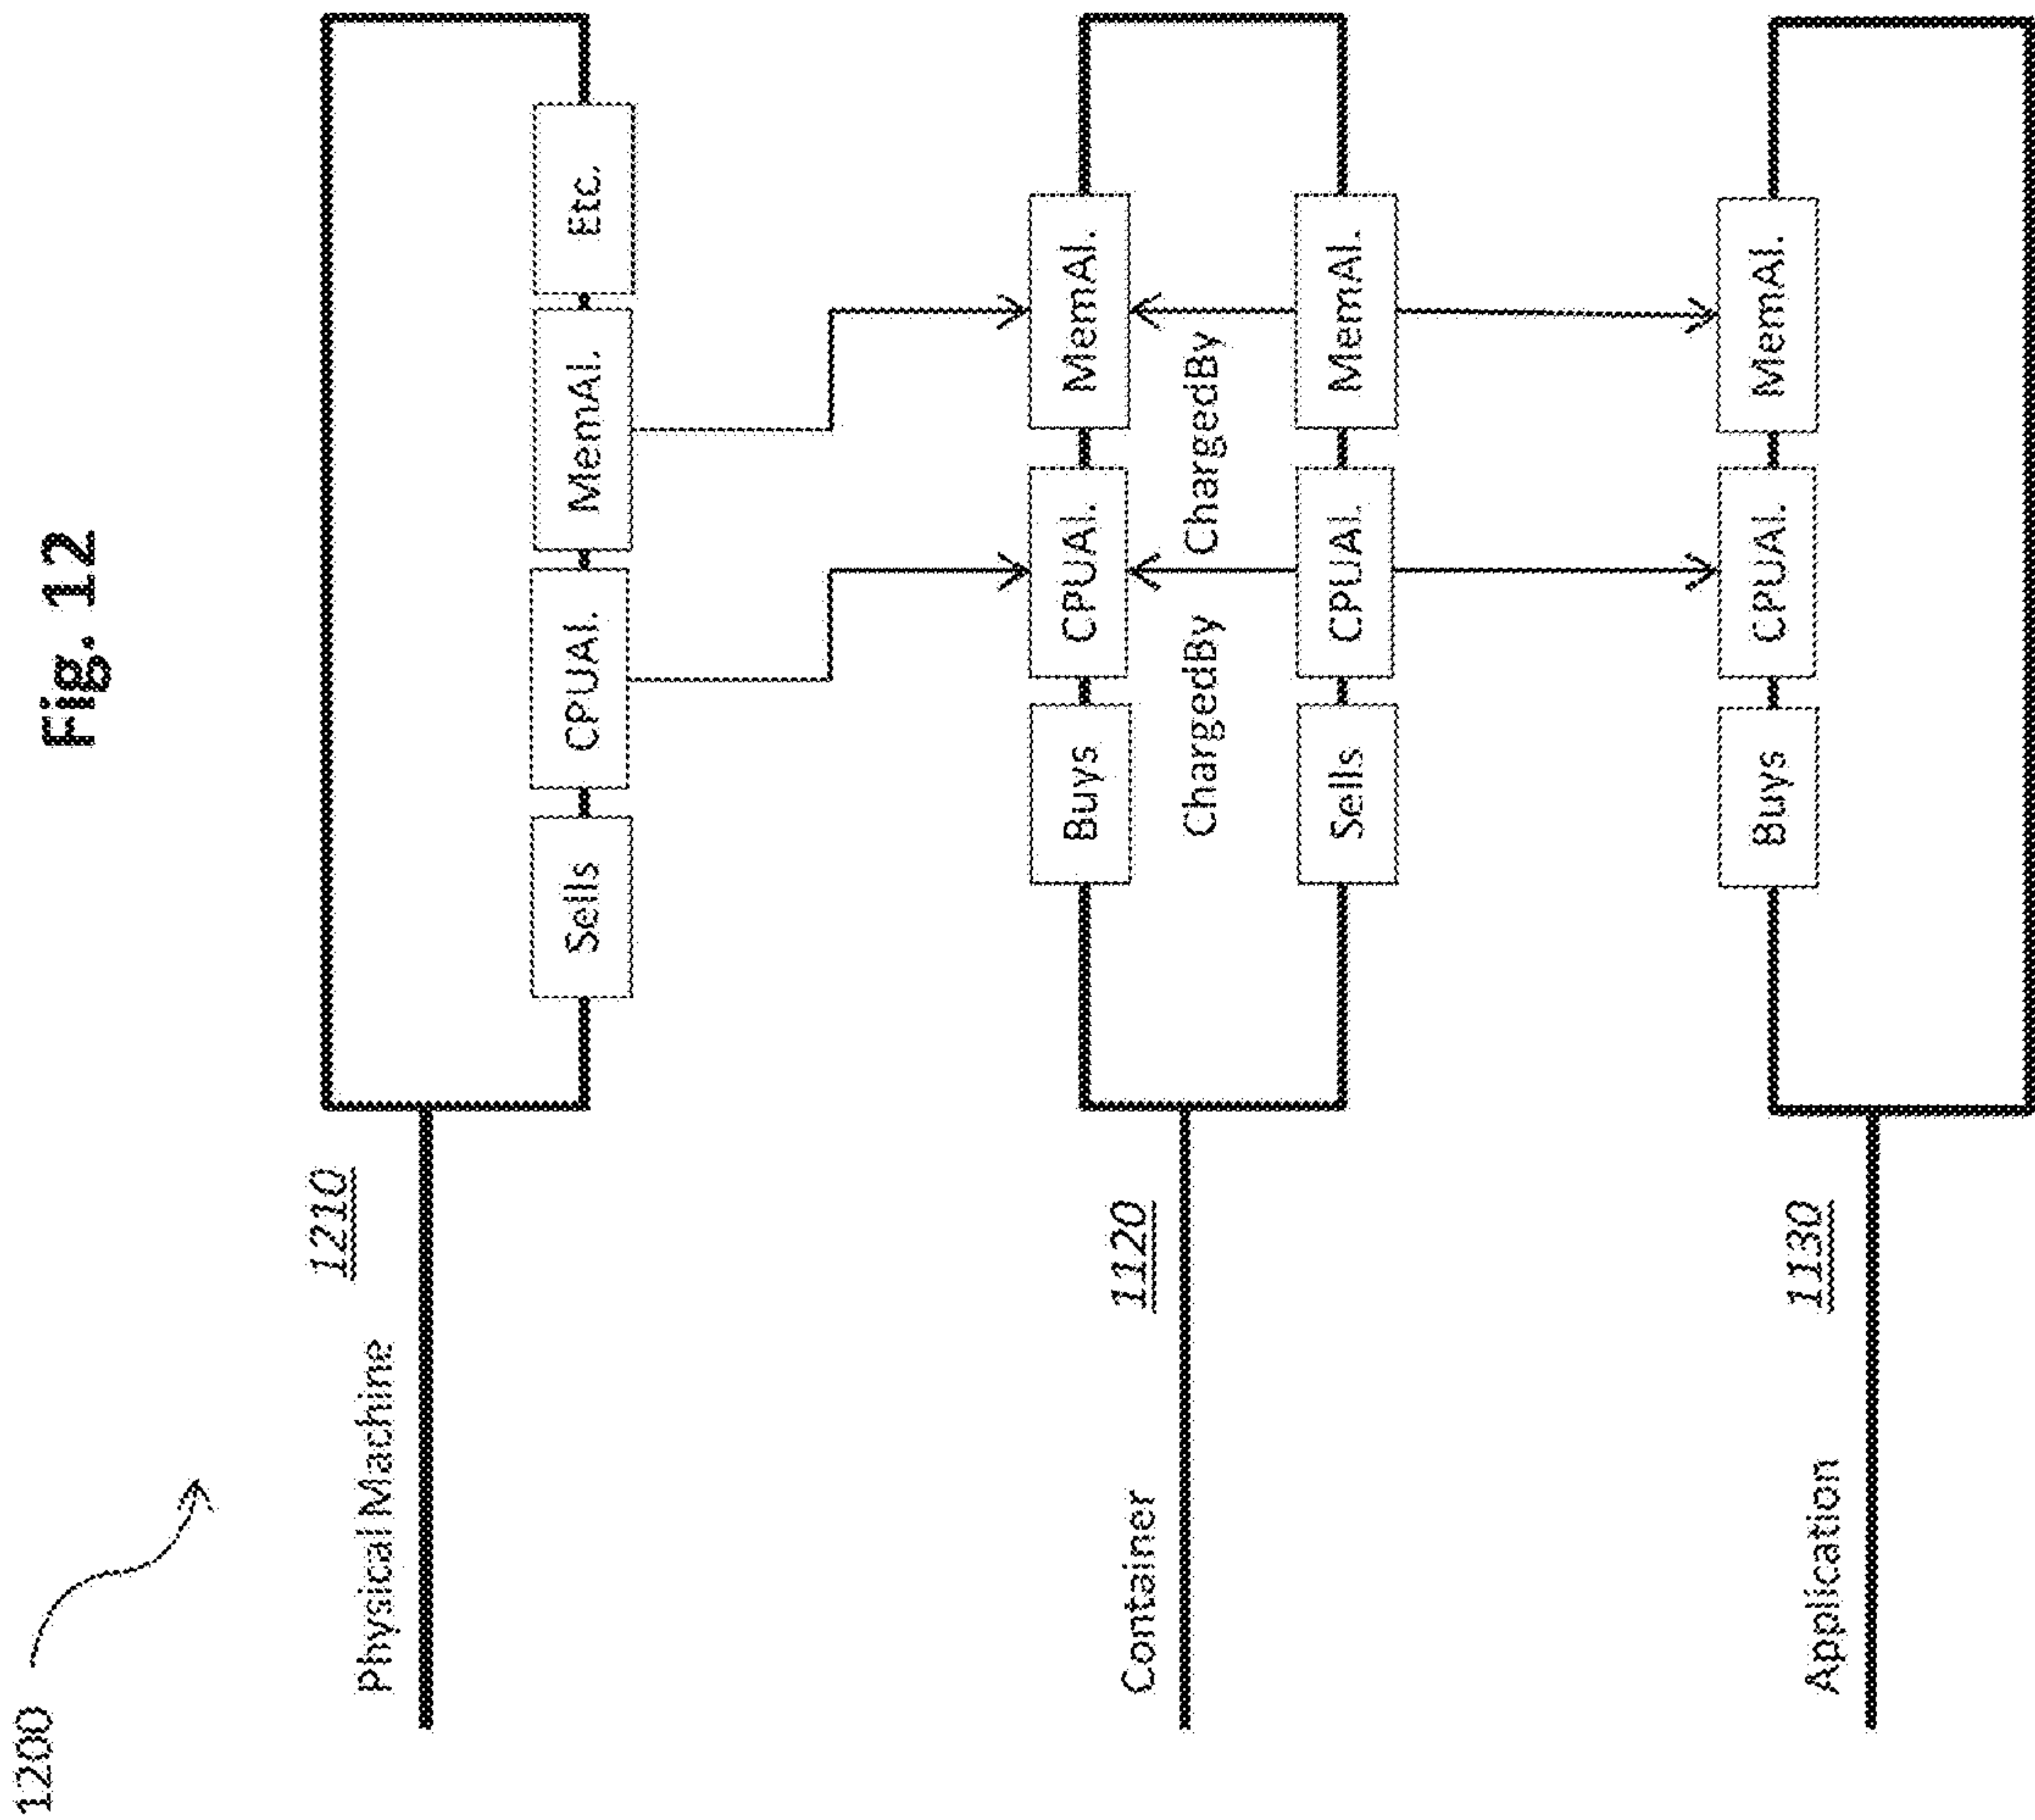


FIG. 10





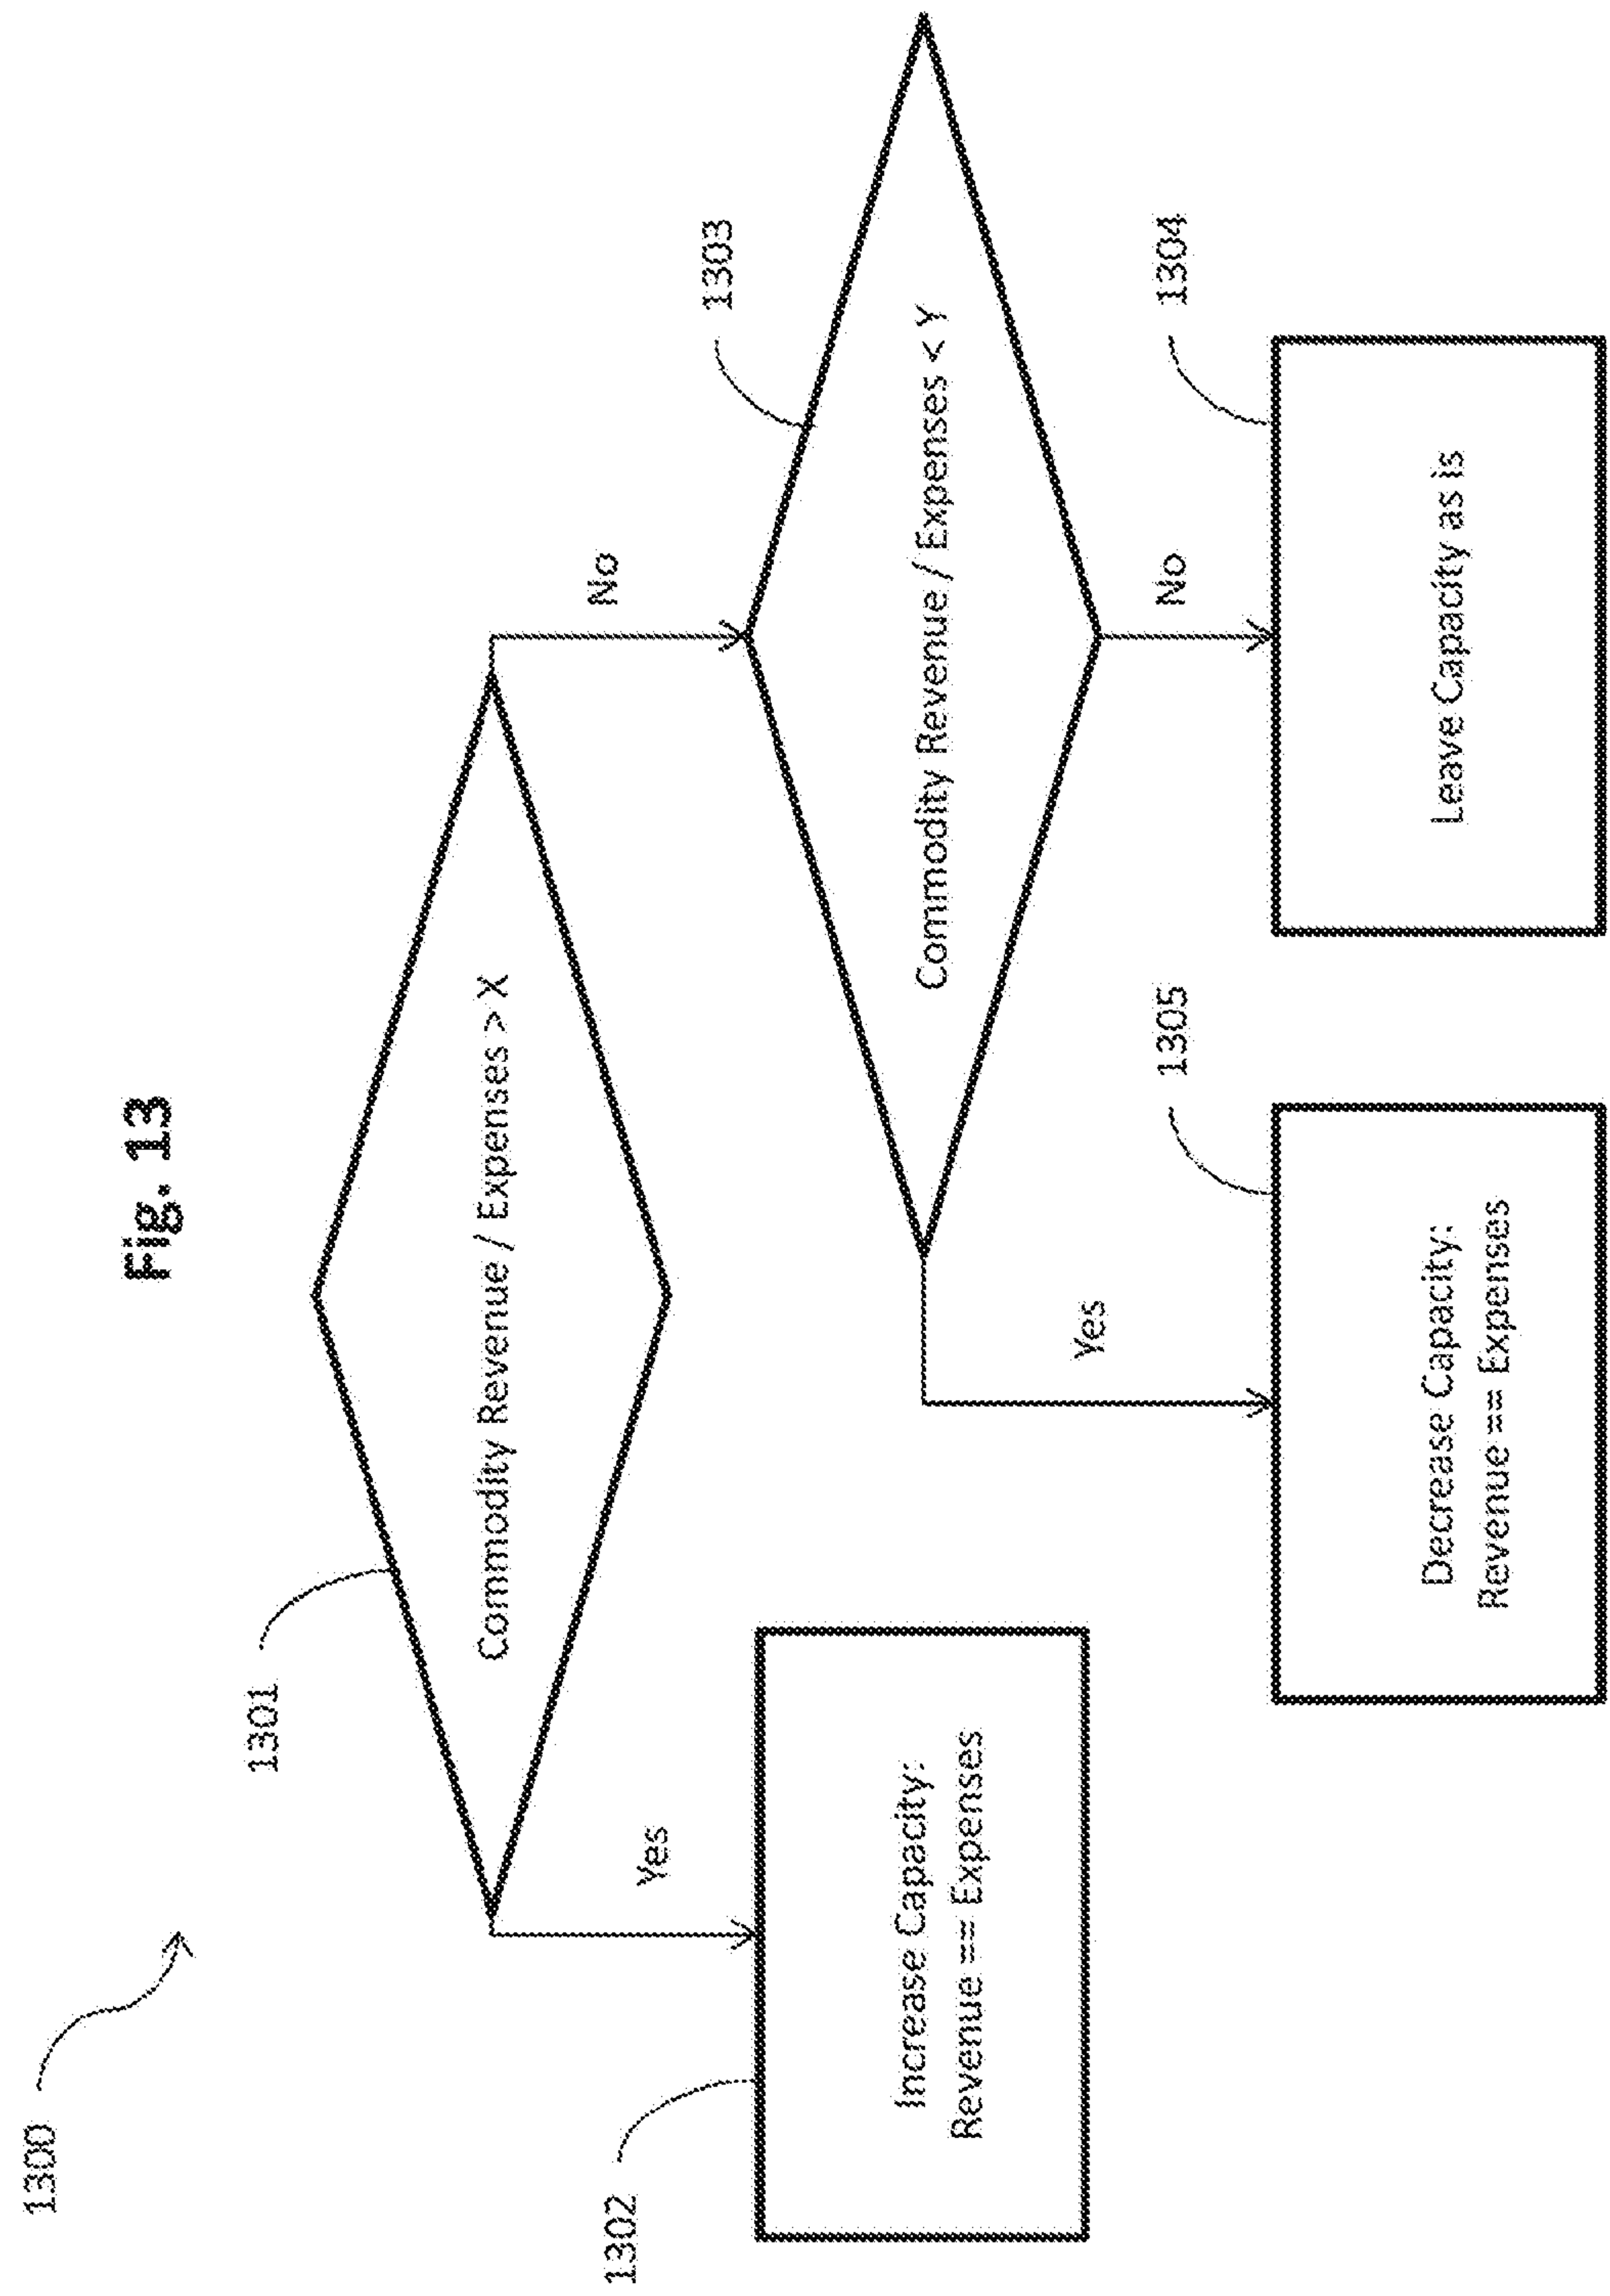


Fig. 14

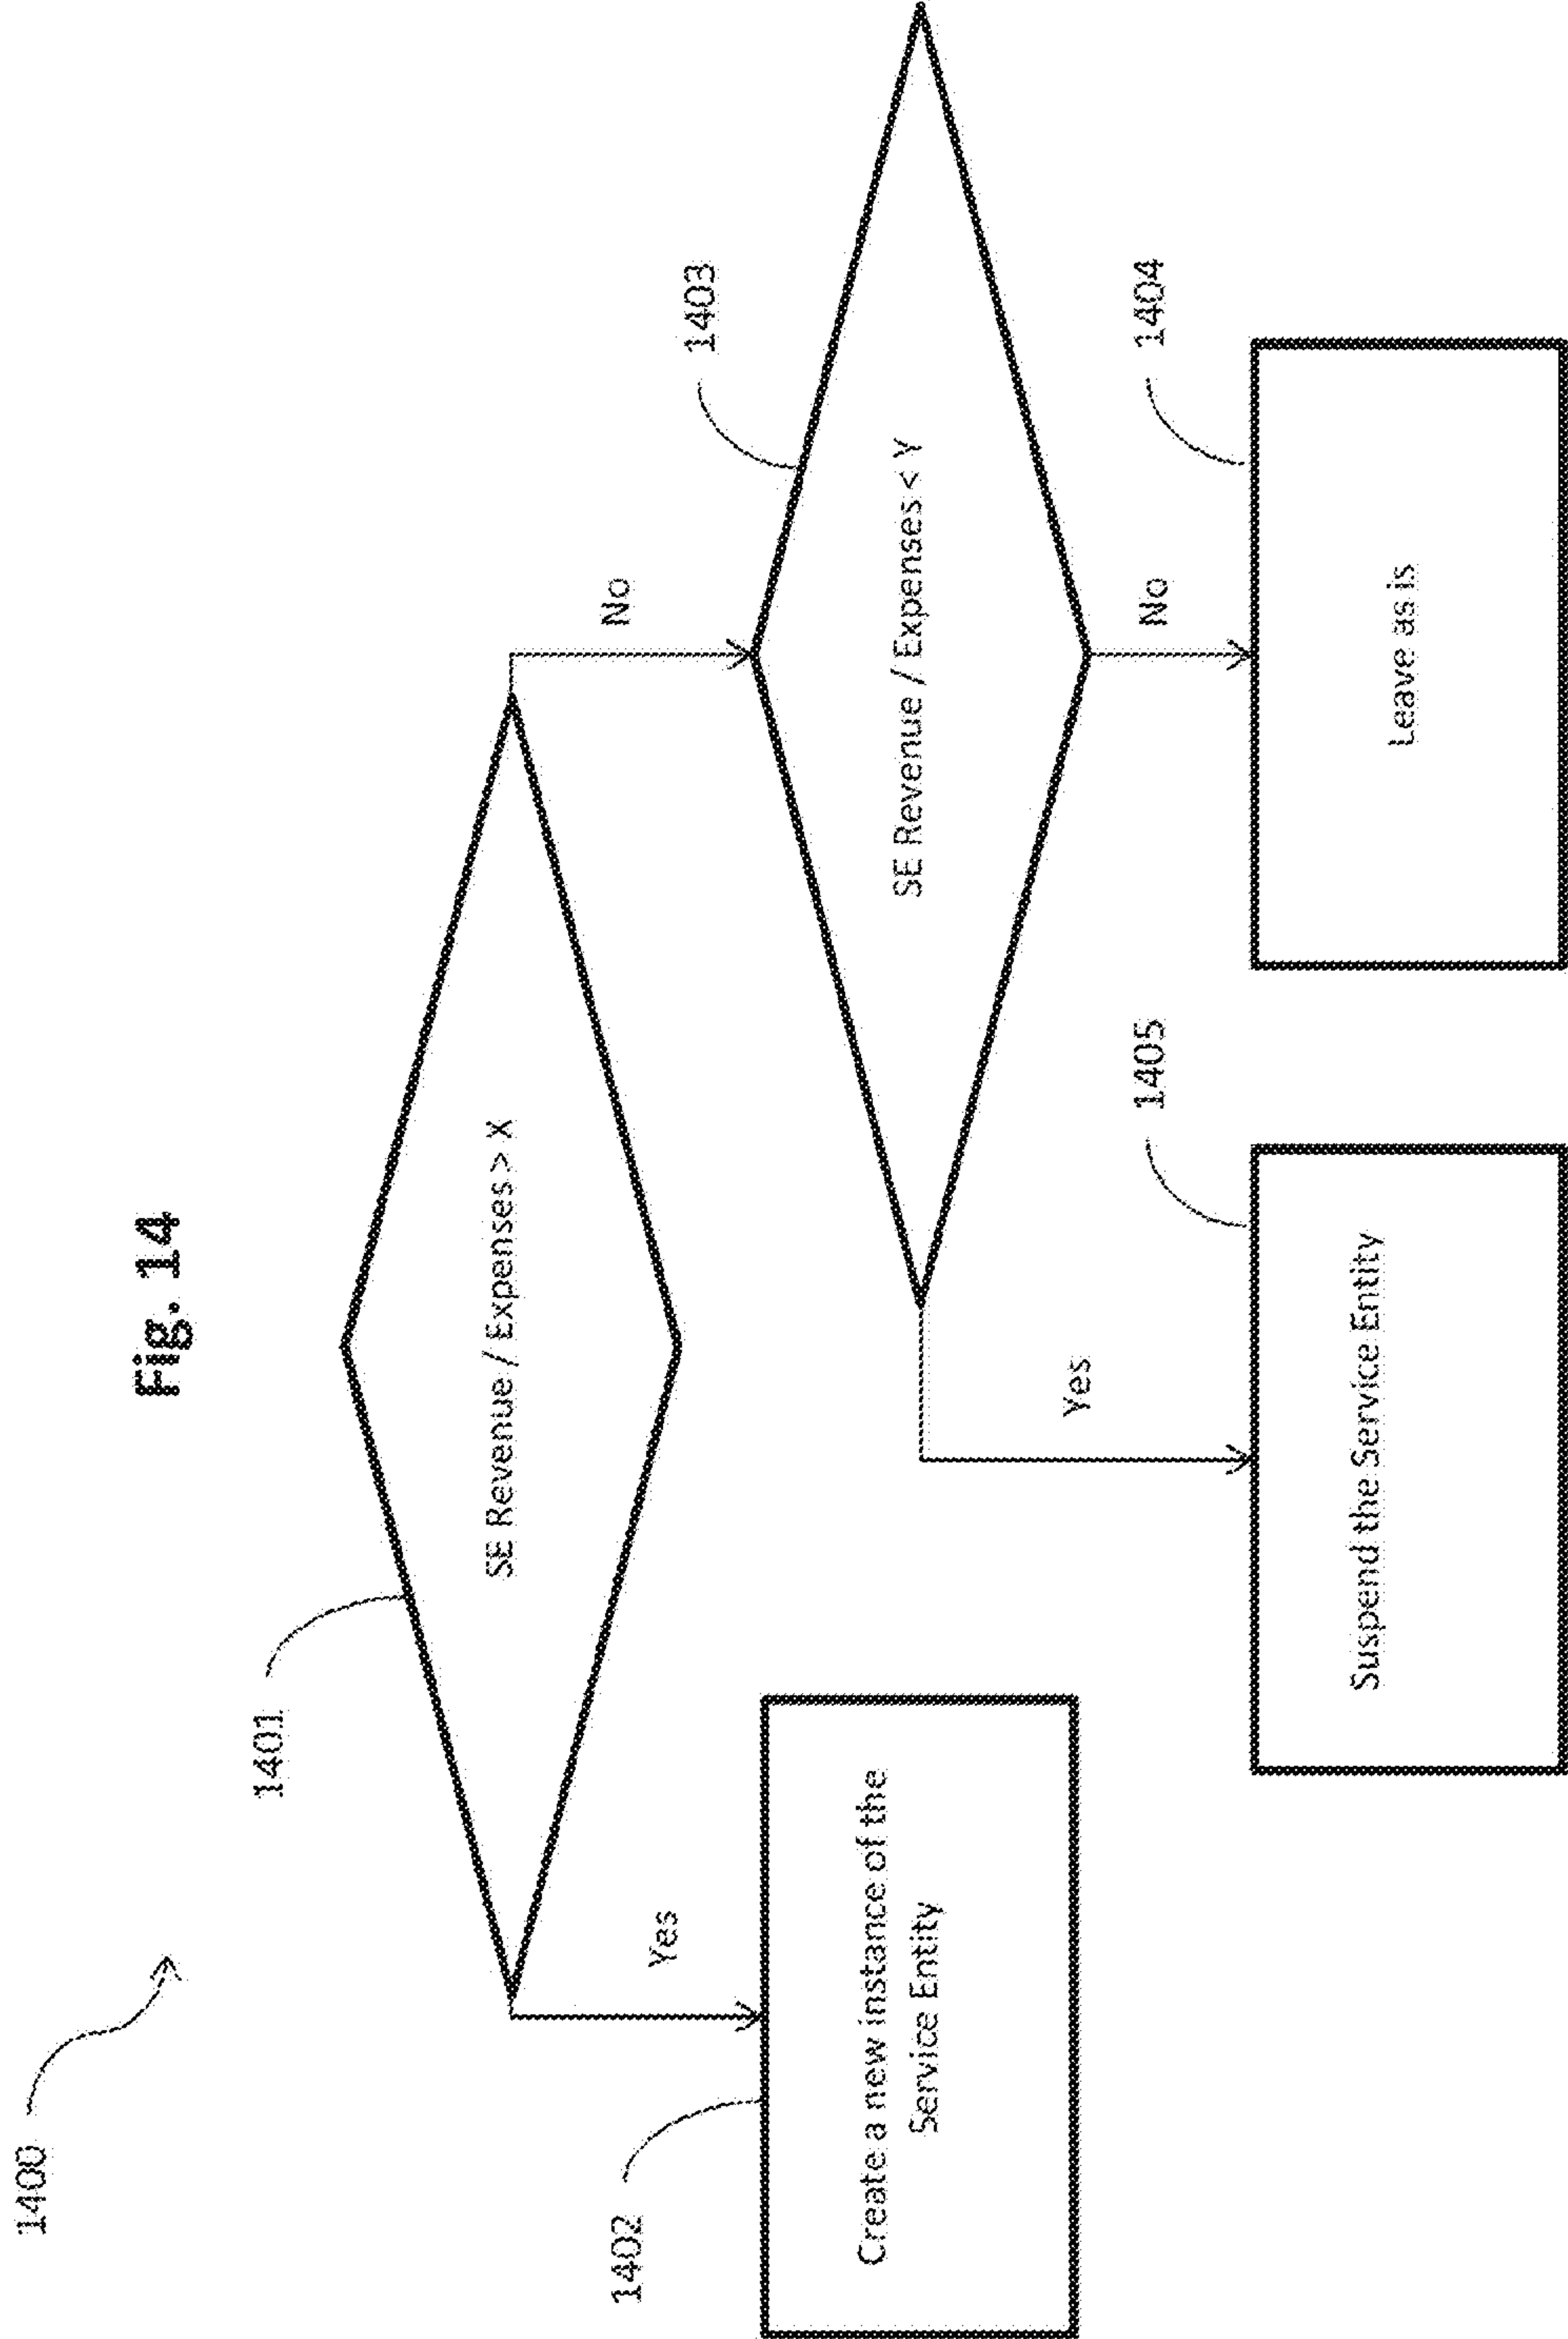


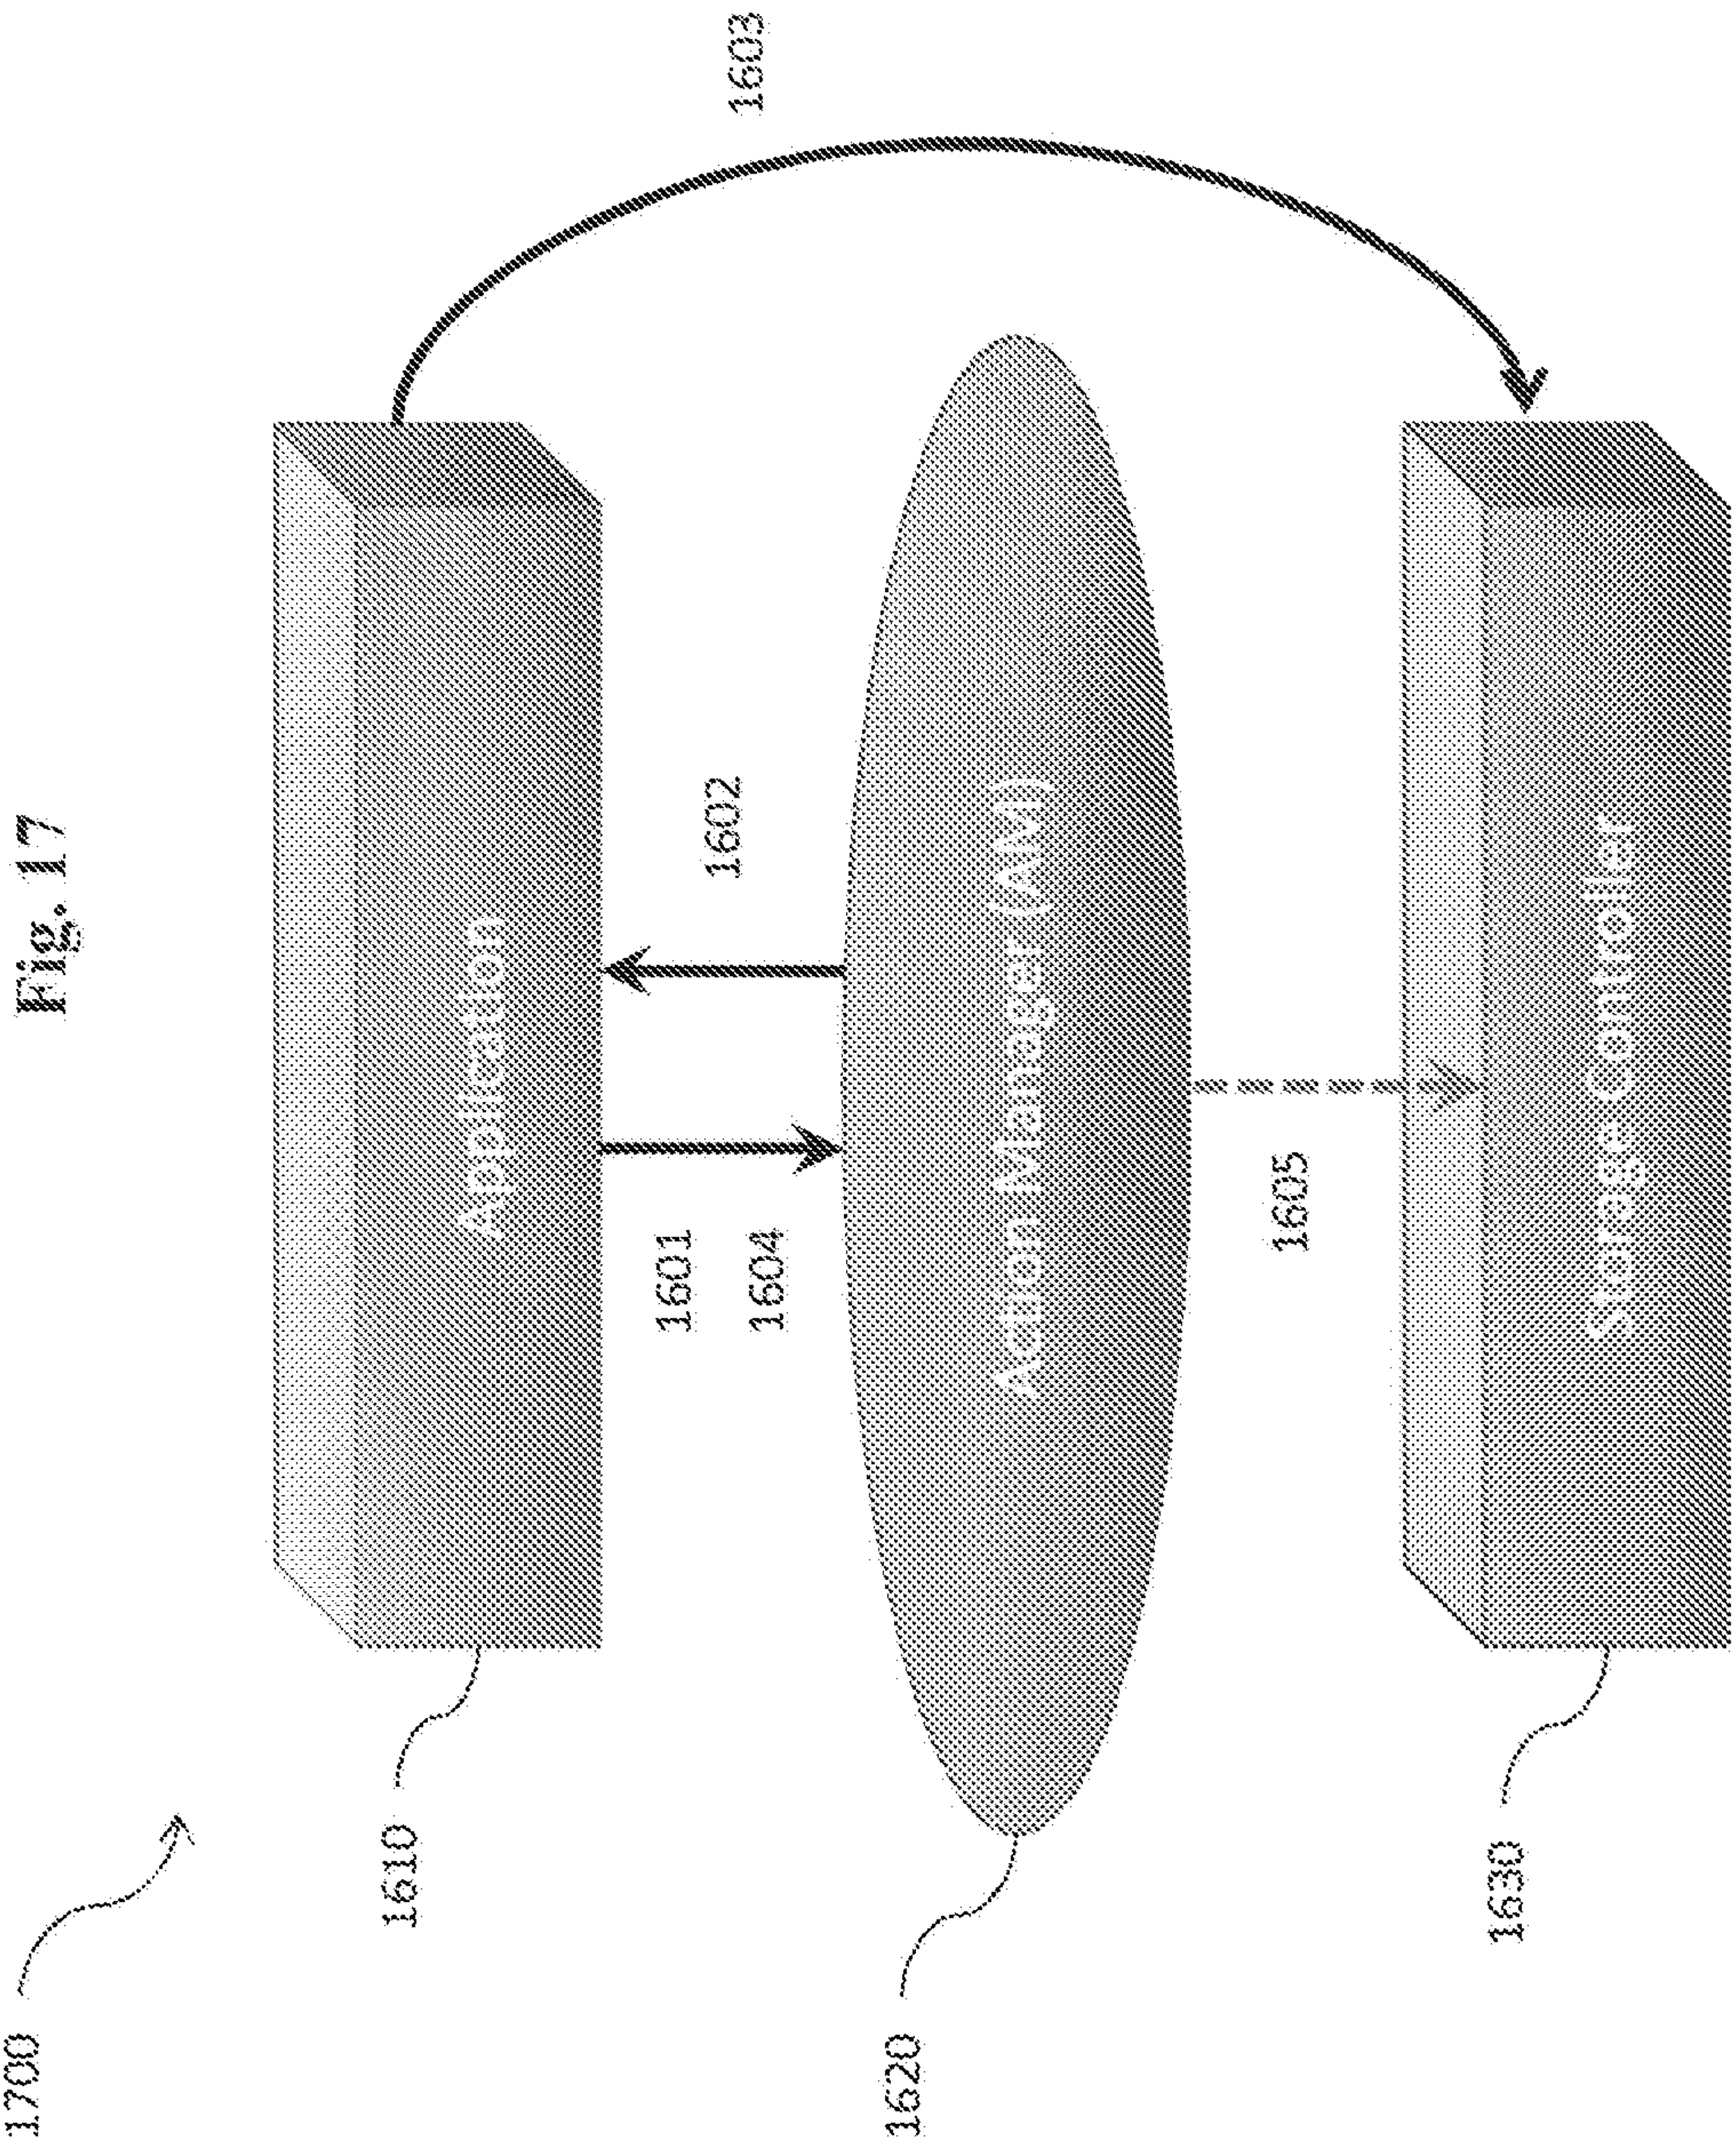
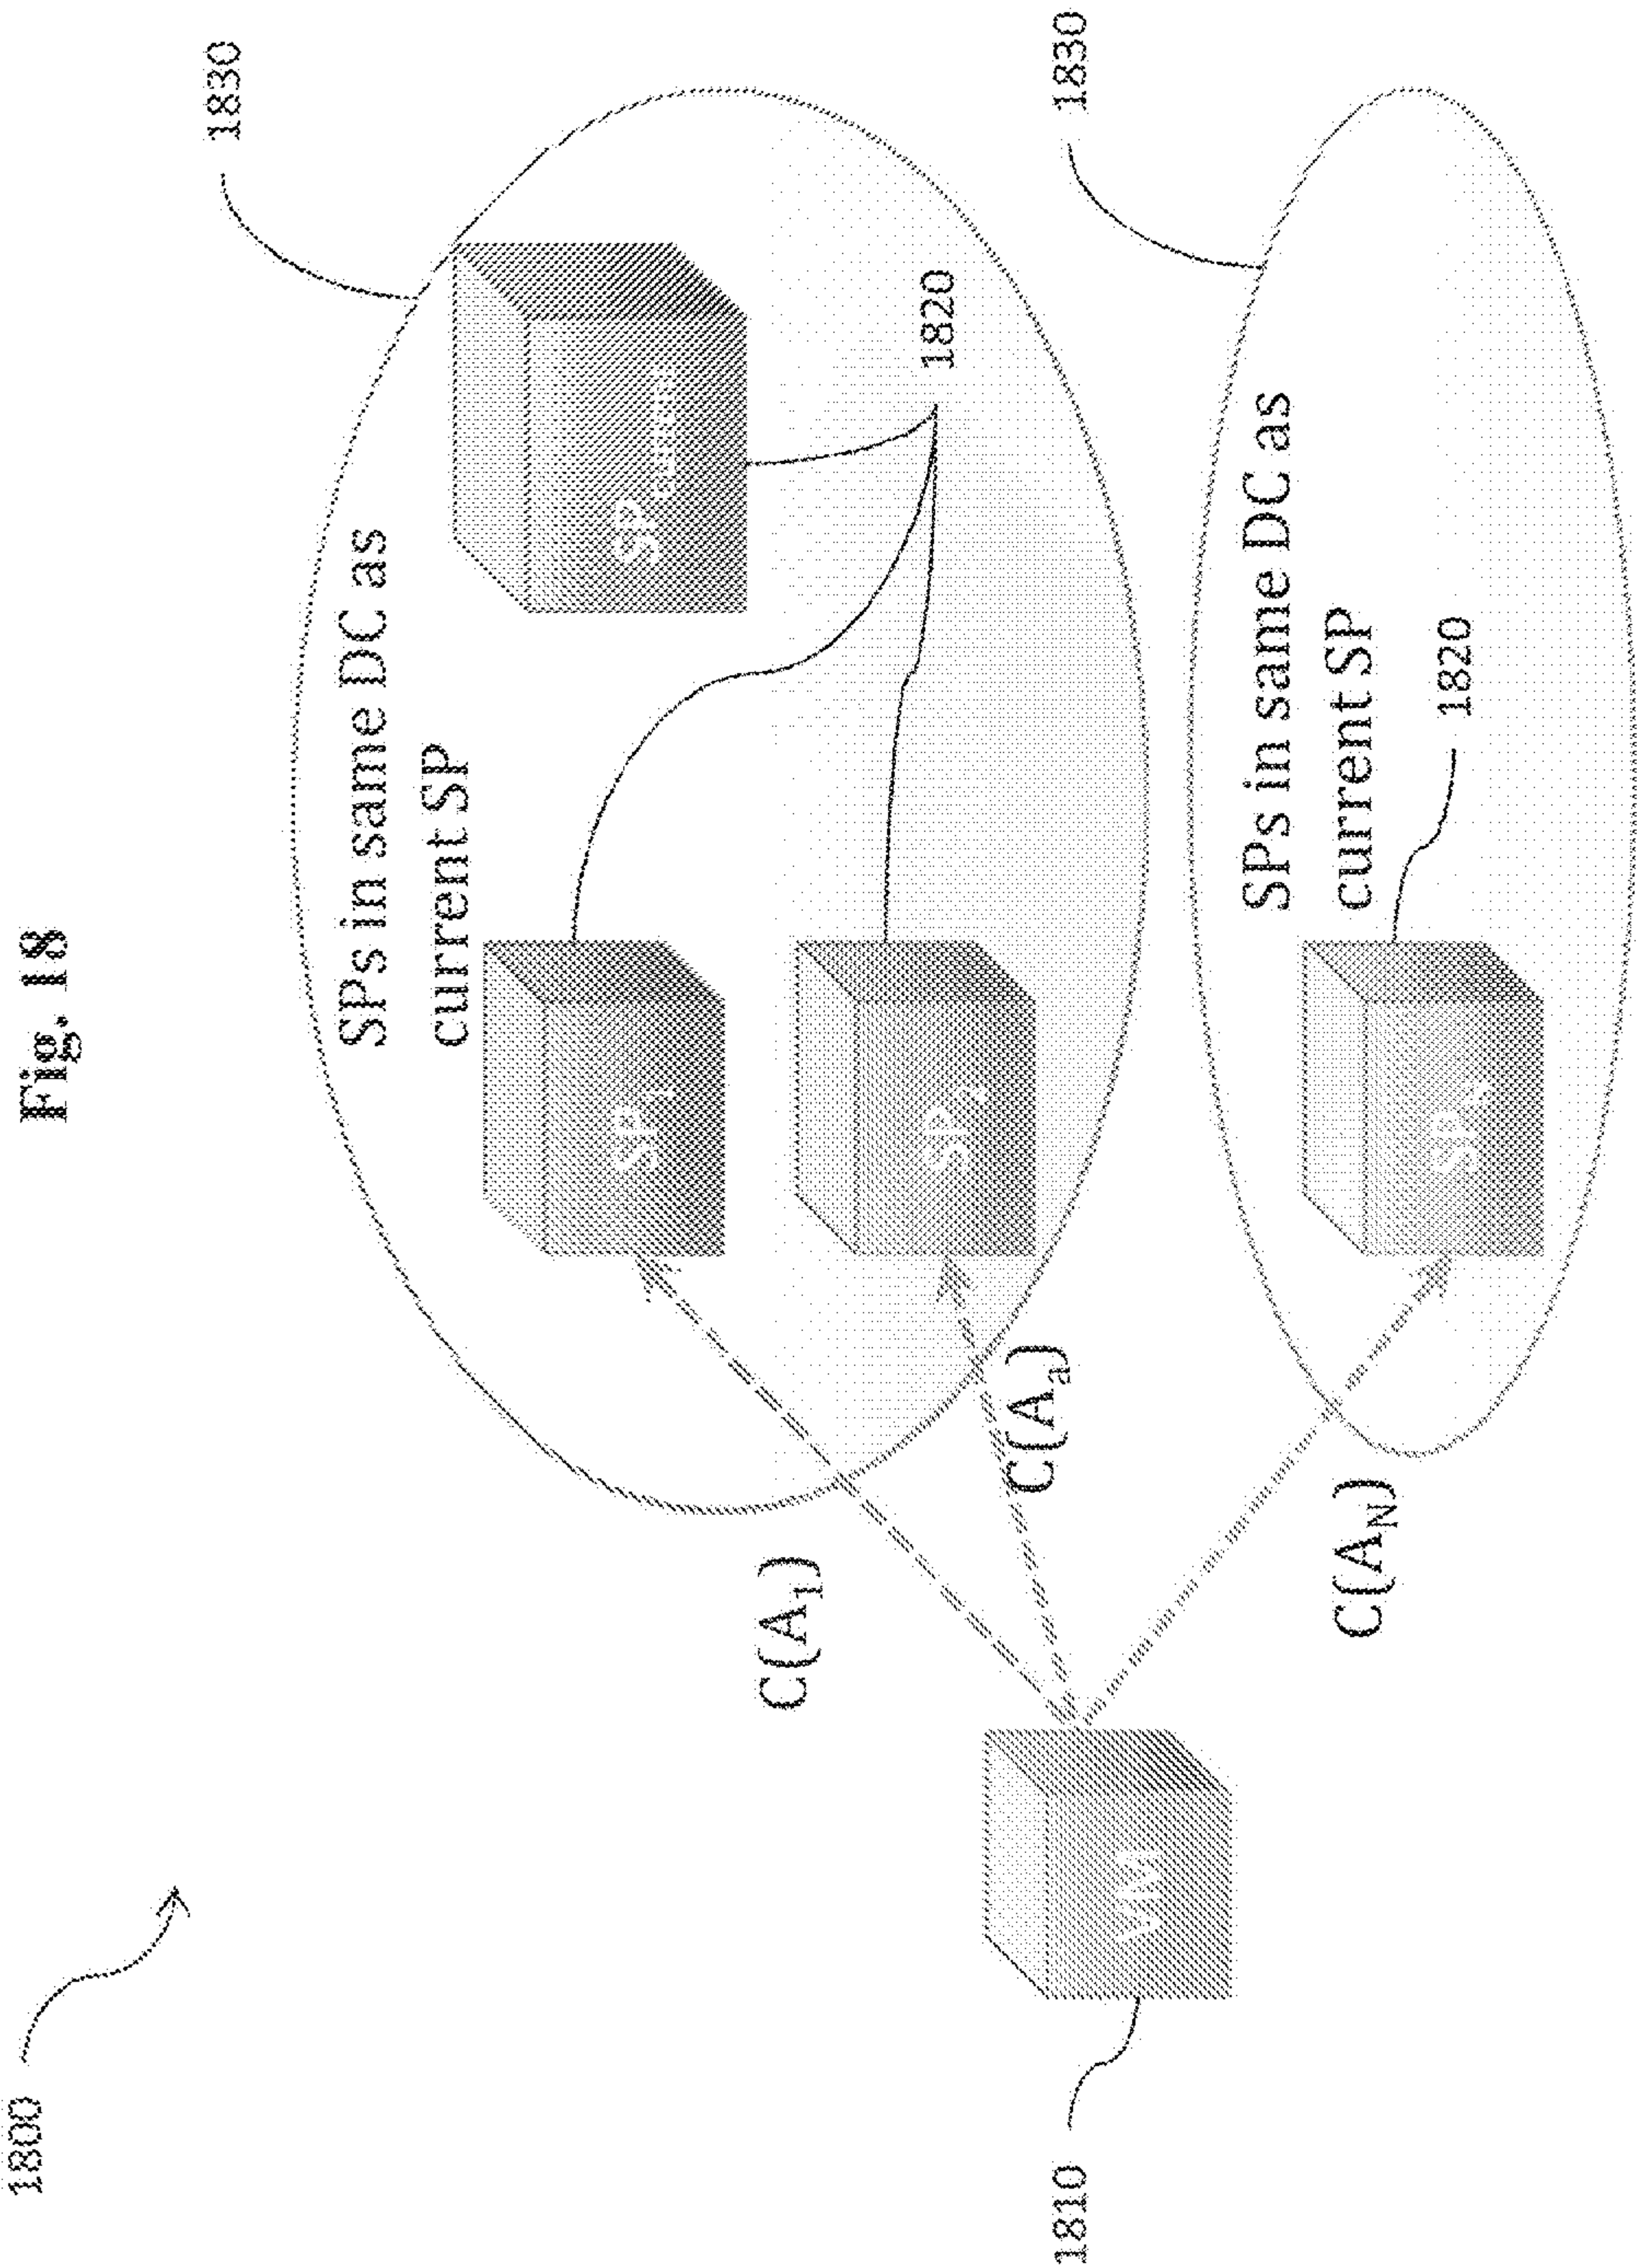


Fig. 15









MANAGING RESOURCES IN CONTAINER SYSTEMS

Matter enclosed in heavy brackets [] appears in the original patent but forms no part of this reissue specification; matter printed in italics indicates the additions made by reissue; a claim printed with strikethrough indicates that the claim was canceled, disclaimed, or held invalid by a prior post-patent action or proceeding.

CROSS-REFERENCE TO RELATED APPLICATIONS

This is a reissue application [claims the benefit] of U.S. application Ser. No. 14/810,284, filed Jul. 27, 2015 (now U.S. Pat. No. 9,888,067) (the “’067 Patent”), which is a continuation-in-part of U.S. Non-Provisional application Ser. No. 14/539,973, filed Nov. 12, 2014[;] (now U.S. Pat. No. 9,852,011) (the “’011 Patent”), which is a continuation-in-part of U.S. application Ser. No. 12/492,890, filed on Jun. 26, 2009 (now U.S. Pat. No. 8,914,511) (the “’511 Patent”). The ’067 Patent also claims the benefit of and priority to U.S. Provisional Application No. 62/077,898, filed Nov. 10, 2014[;], U.S. Provisional Application No. 62/077,852, filed Nov. 10, 2014[;], U.S. Provisional Application No. 62/077,860, filed Nov. 10, 2014[;], and U.S. Provisional Application No. 62/159,883, filed May 11, 2015[; the]. The ’067 Patent also claims the benefit of U.S. Provisional Patent Application Nos. 62/077,898, 62/077,852, and 62/077,860, (each filed Nov. 10, 2014), The disclosures of [each of which] the ’067 Patent and ’011 Patent and U.S. Provisional Application Nos. 62/077,898, 62/077,852, 62/077,860, and 62/159,883 are hereby incorporated by reference in their entirety and for all purposes.

Additionally, this application is related to U.S. patent application Ser. Nos. 14/810,302, 14/810,327, and 14/810,357, each filed Jul. 27, 2015, respectively titled “Managing Resources in Computer Systems Using Action Permits,” “Managing Application Performance in Virtualization Systems,” and “Moving Resource Consumers in Computer Systems,” the entire disclosures of which are incorporated by reference in their entirety and for all purposes.

FIELD

This specification relates generally to managing resources in virtualization systems, and more specifically, but not exclusively, to the control of such resources based on application demands, to the control of access to such resources, and the management of such resources in a computer-implemented container system.

BACKGROUND

Traditional computer system architectures typically include one or more dedicated computer servers for each application being run, and are often designed to include an excessive allocation of resources in order to be able to handle peak demands. Such partitioning of computer systems through dedicated servers and excessive allocation of resources can be costly, inefficient and difficult to scale and manage.

Virtualization, which refers to the abstraction of computer resources from their hardware or software-based physical constructs, is one manner of dealing with the aforementioned problems. One approach to virtualization is based on

one or more virtual machines (VMs), each of which is a software implementation of a computer that executes programs or applications as if it was a physical computer. A virtual machine operates like a physical computer and contains, for example, its own virtual (e.g., software-based) central processing unit (CPU), random access memory (RAM), hard disk storage, and network interface card (NIC). Each virtual machine in a virtualization system generally runs its own guest operating system (OS), and the virtual machines generally share the underlying physical machine resources of the system.

Another approach to virtualization is based on one or more containers, each of which is allocated exclusive access to compute resources, using a separate name space, that it may use to execute applications or programs, as if it was a separate operating system.

There are many potential benefits to operating in a virtualization system versus traditional architectures. For example, by permitting the sharing of hardware among applications workloads, virtualization can be used for improving resource utilization and reducing the need for excess resources to absorb peak traffic. Virtualization can also be used to improve the availability and robustness of applications, by shifting workloads among servers to handle fail-over situations. Similarly, virtualization provides flexible partitioning of applications, deployment, and operations. Notwithstanding the potential benefits, operating in a virtualization system presents several challenges and potential pitfalls, including significant operations management challenges.

For example, virtualization systems perform several input/output (I/O) intensive tasks—often concurrently. When multiple VMs or containers request to execute heavy-storage tasks at the same time (e.g., VM reboots, anti-virus database updates, OS updates, virus scans, and so on), storage controllers can face unrecoverable I/O congestion.

Conventional virtualization does not prioritize actions and account for resource impact of such actions. Managing resources in conventional virtualization systems includes evaluating alternative providers for a service or resource by comparing the attributes of the new service or resource and the current one. For example, when considering moving a VM, or a container, to a new storage device or array, conventional virtualization systems often consider the available storage amount and the data access latency of the new storage location, but ignore the time and resources it takes to move the VM, or container, to the new storage location. The impact of moving a VM can become significant when the amount of associated data to move is relatively large.

In addition, conventional virtualization systems render decisions based on the immediate impact of performing an action, rather than future impact (e.g., benefits) of an action. The virtualization systems may attempt to take into account a variety of parameters, such as how these parameters have evolved in the past, and how they are likely to evolve in the future. These systems, however, generally make decisions now (for the present time) and do not postpone decisions to the future. Moreover, in the past, predictions of future evolution have historically been rarely accurate.

Furthermore, conventional virtualization systems either do not allocate sufficient, or allocate too many, resources to an application that is initially deployed. In some systems, a default configuration is used. However, the default configuration may not be application-specific, may not consider the particular demand profile of the application, and/or cannot account for varying actual demand of the application. In other virtualization systems, modified configurations are

based on best practices for an application type and/or artificial load tests in a simulated production environment. A single configuration generally cannot consider all possible application demands, and artificial load tests do not generally reflect with complete accuracy application demands in the production environment.

As an additional challenge, once an application is deployed, configurations are generally altered only in response to reported degradation in application performance. Performance metrics are collected and analyzed and the configuration can be manually changed to reflect a user's understanding of the correlation between the performance degradation and the existing configuration. Unfortunately, the resulting configuration is static and, again, best suited for a single level of demand. If application demand is less than the target of the resulting configuration, the system's resources will be overprovisioned and result in waste. Alternatively, if application demand exceeds the resulting configuration, the performance of the application is limited. In any case, altering configurations in conventional virtualization systems generally occur only after the application performance has degraded, whereas overprovisioning resources for a particular application is generally not subject to detection.

Additionally, current planning techniques for future application demands involve making assumptions about future changes in infrastructure capacity based on historical infrastructure utilization. For example, if the environment is currently utilized at a rate of 50% and the assumption is that application demand will increase by 20% in the next 12 months, then a conclusion is made that the environment will be utilized at a rate of 60% in 12 months. However, these assumptions are generally based on infrastructure changes—not application demand. Despite any relationship between application demand and infrastructure utilization, these assumptions are generally not based on actual data and can result in overprovisioning or in limiting resources for a particular application.

An alternative virtualization technique can be found in container systems. Container systems provide an operating-system level virtualization in which the kernel of an operating system can allow for multiple isolated user space instances. Stated another way, a container is based on server virtualization that uses a shared operating system. Rather than virtualizing hardware and creating whole virtual machines, each with their own operating systems, containers run atop the shared operating system kernel and file system that looks and feels like a complete, isolated instance of the operating system. Like shipping containers for cargo, these software containers can ship applications across different network-based systems (e.g., cloud computing based systems) and limit the impact of one container's activities on another container.

A container system may include software abstractions to virtualize computer resources (or compute resources) which are used by applications running in the container ("containerized" applications). The container system provides means to provision containers, allocate and control the resources available to a container, deploy and execute applications in the container, and facilitate full use of the container resources by such containerized applications, while isolating them from other applications, sharing the underlying resources. When a containerized application accesses a virtualized container resource (e.g., CPU, memory, storage I/O, Network I/O), the container system maps this access to a direct access of the underlying real resource.

Container systems, like virtual machine systems, provide means for abstracting computer resources (or compute resources), controlling and isolating the allocations of these resources to applications, distributing and migrating applications flexibly, among multiple servers, to support scalable, highly-available, robust and efficient datacenter (DC) and cloud systems. Additional information on containers can be found, for example, at Linux Containers (available at <https://linuxcontainers.org>), [http://en.wikipedia.org/wiki/Docker_\(software\)](http://en.wikipedia.org/wiki/Docker_(software)), and <https://hadoop.apache.org/docs/r2.4.1/hadoop-yarn/hadoop-yarn-site/YARN.html>, each of which is hereby incorporated by reference in its entirety and for all purposes.

Since containers are based on shared operating systems, unlike virtual machines, they do not require guest operating systems and thus avoid some of the overhead of virtual machines. For example, deploying and configuring a container may involve setting data structures to map container resources to server and OS resources. Therefore, deploying and configuring a container can often be accomplished in seconds; deploying a virtual machine and a guest OS and configuring both to run an application may require substantially more time. Studies have shown that container-virtualization can offer significant efficiencies and performance advantages over virtual-machines—e.g., see "An Updated Performance Comparison of Virtual Machines and Linux Containers," by W. Felter et al., IBM Research, Jul. 21, 2014, available at [http://domino.research.ibm.com/library/cyberdig.nsf/papers/0929052195DD819C85257D2300681E7B/\\$File/rc25482.pdf](http://domino.research.ibm.com/library/cyberdig.nsf/papers/0929052195DD819C85257D2300681E7B/$File/rc25482.pdf), the disclosure of which is hereby incorporated by reference in its entirety and for all purposes.

A virtualization system may mix and match virtual machines and containers. For example, containers may run over virtual-machines. Additionally, a group of virtual machines may be containerized, much like any application, and executed by a container.

With lower overheads than VMs, the number of containers sharing a host is often an order of magnitude (e.g., 50-200) larger than the number of VMs (e.g., 5-20). Furthermore, containers perform a faster dynamic of change events (e.g., deploy, delete, increase/reduce resources, and so on). Accordingly, container system management presents challenging scalability and response time problems compared to VMs.

In view of the foregoing, a need exists for an improved resource management system and method for container control in an effort to overcome the aforementioned obstacles and deficiencies of conventional virtualization systems.

SUMMARY

This specification describes new technologies relating to the management of resources and performance in virtualization systems, including container systems and multi-cloud systems. For example, these technologies introduce the use of supply chain economics and other techniques to offer a unified platform to integrate, optimize or improve, and automate resource and performance management in a container system. The economics-based methods can also extend to other virtualization systems for managing application performance.

In general, one aspect of the subject matter described in this specification can be embodied in methods and systems that include the actions of determining, for example, by a Container Manager running on a data processor in a con-

5

tainer system, a computer resource bundle to be purchased for a container of the container system using virtual currency units; identifying multiple computer servers available to the container system; offering the computer resource bundle; determining a purchase price for the computer resource bundle, in virtual currency units, for each of the identified computer servers; automatically selecting, by the Container Manager, one of the identified computer servers based at least in part on the purchase price for the computer resource bundle set for each of the identified computer servers; and allocating the computer resource bundle from the selected one of the multiple computer servers to the container. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Another aspect of the subject matter described in this specification can be embodied in methods and systems that include determining, by a Container Manager running on a data processor in a first container system, a computer resource bundle to be purchased for a container in the first container system using virtual currency units; receiving, from a Proxy Manager of a second container system offering the computer resource bundle, a purchase price for the computer resource bundle in virtual currency units; automatically purchasing the computer resource bundle from the second container system based at least in part on the purchase price received from the Proxy Manager of the second container system; allocating the computer resource bundle from the second container system to the container in the first container system; and dispatching the container from the first container system to execute at the second container system. In some embodiments, dispatching the container includes sending an image of a stateful container to another server or, in the case of stateless containers, initiating a new container on another server and terminating the server on the first server. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

A further aspect of the subject matter described in this specification can be embodied in methods and systems that include the actions of determining, by a Container Manager running on a data processor in a first container system, a computer resource bundle to be purchased for a container in the first container system using virtual currency units; identifying at least a second container system and a third container system; offering the computer resource bundle at respective purchase prices, in virtual currency units; receiving, from respective Proxy Managers of the second container system and the third container system, a first purchase price for the computer resource bundle for the second container system and a second purchase price for the computer resource bundle for the third container system; automatically selecting one of the second container system and the third container system based at least in part on the respective purchase prices for the computer resource bundle for the second container system and the third container system; and allocating the computer resource bundle from the selected one of the second container system and the third container system to the container in the first container system. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Still another aspect of the subject matter described in this specification can be embodied in methods and systems that include the actions of allocating a computer resource bundle to an application executed by a container system; determining, by an Application Element Manager running on a data processor in the container system, the value of a service

6

level agreement parameter for the application based on the allocated computer resource bundle; comparing the determined service level agreement parameter level for the application to a threshold service level agreement parameter level; automatically modifying the allocation of computer resources to the application depending on whether the identified service level agreement parameter level for the application is below or above the threshold service level agreement parameter level; and repeating the determining, comparing and automatically modifying steps until the operation of the application in the container system is suspended or terminated. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Yet another aspect of the subject matter described in this specification can be embodied in methods and systems that include the actions of purchasing a computer resource bundle, using virtual currency units, for an application running at a container in a container system; allocating the purchased computer resource bundle to the application; determining, by a Container Manager running on a data processor in the container system, the value of a service level agreement parameter for the application; automatically modifying the purchase of the computer resource bundle based at least in part on a comparison of the determined value of a service level agreement parameter to a predetermined threshold service level agreement parameter level; and repeating the determining, comparing and automatically modifying steps until the operation of the application in the container system is suspended or terminated. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Another aspect of the subject matter described in this specification can be embodied in methods and systems that include network-aware placement of containers to minimize input/output (I/O) latency. Some methods include the actions of determining a target I/O latency metric, for storage or network communications of an application running in a container system; determining, by a container manager, a first I/O pathway that meets the target latency and an amount of I/O bandwidth, over the first I/O pathway, to be acquired for the container; allocating the determined amount of I/O bandwidth, over the determined first I/O pathway to the application; automatically allocating the determined amount of I/O bandwidth to the application, from the second I/O pathway, based at least in part on the determined I/O bandwidth utilization of the first I/O pathway. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Another aspect of the subject matter described in this specification can be embodied in methods and systems that include the actions of determining, by a Container Manager running on a data processor in a container system, an amount of input/output (I/O) bandwidth—for example, for storage or network communications—to be acquired for a container of the container system; identifying at least a first computer server and a second computer server in the container system, each offering at least the determined amount of I/O bandwidth; allocating, from the first computer server, the amount of I/O bandwidth to the container; determining the I/O bandwidth utilization of the first computer server; automatically allocating the amount of I/O bandwidth to the container, from the second computer server, based at least in part on the determined I/O bandwidth utilization of the first computer server. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Yet another aspect of the subject matter described in this specification can be embodied in methods and systems that include the actions of determining, by a Container Manager running on a data processor in a container system, an amount of I/O bandwidth—for example, for storage or network communications—to be acquired from a first computer server for a container running on the first computer server; allocating, from the first computer server, the amount of I/O bandwidth to the container for a first period of time; after the first period of time has elapsed, determining that the I/O bandwidth utilization of the first computer server is greater than a threshold limit; automatically identifying at least a second computer server in the virtualization system offering at least the determined amount of I/O bandwidth of the first computer server; moving the container from the first computer server to the identified second computer server; and allocating from the second computer server, the determined amount of I/O bandwidth, to the container, for a second period of time. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Yet another aspect of the subject matter described in this specification can be embodied in methods and systems for fractional resource allocation that include the actions of determining, by a Container Manager running on a data processor in a container system, a resource bundle to be acquired from a first computer server for a container running on the first computer server; allocating, from the first computer server, the determined resource bundle for a first period of time; after the first period of time has elapsed, determining that the utilization of the first computer server exceeds a threshold limit; determining the excess resources corresponding to the excess utilization; automatically identifying a second computer server to replicate the container; replicating the container at the identified second computer server; allocating from the identified second computer server, a resource bundle equal or lower than the determined excess resources, to the replicated container, for a second period of time. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Yet another aspect of the subject matter described in this specification can be embodied in methods and systems that include the actions of determining, by a Container Manager running on a data processor in a container system, an amount of I/O bandwidth—for example, for storage or network communications—to be acquired from a first computer server for a container running on the first computer server; allocating, from the first computer server, the amount of I/O bandwidth to the container for a first period of time; after the first period of time has elapsed, determining that the I/O bandwidth utilization of the first computer server exceeds an I/O bandwidth threshold limit; determining the excess I/O bandwidth corresponding to the excess I/O bandwidth utilization; automatically identifying a second computer server in the container system to replicate the container; allocating from a second computer server, an amount of I/O bandwidth equal or lower than the determined excess I/O bandwidth, to the replicated container, for a second period of time; and using a Load Balancer to divide the processing of I/O bandwidth between the container, at the first computer server, and the replicated container at the second computer server. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Still another aspect of the subject matter described in this specification can be embodied in methods and systems that

include the actions of determining, by an Application Element Manager running on a data processor in a container system, a software license to be purchased using virtual currency units by an application operating in the container system; determining, by a Licensing Element Manager running on the data processor, the number of copies of the software licenses available for purchase in the container system; determining, by the Licensing Element Manager, the purchase price for a copy of the software license; automatically purchasing, by the Application Element Manager, a copy of the software license for the application based at least in part of the determined purchase price for a copy of the software license; and allocating the purchased copy of the software license to the application. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

A further aspect of the subject matter described in this specification can be embodied in methods and systems that include the actions of allocating, from a computer server in a container system, a computer resource bundle to a container operating in the container system; determining, by a Container Manager running on a data processor in the container system, a productivity score of the container; determining that the productivity level of the container is below a threshold productivity level; and automatically initiating, by the Container Manager, at least one of a suspension sequence and a termination sequence for the container based at least in part on the determination that the productivity level of the container is below a threshold level. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Another aspect of the subject matter described in this specification can be embodied in methods and systems that include the actions of allocating a computer resource bundle to a computer server operating in a container system; determining, by a Computer Server Element Manager running on a data processor in the container system, a productivity score of the computer server; determining that the productivity level of the computer server is below a threshold productivity level; and automatically initiating, by the Computer Server Element Manager, at least one of a suspension sequence and a termination sequence for the computer server based at least in part on the determination that the productivity level of the computer server is below a threshold level. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Yet another aspect of the subject matter described in this specification can be embodied in methods and systems that include the actions of determining a computer resource bundle to be purchased for a containerized container running at a parent container in a container system; purchasing the determined computer resource bundle; and allocating the purchased computer resource bundle to the containerized container. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

As further discussed herein, the subject matter includes methods to place containers at hosts and allocate computer resources (or compute resources) to the containers based on a supply-chain marketplace for buying/selling resources; migrate containers to hosts based on purchasing resources in a supply chain marketplace; automatically replicate an overloaded container, allocate resources to the replicated containers, and partition the container workload among its replicated containers; automatically replicate an overloaded container, allocate resources to the replicated containers, and

partition the I/O stream among the replicated containers using a load balancer; allocate resources to containers to handle the workload of an application and assure a target Service Level Agreement (SLA) performance metric; allocate software licenses to applications running in containers; and select a production system and to dispatch a container from a development system to a production system; assign a productivity score to a host running a container system, based on its virtual Profit and Loss (P&L) performance, remove containers from a host with low productivity and set it into standby mode, or shut it down to save electricity. Respective methods to shift a host from standby mode to execute containers systems, when the productivity score is sufficiently high, are also disclosed.

As also discussed herein, the subject matter further includes methods to place highly communicating containers “close” to one another to minimize their I/O latency and maximize its bandwidth. In some embodiments, “close” can be measured by the number of switches along the I/O pathway. Additional subject matter disclosed includes methods to cluster containers with related applications into container points of delivery (cPODs) and migrate an entire cPOD between clouds through an atomic action. Further disclosure includes methods and systems to containerize a container to create a supply chain hierarchy of nested containers and use it to provide scalable resource allocation whereby a parent container “sells” its resources, using virtual currency, to its daughter containers.

These and other embodiments can optionally include one or more of the following features. Virtual currency units can be used to control the purchase of computer resources. Computer resources can be allocated in the event a budget of the acquiring computer element is sufficient. The allocation of computer resources can be in exchange for the payment of virtual currency units. Virtual currency units can be converted into real currency, such as United States dollars, or euros. A performance metric of an application executing in a container can be determined, and the allocation of computer resources is based on such performance metric. A performance metric of an application executing in a container can be determined, and the allocation of virtual currency units to the application or the container can be based on such performance metric. The computer resources to be allocated can include at least one of an allotment of computer memory, an allotment of program execution scheduling for one or more central processing units, an allotment of storage interface bandwidth, an allotment of network or I/O interface bandwidth, and an allotment of a storage array.

Additionally, according to various embodiments, a computer resource provider can be selected based on the price offered for the computer resources. The price offered for the computer resources can be determined by an auction process, or by the utilization of the computer resources (e.g., congestion pricing), or by a service level metric provided by the resources, or by the cost of the resource, or any combination of these and other pricing factors. The amount of computer resources purchased can be dependent on a computed computer resource utilization. A proxy manager of a second container system optionally controls the allocation of computer resources to a first container system. The foregoing features are optional and not exhaustive, and they may be separately realized or realized in various combinations.

Moreover, particular embodiments of the subject matter described in this specification can be implemented to realize one or more of the following advantages. Computing resources may be monitored and managed in container

systems based on economic and non-economic principles, both in single-cloud and multi-cloud systems. Such management of resources may be used to control the delivery of service level agreements, and the distribution of software licenses, in container systems. They also may be used to suspend or terminate computing resources in container systems. These advantages and features are optional and not exhaustive, and they may be separately realized or realized in various combinations.

Another aspect of the subject matter described in this specification can be embodied in systems and methods that assure application performance by matching the supply of resources (e.g., application resources, VM resources, or physical resources) with the fluctuating demand placed on the application. For example, the systems and methods disclosed herein can be used to ensure that the application is allocated sufficient resources when it is initially deployed to handle anticipated demand; dynamically alter the resources allocated to the application during operation by matching the resource requirements to the actual measured application demand; and predict future resource requirements based on planning assumptions related to future application demand.

In some embodiments, the systems and methods described herein apply to applications that scale vertically—e.g., to scale the additional resources that are allocated to the existing application components. Additionally and/or alternatively, the systems and methods described herein apply to applications that scale horizontally. Horizontally scaling applications are scaled by provisioning additional application components and dividing the work to be done among the old and new components (compared to adding resources to existing components).

In one embodiment, the application component that horizontally scales buys additional resources from a new provider (e.g., virtual server)—rather than the current provider—and is provisioned along with a new application component. The new application component can purchase the required resource from the new provider. Conversely, if the application component can no longer afford the resources, the provider and the application component can be decommissioned. Alternatively, if the application component can no longer afford the resources, the resources can be removed.

In some embodiments, the system can be used in real time in order to efficiently allocate the required resources to applications based on a balance of supply and demand. The system is able to scale an application either vertically (reconfigure) or horizontally (auto-scale) such that more detailed control can be achieved by considering target Quality of Service (QoS) metrics such as transactions throughput rate or transaction response time targets defined for these measures. These targets represent the desired service level of the application. The application may perform below its desired service level and not achieve the defined targets for these QoS metrics. If the application achieves the target, it is considered to perform at its desired service level. An application that is below its desired service level can be brought into compliance by either adding more resources (vertical scale) or more application components (horizontal scale).

The ability of an application to buy resources is a function of its revenue and the expenses based on the price of the resources. In the scenario where an application is below its desired service level, the system can allocate additional budget to the application, providing a greater power to buy resources. Therefore, QoS metrics can be used to gain additional granularity in the allocation of resources.

In other embodiments, the system can also be used when deploying new applications in order to determine the required or preferred resource allocation. By simulating application demand (e.g., transactions), the system can determine the required or preferred number of application components and the appropriate allocation of resources.

Additionally and/or alternatively, the system can be used when planning for future application needs. By simulating future application demand (e.g., transactions), the system can determine the required number of application components and the appropriate allocation of resources.

Moreover, particular embodiments of the subject matter described in this specification can be implemented to realize one or more of the following additional advantages. Computing resources may be monitored and managed in virtualization systems based on economic and non-economic principles, both in single-cloud and multi-cloud systems. Such management of resources may be used to control the delivery of service level agreements, and the distribution of software licenses, in virtualization systems. They may also be used to suspend or terminate computing resources in virtualization systems. These advantages and features are optional and not exhaustive, and they may be separately realized or realized in various combinations.

As also discussed herein, the subject matter further includes methods relating to the management of resources in virtualization systems, including multi-cloud systems. These technologies introduce the use of economic-based cost analysis, and other techniques to offer a comprehensive framework to manage the number and rate of actions and prioritize resource management in a virtualization system.

Moreover, particular embodiments of the subject matter described in this specification can be embodied in methods and systems that regulate access of consumers (e.g., applications, containers, or VMs) to resources and services (e.g., storage). In one embodiment, this regulation occurs through the use of access or action permits, referred to as permits that the consumer acquires from an intermediate entity—an Action Manager (AM)—prior to accessing the resource or service. Regulating access includes, for example, controlling one or more of the number of concurrent accesses to a particular resource, the rate at which consumers access the resource, the total number of consumers in a group of consumers accessing the resource, and the total rate at which a group of consumers accesses a resource. According to various embodiments, similar regulation is applied to a group of resources (rather than a single resource).

In some embodiments, there can be one type of permit per provider. The AM can sell multiple types of permits, regulating access to a number of resources. Each permit can be associated with a predetermined price. Additionally and alternatively, this price can be dynamically adjusted, for example, taking into consideration the availability of permits the AM possesses. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

In another embodiment of the invention, the consumer acquires the permits from the resource or service itself.

Another aspect of the subject matter described in this specification can be embodied in methods and systems that formulate and evaluate the option to move a consumer to a new provider. According to various embodiments, “formulating” can include the attributes taken into account when considering the option to move to the new provider. The cost of moving can be part of the comparison between two or more alternatives (e.g., keeping a VM in an existing infrastructure or moving the VM to an external cloud provider).

Cost can be expressed in actual currency or any unit suitable for the comparison. For example, moving time can be expressed in a real value that quantifies the cost of the VM downtime. In contrast, if there is a strict limit on acceptable downtime, the cost of moving the VM can be expressed in terms of a cost metric of downtime.

According to various embodiments, “evaluating” includes making the decision (e.g., initiating an action based on the decision) and determining the right time to take the action. Compared to other economics-based decision-making systems, one embodiment described herein postpones the decision for the future, advantageously waiting for a sufficient amount of time until the decision-maker is convinced that the decision is the right one. Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

Moreover, particular embodiments of the subject matter described in this specification can be implemented to realize one or more of the following advantages. Computing resources may be monitored and managed in virtualization systems based on economic and non-economic principles, both in single-cloud and multi-cloud systems. Such management of resources may be used to control the delivery of service level agreements based on access control, and the distribution of software licenses, in virtualization systems. They may also be used to suspend or terminate computing resources in virtualization systems. These advantages and features are optional and not exhaustive, and they may be separately realized or realized in various combinations.

Another aspect of the subject matter described in this specification can be embodied in methods and systems that include determining, by a container manager running on a data processor, a computer resource bundle to be purchased for a container using virtual currency units, wherein the container is a software implementation of a computer that runs on top of a shared operating system and executes one or more applications on the shared operating system; identifying multiple computer servers available to offer the computer resource bundle; determining a purchase price for the computer resource bundle, in virtual currency units, for each of the identified computer servers; automatically selecting, by the container manager, a first one of the identified multiple computer servers based at least in part on the purchase price for the computer resource bundle for each of the multiple computer servers; and allocating the computer resource bundle from the first one of the multiple computer servers to the container.

In some embodiments, the method further includes initiating, following an increase in the purchase price for the allocated computer resource bundle, a credit check process by the container manager, wherein the credit check process results in an increase in budget of virtual currency units available for usage by the container; and determining, following the initiation of the credit check process, whether the container is to be moved from the first one of the multiple computer servers to a second one of the multiple computer servers, wherein the moving determination is based at least in part on the purchase price for the computer resource bundle offered by the second one of the multiple computer servers.

In some embodiments, the method further includes maintaining a budget, in virtual currency units, for the container; determining whether the budget of virtual currency units for the container is sufficient to pay the purchase price for the computer resource bundle for the selected first one of the multiple computer servers; allocating the computer resource bundle from the first one of the multiple computer servers to

13

the container only if the budget of the virtual machine has been determined to be sufficient to pay the purchase price for the computer resource bundle for the first one of the multiple computer servers; and transferring virtual currency units from the budget for the container to a budget for the first one of the multiple computer servers based on the purchase price for the computer resource bundle determined for the first one of the multiple computer servers.

In some embodiments, the method further includes maintaining a budget, in virtual currency units, for an application executing on the container; determining a purchase price, in virtual currency units, for at least a part of the computer resource bundle allocated to the container; determining whether the budget of virtual currency units for the application is sufficient to pay the purchase price for the at least a part of the computer resource bundle; allocating, to the application executing on the container, the at least a part of the computer resource bundle if the budget for the application is sufficient to pay its determined purchase price; and deducting virtual currency units from the budget for the application based on the determined purchase price for the at least a part of the computer resource bundle.

In some embodiments, the computer resource bundle includes one or more computer resources, the method further comprising: determining a performance metric level of an application executing on the container; and adjusting the allocation of at least one of the one or more computer resources to the application based on a comparison of the determined performance metric level to a predetermined threshold performance metric level.

In some embodiments, the performance metric level of the application executing on the container is correlated to at least one of the transaction rate for the application, the number of transactions performed by the application, and the completion time or delay for the application.

In some embodiments, the computer resource bundle includes one or more computer resources, the method further comprising: determining a performance metric level of an application executing on the container; and adjusting a budget of virtual currency units for the application based on a comparison of the determined performance metric level to a predetermined threshold performance metric level.

In some embodiments, the computer resource bundle includes at least one of an allotment of computer memory, an allotment of application execution scheduling for one or more central processing units, an allotment of storage interface bandwidth, an allotment of network interface bandwidth, and an allotment of a storage array.

In some embodiments, allocating the computer resource bundle from the selected one of the multiple computer servers to the container comprises moving the container to the selected one of the multiple computer servers from a different computer server.

In some embodiments, the computer resource bundle includes a plurality of computer resources, and wherein the automatically selecting one of the multiple computer servers includes selecting the computer server having the lowest combined purchase price for the plurality of computer resources.

In some embodiments, the allocated computer resource bundle includes at least a first computer resource and a second computer resource, the method further comprising de-allocating the first computer resource based at least in part on a comparison of a determined usage of the first computer resource to a determined usage of the second computer resource.

14

In some embodiments, the computer resource bundle includes multiple computer resources, and wherein the determining a computer resource bundle to be purchased is based at least in part on a computed utilization of one of the multiple computer resources.

In some embodiments, the method further includes increasing the amount of the computer resource bundle to be purchased if it is determined that the computed utilization of the one of the multiple computer resources is below a predetermined threshold.

In some embodiments, the method further includes decreasing the amount of the computer resource bundle to be purchased if it is determined that the computed utilization of the one of the multiple computer resources is above a predetermined threshold.

In some embodiments, determining a purchase price for the computer resource bundle includes receiving one or more auction bids for purchase of the computer resource bundle.

In some embodiments, the method further includes maintaining a budget, in virtual currency units, for the first one of the multiple computer servers; determining a purchase price, in virtual currency units, for a second computer resource bundle for a provider of the second computer resource bundle; determining whether the budget of virtual currency units for the first one of the multiple computer servers is sufficient to pay the purchase price for the second computer resource bundle; allocating, to the first one of the multiple computer servers, the second computer resource bundle if the budget for the first one of the multiple computer servers is sufficient to pay the determined price; and deducting virtual currency units from the budget for the first one of the multiple computer servers based on the determined purchase price.

In some embodiments, the second computer resource bundle includes at least one of an allotment of electrical power, an allotment of air temperature cooling resources, an allotment of physical space, an allotment of operational expense, and an allotment of capital expenditure.

In some embodiments, the method further includes converting virtual currency units to government issued monetary currency based on a predetermined exchange rate.

Another aspect of the subject matter described in this specification can be embodied in methods and systems that include determining, by a container manager running on a data processor in a first container system that includes at least one container, a computer resource bundle to be purchased for a first container in the first container system using virtual currency units, wherein the first container in the first container system is a software implementation of a computer that runs on top of a shared operating system and executes one or more programs or applications on the shared operating system; allocating the computer resource bundle from a first computer server in the first container system to the first container; receiving, from a proxy manager of a second container system offering the computer resource bundle, a purchase price for the determined computer resource bundle in virtual currency units, wherein the second container system also includes at least one container that runs on top of a second shared operating system and executes one or more programs or applications on the second shared operating system; determining whether the determined computer resource bundle is to be purchased from a second computer server in the second container system based at least in part on the purchase price received from the proxy manager of the second container system; allocating the computer resource bundle from the second

15

computer server in the second container system to the first container; and dispatching the first container from the first container system to execute at the second container system.

In some embodiments, the method further includes initiating, following an increase in the purchase price for the allocated computer resource bundle, a credit check process by the container manager, wherein the credit check process results in an increase in budget of virtual currency units available for usage by the first container.

Another aspect of the subject matter described in this specification can be embodied in methods and systems that include allocating one or more computer resources to a plurality of containers in a container system, wherein each container is a software implementation of a computer that runs on top of a shared operating system and executes one or more applications on the shared operating system, and wherein the plurality of containers operate as components of a cluster for running related applications; determining, by a container manager running on a data processor of the container system, a target latency for the plurality of containers; determining, by the container manager, an actual latency for the plurality of containers; comparing the target and actual latency for the plurality of containers; and determining that the plurality of containers is to be moved from a first computer server to a second computer server based at least in part on the comparison of the target and actual latency.

In some embodiments, the method further includes initiating a migration of the plurality of containers from the first computer server to the second computer server.

In some embodiments, initiating comprises an atomic migration of the plurality of containers.

Additional details of one or more embodiments of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the descriptions contained herein, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of an example container environment in which resources are managed.

FIG. 2 is a block diagram of an example software system for managing resources in a container system.

FIG. 3 is a flow diagram of an example process for using a platform manager in a container system.

FIG. 4 is an example model for service provision and consumption in a supply chain container system.

FIG. 5 is a flow diagram of an example process for deploying a new consumer element with a provider element in a container system.

FIG. 6 is a flow diagram of an example process for delivering service level agreement targets through resource allocation in a container system.

FIG. 7 is a flow diagram of an example process for economic-based I/O scheduling in a container system.

FIG. 8A is an example purchase order data structure for use in purchasing services from a provider element manager in a container system.

FIG. 8B is an example service confirmation data structure for use in confirming or rejecting the purchase of services from a provider element manager in a container system.

FIG. 9 is an example process for managing the states of system elements in a container system.

16

FIG. 10 is a block diagram of an example multi-domain software system environment for managing virtualized resources.

FIG. 11 is a block diagram of an example virtualization environment which illustrates supply chain relationships between service entities and resources.

FIG. 12 is a block diagram of another example virtualization environment which illustrates supply chain relationships between service entities and resources in a container system.

FIG. 13 is a flow chart illustrating a process for resource scaling in the virtualization environment of FIG. 11.

FIG. 14 is a flow chart illustrating a process for service entity scaling in the virtualization environment of FIG. 11.

FIG. 15 is a block diagram of yet another example virtualization environment which illustrates the supply chain relationships between service entities and resources in a virtualization environment and can cooperate with the processes described in FIGS. 13-14.

FIG. 16 is a block diagram of an example virtualization environment in which resources are managed by an action manager.

FIG. 17 is a block diagram illustrating the data flow for managing resources in the virtualization environment of FIG. 16.

FIG. 18 illustrates an exemplary block diagram of a virtualization environment in which a virtual machine is determining whether to take an action.

Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

FIG. 1 is a block diagram of an example container system or environment 100 in which resources are managed. The example system 100 includes two servers 102 and 104 that run respective container systems 110 and 112. The container system 110, at the server 102, allocates computer resources (or compute resources) of the server 102—e.g., CPUs, memories, storage volume, storage, and/or network I/O pathways and bandwidth—to two containers 120 and 122. Similarly, the container system 112 at the server 104 allocates resources of the server 104 to containers 124 and 126. The containers 120, 122, 124 and 126 execute respective containerized applications 130, 132, 134, and 136.

As previously discussed, container systems admit flexible organizations. In the example system 100, the servers 102 and 104 may be physical machines with physical computer resources (or compute resources). Alternatively, the server 102 may be a virtual machine with virtualized resources while the server 104 is a physical server. The containers 120, 122, 124 and 126 may be distinct containers, or replicated copies of a single container. In some embodiments, a group of containers may be clustered into a container-Point-of-Delivery (cPOD) system, to run related applications. For example, a multi-tier Web service may include a containerized Web server (shown as the application 130), a containerized application server (shown as the application 134), and a containerized database server (shown as the application 136). The Web server provided by the application 130 can maintain significant level of communications with the application server provided by the application 134. The I/O pathway between the applications 130, 134 traverses the application 130, the container 120, the container system 110, an operating system 106, a network interface card (NIC)

17

140, a data network 160, a NIC 142, an operating system 108, the container system 112, the container 124, and the application 134.

In this example, the portion of the aforementioned I/O pathway that includes the NIC 140, the data network 160, and the NIC 142 traverses network switches and links, and can thus result in significant I/O latency as well as bandwidth limitations. A container manager 236—considered below and shown, for example, in FIG. 2—can migrate (ship) the container 120, with the application 130, from the server 102 to the server 104. This migration replaces the I/O pathway from the application 130 to the application 134 with a pathway that includes the application 130, the container 120, the container system 112, the operating system 108, the container system 112, the container 124, and the application 134. Advantageously, this modified I/O pathway entirely can be handled by the server 104 through memory transfers. This in-memory I/O pathway can support very high memory transfers bandwidth and very low latency, thus, improving the cPOD performance.

Although a specific environment 100 including the two servers 102 and 104 is shown in FIG. 1 and described above, it will be understood that the environment 100 is illustrative only. For example, the environment 100 may include more than two servers, and each of the servers 102 and 104 may be associated with any number of containers as desired.

The principles described herein may be applied regardless of the particular application or applications being run in the container system.

FIG. 2 is a block diagram of an example software system 200 for managing resources in container systems, such as the container system 100. According to various embodiments, the software system 200 may be used to allocate server and I/O resources such as CPU, memory, flash storage, hard drive storage and I/O bandwidth) to containers. The software system 200 also may be used, for example, to monitor, detect, and handle congestion conditions at a resource (e.g., I/O pathway, memory, and so on) and to move containers among available servers to optimize or improve application performance and resource utilization.

The software system 200 monitors, controls, and otherwise interacts with various managed container system elements (also referred to herein as service elements or computer elements) through respective instrumentation. As used herein, in the context of computers, the term “instrumentation” refers generally to any software and/or hardware that provides an ability to monitor, control, or otherwise interact with a computer element, such as to detect operations events, reconfigure parameters, diagnose errors, and write trace information. For example, when a computer application contains instrumentation code, the computer application may be managed using a management tool.

Several example container system elements are shown in FIG. 2 as part of an Information Technology (IT) Container Stack (ITCS) 202, including applications components 210, container systems 212, servers 214, storage systems 216, networks 218, and operating resources 220 (such as power supplies, cooling systems, and rack space). In some embodiments, the ITCS 202 may include, for example, a proper subset or a proper superset of these container system elements 210, 212, 214, 216, 218, and 220.

As shown, the software system 200 includes a platform layer 230, which provides an infrastructure to manage, for example, the I/O flows in a container system (such as the example container system environment 100 shown in FIG. 1). The platform layer 230 includes element managers 234, 236, 238, 240, 242, and 244. More particularly, the platform

18

layer 230 includes an application manager 234, a container system manager 236, a server manager 238, a storage manager 240, a network manager 242, and an operations manager 244. These element managers 234, 236, 238, 240, 242, and 244 use management instrumentation of respective elements to monitor and control the respective elements of the ITCS 202.

For example, the server manager 238 may use built-in management instrumentation, such as Management Information Bases (MIBs) of the server it is managing, to monitor the server's CPU, memory, and I/O interfaces (such as a Host Bus Adapter (HBA) and NICs) and to control their operational parameters. The server manager 238 may access such management instrumentation using standardized protocols (such as Simple Network Management Protocol (SNMP)) or specialized mechanisms. In some embodiments, a proper superset or only a proper subset of these element managers 234, 236, 238, 240, 242, and 244 may be desired or needed in certain environments. For example, when the containers do not access storage, the use of a storage manager 240 may not be needed. Additionally, for example, an operating system element manager (not shown) may be included as part of platform layer 230.

As also shown, the platform layer 230 also includes one or more types of modeling databases 245. As discussed in more detail below, the databases 245 may include supply chain modeling (SCM) databases 246 and operations databases 248. The platform layer 230 also includes a platform manager 250, which, as explained in greater detail below, can be responsible for general provisioning, initializing, and management tasks.

The software system 200 shown in FIG. 2 also includes a functional management layer 252, which includes user interface (UI) software 260 for use by administrators or other users to monitor and control a container system (such as the example container system environment 100 shown in FIG. 1). For example, an administrator may use the UI software 260 to set proactive automation policies to optimize or improve performance and resource utilization, detect and resolve operational problems and performance bottlenecks, allocate priorities and usage charges to different applications, and plan capacity expansions.

The functional management layer 252 also includes a collection of functional managers 272, 274, 276, and 278, which are used to enable users to monitor, control, and automate the underlying automated management mechanisms of container systems according to the principles described herein. The software system 200 may alternatively include, for example, a proper subset or a proper superset of these functional managers.

As shown in FIG. 2, the functional management layer 252 includes an application manager 272, which, for example, enables users to select or configure respective parameters of a computer agent or process to partition application components among different containers, allocates virtual budgets to applications based on the business value of their services, as described in greater detail below, and specifies the resources required by the applications. The application manager 272 uses the parameters to create respective records in the operations databases 248. The platform manager 250 uses the operations records to initialize respective application managers 234, which use the operations records to deploy the applications 210, according to the principles described below. Additional functions of monitoring and controlling applications may be incorporated into the application manager 272.

The functional management layer **252** also includes a performance manager **274**, which allows users to monitor and control the delivery of Service Level Agreements (SLAs) to applications. For example, a user of the software system **200** can specify target SLA parameters—such as latency or transaction rate—of one or more particular applications. The SLA parameters are used by the software system **200** to adjust the performance of the applications using the principles described below. A user can also monitor the SLA parameters value, as well as the respective virtual payments made by an application, thereby correlating the application's budget with its SLA performance. Additional functions of monitoring and controlling the performance of applications, as well as the other elements of the ITCS **202**, may be incorporated into the performance manager **274**.

A capacity manager **276** monitors relationships between the supply and demand of resources in the ITCS **202**. For example, the capacity manager **276** may monitor the relationships over a predetermined time period, which can range from short term (such as a few minutes or one hour) to long term (such as one day, week, month or year). In some embodiments, the capacity manager **276** maintains full accounting of revenues and costs and provides monitoring of these accounts and notifications upon certain accounting events. The capacity manager **276**, by itself or with the assistance of an incorporated or separate Return-on-Investment (ROI) manager (not shown), enables a user to monitor the ROI of the elements in the ITCS **202**. The ROI is defined as revenue divided by cost, where revenue is the income from virtual payment collected by a selected element and cost is the virtual payments by the element for the resources that the element uses.

For example, a large ROI may indicate to the capacity manager **276** that there is excess demand over supply of the element capacity, and a sustained high ROI may thus indicate insufficient capacity. The capacity manager **276** compares a monitored ROI with specific and potentially predetermined ROI targets, which may be configured by an administrator or other user, to recommend capacity increases of particular elements to meet demand. According to the supply chain economic principles described below, the ROI of an element in the ITCS **202** may be considered as a central metric of economic value.

The ROI may be calculated at any appropriate time and for any appropriate duration over which revenue and cost are considered. Thus, the principles described herein provide an accounting framework to quantify and measure the value generated by components of the ITCS **202**. For example, at the bottom of the ITCS **202**, there are raw resources that generate real (non-virtual) costs, such as monetary costs that are paid to an electric company. At the top of the ITCS **202**, there are applications that play roles in generating real (non-virtual) revenues, such as monetary sales revenue received from customers. It is possible to treat one or more of the system elements **210**, **212**, **214**, **216**, **218**, and **220** as virtual Profit and Loss (P&L) entities, generating revenues through payments by its consumers, and paying the costs of services it consumes. The use of virtual currency pricing and payments, as described herein, to distribute a share of these revenues to cover costs increases the efficiency and overall ROI of the entire system.

A back-charging manager **278** monitors and accounts for the virtual cash flows between elements in the ITCS **202** and enables users to flexibly compute financial metrics of interest. For example, users can monitor metrics describing the allocation of application budgets to acquire supply chain

resources, the allocation of a resource among the workloads of different applications, the ROI efficiency of different resources, and application budgets required to deliver particular SLAs. These metrics and other parameters may be used to support policies on budgeting applications, adjusting the budgets to represent changing prices, capacity, and demand of resources along the supply chain, and converting virtual currency used within the software system **200** to real currency (such as United States dollars, or euros) that is generated by the business units who own the applications and that may be used to pay for IT resources.

The platform manager **250** can manage a container system using any suitable means described herein, including using a process **300** as shown in FIG. 3. Turning to FIG. 3, a flow diagram of the example process **300** for using the platform manager **250** in a container system (such as the container system **100**) is shown. According to various embodiments which implement process **300**, the platform manager **250** initializes, or launches, the functional managers **272**, **274**, **276** and **278** of the functional management layer **252** for a specific container environment (step **302**). The platform manager **250** discovers the managed container system elements of the ITCS **202** in the container environment (step **304**). This discovery is handled, for example, through standard processes to get configuration data from the container system, OS, server, network, and storage systems.

The platform manager **250** also initializes, or launches, an element manager (such as one or more of element managers **234**, **236**, **238**, **240**, **242**, and **244**, described above) for each group of respective elements of a given class of elements that have been discovered (step **306**). For example, the platform manager **250** may detect a Dell® server and a SUN® server, and the corresponding groups of respective elements may both be assigned respective element managers. The platform manager **250** configures the element managers to monitor and control the respective elements via respective management instrumentation.

The platform manager **250** populates and initializes the platform modeling databases **245**—for example, the supply chain modeling databases **246** and the operational databases **248** (step **308**)—and starts monitoring certain potential changes of the managed environment (step **310**). For example, the container system **100** may be monitored to determine if there have been any container changes, such as any added, deleted, or migrated container (decision block **312**). If a container change has been detected, the platform manager **250** again initializes the element managers as described above.

If no container changes have been detected, the presence of container systems is evaluated to determine if there have been any container system changes, such as any added or deleted container system (decision block **314**). If a container system change has been detected, the platform manager **250** again discovers the managed container system elements of the ITCS **202** in the container environment as described above. Otherwise, the platform manager **250** evaluates whether there have been any major network changes (decision block **316**), in which case the platform manager **250** similarly re-discovers the managed container system elements of the ITCS **202** in the container environment as described above. For example, the platform manager **250** may discover loss or gain of network I/O pathways, congestion or under-utilization of an I/O pathway, low or excessive latency of an I/O pathway, or packet losses along an I/O pathway. Otherwise, the platform manager **250** evaluates whether there have been any major storage changes (decision block **318**). For example, the platform

21

manager **250** may discover storage I/O congestion, or alternate I/O pathways that would provide better (i.e., lower) access latency. If major storage changes have been detected, the platform manager **250** again discovers the managed container system elements of the ITCS **202** in the container environment as described above.

If no container, container system, network, or storage changes have been detected, the platform manager **250** determines whether to continue monitoring of the same (decision block **320**). If the platform manager **250** decides to continue monitoring, the platform manager **250** again starts the monitoring of potential changes of the managed environment. Otherwise, the process **300** ends (end block **322**).

The order of steps in the example process **300** described above is for illustration purposes only, and can be done in different orders. For example, the platform manager **250** may evaluate whether there has been any major storage changes (decision block **318**) before determining whether there has been any major network changes (decision block **316**). Moreover, additional steps may be included, for example, to protect the software system **200** against its own failures. Such additional steps may include, for example, inserting between steps **308** and **310** described above the steps (not shown) of creating a mirror and backup copies of the platform image (including the databases **246** and **248**), running a second instance of the software system **200** in standby mode and monitoring the primary instance of the software system **200**, and switching to the standby instance of the software system **200** upon detecting the failure of the first instance of the software system **200**.

According to various embodiments, the software system **200** described above can operate using a supply chain software model of the ITCS **202** that it manages. In other words, each container system element **210**, **212**, **214**, **216**, **218**, and **220** of the ITCS **202** is modeled as a provider and a consumer of services. For example, FIG. **4** is an example model **400** for service provision and consumption in a supply chain container environment. According to various embodiments as shown in FIG. **4**, which includes references to the container system elements of the ITCS **202** shown in FIG. **2**, the server **214** may consume services of the operating resources **220**, including, for example, power **402**, cooling **404**, physical space **406**, a share of capital expenditure (CAPEX) costs **408**, and a share of operating expenditure (OPEX) costs **410**. The server **214** further may consume the resources of the networks **218**, including, for example, a local area network (LAN) **420** and a storage area network (SAN) **422**.

However, the server **214** may provide the container systems **212** with various physical resource services, including, for example, CPU bandwidth **430**, memory **432**, network I/O bandwidth **434**, and storage I/O bandwidth **436**. The container systems **212** may also consume storage resources **438** from the storage element **216**, and, in turn, may offer services (such as services **440** and **442**) to the application **210**. The application **210**, on the other hand, may offer services to respective business activities of one or more business units **450**.

According to various embodiments, the allocation of resources and the processing of workloads through the supply chain, as described above, may be performed through the use of virtual currency. In these cases, supply chain elements use virtual currency to pay for the services they consume and to price the services they offer. For example, a selected application **210** may receive a budget from its business users reflecting the business value of the services that it offers. The application **210** may shop for a container

22

system **212** that offers the lowest priced processing services that the application **210** requires, and may use its virtual budget to pay for these services. The container system **212**, in turn, may use its income of virtual currency to pay for the services offered by the server **214**, the network **218**, and the storage system **216**. Each of the container systems elements **210**, **212**, **214**, **216**, **218**, and **220** of the ITCS **202** may price their services in virtual currency to reflect their costs, and additionally, or alternatively, to balance supply and demand.

According to various embodiments, resource pricing may also be based one or both of capacity or performance characteristics. For example, the server **214** may offer multiple types of processors or CPUs, each with respective clock rates and other characteristics, at different prices. Similarly, for example, storage I/O resources in the storage system **216** and network I/O resources in the network **218** may be priced according to their bandwidth and latency characteristics. This manner of pricing can take into account that, as noted above, I/O pathways internal to a server (i.e., interconnections of containers co-located with a single server, e.g., the containers **120** and **122** as shown in FIG. **1**) typically offer higher bandwidth and lower latency than I/O pathways between containers located at different and distinct servers (e.g., the containers **120** and **124** as shown in FIG. **1**). Thus, for example, one or more of the components and resources associated with internal I/O pathways (or the aggregate of such components and resources) may be priced lower than components and resources (alone or in the aggregate) for pathways traversing switches and/or involving multiple servers. Alternatively, for example, components and resources associated with such internal I/O pathways may be priced higher to account for an expected increase in performance and thus value to the acquiring entity.

The supply chain model of the ITCS **202** is primarily maintained by the supply chain model databases **246** shown in FIG. **2**. According to various embodiments, the supply chain model databases **246** may include one or more financial databases to debit and credit the respective accounts of customers and providers to reflect the transfer of virtual payments, as discussed in greater detail below. It will be understood, however, that non-monetary transactions may be entered into between a consumer and a provider.

The supply chain model databases **246** may be object-relationship databases, such that elements of the supply chain are modeled as objects corresponding to services to be offered. As used herein, the term “objects” refers to data structures including data fields and methods. Examples of service objects include simple and composite service objects.

According to various embodiments, simple service objects—or objects relating to the provision of a single type of service—may include the following types of attributes:

<service-identifier, units, used, available, duration, price(>).

The “service-identifier” attribute may itself include the following types of attributes as descriptors of the service that may be used for a particular class of services: <name, type, description, element manager>. For example, a CPU service provided by a Dell® server with an Intel iQ9550® processor managed by an element manager ServerEM015 may be assigned the following identifier: <Dell4, CPU, iQ9550, ServerEM015>.

The “units” attribute may measure the quantity of service, such as 5 Mhz (CPU), 2 GB (memory) or 10 Mbps (net I/O).

The “used” attribute may refer to the amount of the service or resource capacity that is already committed.

The “available” attribute may refer to the amount that remains to meet new demands.

The “duration” attribute may indicate the period of time over which service is to be rendered.

The “price(demand)” attribute may refer to a method whose input is the demand by a service consumer, for a number of service units it requires, which computes the price in virtual currency units, as set by the service provider. For example, the simple service object <<Dell4, CPU, iQ9550, ServerEM015>, 0.1 Ghz, 0.8 Ghz, 2 Ghz, 1 hr, price(x)>, where $\text{price}(x)=1/(2-0.1x)^2$, may be used to describe a CPU service named Dell4, providing an Intel processor of type Q9550 for one hour in units of 0.1 Ghz. In this case, a request for 0.5 Ghz (5 units) of this CPU service will be priced at $\text{price}(5)=1/2.25=\$0.44$ per one hour of use.

According to various embodiments, the pricing functions used by simple service objects can be flexibly adapted by element managers to reflect different pricing goals and mechanisms. For example, a server may be shared by 10-100 containers, which preferably utilize no more than 50% of its capacity to avoid congestion. In this case, the percentage of average demand to capacity of a given server resource preferably falls between 0.5%-5%.

Consider a commodity service, defined as one where this ratio is very small. With supply far exceeding demand, prices will drop to reflect costs. Thus, a commodity service may be priced at fixed cost-based price. For example, suppose the percentage of average demand to capacity for CPU usage by a container is 0.2%. In such a scenario, the shifting of a container among servers would have negligible impact on the quality of CPU services seen by the containers. CPUs can therefore be priced at a fixed level to merely reflect the costs of providing CPUs. In general, a commodity service may be priced at a fixed level, independently of demand. However, when the ratio of average demand to capacity is sufficiently large, arriving demands may easily deplete the supply absent pricing control, thus requiring higher prices to balance the supply and demand.

A sample pricing function that provides such pricing control is:

$$\text{price}[x]=\text{cost}/(1-(U+x)/C)^4$$

where

C=capacity of the resource;

U=amount of resource used; and

x=new demand.

Such a pricing function is proportional to costs, penalizing high utilization. When the utilization $u=(U+x)/C$ approaches its limit of one, prices increase rapidly, preventing all but the highest budget applications from accessing the resource. For example, suppose containers require, on average, 2% of the CPU capacity of servers, but 20% of their storage I/O capacity. In this scenario, a container wanting to deploy with a server supporting three containers will see the following CPU and storage I/O prices:

$$\begin{aligned}\text{price}_{CPU}[0.02C]&=\text{cost}_{CPU}/(1-0.08C/C)^4=\text{cost}_{CPU}/0.92^4=1.4*\text{cost}_{CPU} \\ \text{price}_{I/O}[0.2C]&=\text{cost}_{I/O}/(1-0.8C/C)^4=\text{cost}_{I/O}/0.2^4=625*\text{cost}_{I/O}.\end{aligned}$$

Thus, in the above-described scenario, CPU is priced at a relatively small multiplier of the cost base of CPU, while the storage I/O is priced at a relatively large multiplier of the cost base of I/O. Although specific pricing considerations and mechanisms have been described, a large variety of pricing functions may be used according to other embodiments to best reflect specific use considerations.

Composite service objects, which are objects that include more than one service object and which relate to the provision of multiple types of services, may take the following form according to various embodiments:

<service-identifier, service-1, service-2 . . . , service-n>,

where service-k is either a simple or composite service object and is referred to as a component of the composite service. In some embodiments, the “duration” attributes of all components of a composite service are identical, and their common value is called the duration of the composite service. For example, a hardware server may be described by the following composite service object:

<<server-1, Server, LS41>, CPU4, Memory-2, NIC-3, NIC-4, HBA-2>

where Memory-2, NIC-3, NIC-4 and HBA-2 indicate respective simple service objects associated with respective memory services, LAN-interface services provided by two NICs, and SAN I/O services provided by HBA-2. The HBA-2 may itself be described by a simple service object as follows:

<<HBA-2, FC-HBA, Emulex, LP11000-M4>, 0.1 Gbps, 1.1 Gbps, 2.9 Gbps, 1 hr, price(x)>.

This service object indicates that the duration of the composite service is one hour, as the durations of all components of a composite service are identical.

In some embodiments, the price of a composite service is defined as the sum of the prices of all its components. For example, the price of a server object is the sum of the prices of the units of CPU, memory, network I/O and storage I/O required by a consumer.

The supply chain model databases 246 are maintained by element managers (such as element managers 234, 236, 238, 240, 242, and 244 shown in FIG. 2), which handle the service objects corresponding to the respective elements that they manage. As explained above with respect to the sample process 300 shown in FIG. 3, according to various embodiments, an element manager is initialized by the platform manager 250, and subsequently the element manager proceeds to populate the supply chain model databases 246 with respective service objects it is responsible for. Once the supply chain model databases 246 have been updated, the element manager continues to update the dynamic attributes of its respective service objects (such as the “used” and “available” attributes). For example, a server manager 238 that is responsible for managing HBA resources will initialize the supply chain model databases 246 with corresponding simple service objects relating to the HBA. The server manager 238 will then monitor and update the “used” and “available” attributes of this simple service object by periodically accessing the HBA instrumentation.

As mentioned above, the supply chain economy matches consumers and providers of resources or services by using pricing and budgeting. According to various embodiments, demand for services is matched to supply through a shopping model. A consumer element manager (such as one of element managers 234, 236, 238, 240, 242, and 244 shown in FIG. 2), desiring services from a provider element manager, queries the supply chain model databases 246 in search of the best priced provider or providers of the desired services. The query specifies requirements and the service or services the element manager is requesting. For example, a query may take the following form:

Query: Server, CPU.units=50 Mhz, Memory.units=4 GB, StorageI/O.units=200 Mbps, NetworkI/O.units=100 Mbps.

Such a query may retrieve records of composite service objects of the servers 214 offering the respective CPU, memory, storage I/O and network I/O capacity at the lowest

25

price. Once the consumer element manager acquires these records of lowest-priced service objects, it can proceed to extract the identities of the element managers posting these service offerings. The consumer element manager may then pursue direct interactions and contract with one or more
 5 respective provider element managers to acquire and pay for the desired services. There exists the possibility that multiple consumers may query the supply chain model databases **246** simultaneously for similar services, and thus potentially interfere with each other's shopping processes. Such interference may be avoided, for example, by providing standard locking mechanisms to maintain atomicity of the query and purchase transactions.

Moreover, various embodiments may use an auction, or bidding model, rather than a shopping model, to match demand and supply. For example, consumer element managers may post respective bids for services in a bidding database, which a provider element manager may then query for the highest bid price offered for its services and contract to serve it. The shopping model is generally preferred to bidding in situations where consumers' demands arrive asynchronously and unpredictably. In such cases, an arriving consumer can find the low-cost provider by searching the supply chain model databases **246**. In contrast, a bidding process requires providers to poll, whether constantly or at intervals, the bidding database to detect arrivals of new bids, while bidding consumers may be required to wait until enough providers have polled the bidding database and accepted the bids, and thus contract with providers based at least in part on chance. There are various situations where bidding may offer benefits over shopping, and those situations may be handled using the principles described herein.

FIG. 5 is a flow diagram of an example process **500** for deploying a new consumer element (such as a container) with a provider element (such as a server) in a container system that is used according to various embodiments for balancing the demand and supply of services. According to various embodiments, the dynamic load balancing approach illustrated by example process **500** provides an effective solution to several of the resource management problems described above. For example, process **500** may be used to improve the balancing of demands by containers and the supply of server resources; it may also be used to balance the resource bundle allocated to a container, e.g., to match the amount of CPU, memory and storage I/O bandwidth allocated to the container, in order to improve the use of its virtual budget to best service its resource demands.

As shown in FIG. 5, once the relevant consumer element managers and provider element managers are running, having been initiated by the platform manager **250**, a consumer element manager shops for lowest cost provider for a bundle of services by querying the supply chain model databases **246** as described above (step **502**), and contacts the provider element manager to buy services (step **504**). In the case of a container consumer, for example, the bundle of services to be purchased may include CPU, memory, and storage I/O.

The provider element manager determines whether the consumer budget is sufficient to pay the price for the requested provider services (decision block **506**). If it is determined that there is sufficient budget, the provider element manager deploys the consumer at the provider, which proceeds to process its workload (step **508**). For example, CPU and memory resources that have been purchased may be allocated to a container by the underlying scheduler of the container system, which may include the use of a traditional operating systems scheduling algorithm. The server element manager configures the scheduler

26

parameters to accomplish fairly accurate allocation of the CPU and memory. Memory may be allocated by specifying an amount of memory to be provided. The container system can allocate physical memory, based on these specifications, or support virtual memory mechanisms that permit over
 5 100% utilization of physical memory. Additionally, the CPU may be allocated by configuring reservations and shares parameters of the scheduler. For example, reservations may be used to allocate a reserved CPU slice, using a time-shared round-robin scheduler, while shares allocate the remaining CPU bandwidth through a Weighted Fair Queuing scheduler. CPU reservations and shares may be viewed as separate services, and may be individually priced according to supply and demand. For example, a low-priority application may be
 15 unable to buy reservations, and may thus need to settle for shares, which may be priced lower. A high-priority, mission-critical application, on the other hand, may have sufficient budget to afford sufficient reservations to support its needs.

Otherwise, if it is determined that there is not sufficient budget, the consumer element manager initiates a credit check process to decide whether the consumer can increase its budget or sufficiently lower its service demands, and thus continue to run (decision block **510**). For example, suppose the consumer is a container whose budget is short of paying the cost of a provider server. In that case, the container may use credit it has accumulated to pay for the service, obtain additional budget from the applications it serves, or reduce its demand for services and the corresponding price to the point where it can afford to pay. If one or more of these scenarios is possible, the consumer uses credit, increases its budget and/or lowers its service demands (step **512**), and the provider element manager thus deploys the consumer at the provider as described above. Otherwise, if none of these options is available, the consumer is suspended and then will either terminate or re-launch when adequate budget becomes available to it (step **514**), as described in greater detail below.

After the provider element manager deploys the consumer at the provider, the provider element manager or the consumer element manager monitors consumer resource usage and adjusts allocation of resources to optimize or improve the use of the consumer's budget (step **516**). For example, the provider element manager may find that the consumer is using only 20% of one service it bought, while using 90% of another service it bought. In that case, the provider element manager may reduce the allocation of the first service and use the corresponding released budget to increase the allocation of the second resource.

Upon completion or termination of the consumer service period, the provider element manager notifies the consumer element manager (step **518**), which may proceed to shop for a new provider offering lowest cost services to meet the consumer's needs (step **520**). The consumer element manager determines whether the price of the new provider found is lower than the price of the old provider (where the consumer resides at the time), or according to some embodiments, whether it is lower by a threshold amount (decision block **522**). Assuming it is, the consumer element manager moves the consumer to the new provider, in which case it may also adjust the budget to reflect the price of moving, if any (step **524**). Namely, according to various embodiments, a price of moving may be factored into the decision making process for whether the consumer should be moved to the new provider, and such price may be subtracted or deducted from the available budget. Otherwise, if the consumer element manager decides to keep the consumer with the old provider, it does not adjust the budget to reflect the price of

moving. In either case, the provider element manager (of the new or old provider) checks to see if the consumer budget is sufficient to pay for the provider as described above.

According to various embodiments, the process of shopping for a new provider **520** may depend on specific characteristics of the consumer, the resource, and/or the provider. For example, the containers **120** and **124** shown in FIG. **1** may need to exchange high-bandwidth latency-sensitive communications through a congested switch in the network **160**. Further to the discussion above, internal I/O pathways (including at either the server **102** or the server **104**) may offer higher bandwidth and lower latency, and thus result in improved performance. Therefore, according to various embodiments, such internal I/O pathways may be priced lower than I/O pathways involving, for example, multiple servers **102** and **104** and network **160**.

As an example, in the step **520** described above and shown in FIG. **5**, the consumer element manager may determine that it would be more economical or efficient to move a consumer element from the server **102** to the server **104** based on reduced I/O pathway pricing. For example, the consumer element manager may discover that the container **120** should be moved to the server **104** to obtain one or more resources and communicate with one or more other elements located at the server **104**. This can be the case where, for example, it is determined at the step **522** that the overall price of providing container **120** with necessary resources is reduced at least in part because of a lower price of the I/O pathway should container **120** be moved to server **104**. In that case, at step **524**, the container **120** may be moved to server **104** so that the I/O pathway becomes more (or entirely) local to server **104**, thus benefiting from higher expected bandwidth capacity and lower latency.

According to various embodiments, at step **524**, the budget of the consumer element (e.g., container **120**) may also be adjusted (e.g., increased or decreased) based at least in part in such change in pricing. As indicated above, in an alternative embodiment, the pricing of resources (e.g., associated with the I/O pathway) may be increased to account for performance improvement that would result from movement of a consumer element to another server and the resulting localization.

According to other embodiments, the process of shopping for a new provider **520** may depend on functional characteristics of the consumer or provider. For example, the server **102** may be used to support development of containerized applications. The server **104**—the provider, for example—may be used for testing the containerized application **130**—the consumer, in this example. The process **500** may be used to select a new provider (the server **104**), from among a group of servers providing rests of containerized applications, to run tests (consumer) of the containerized application **130**. Similarly, the server **104** may be a production system running containerized applications and the process **500** may be used to dispatch the containerized application **130**, and its container **120**, from the development server **102** to the production server **104**.

The order of steps in the example process **500** described above is illustrative only, and can be done in different orders. Moreover, it is contemplated that modifications and extensions of the process **500** will be used according to various embodiments. For example, a consumer may need to contract with two or more providers to be deployed, as in the case of a container that needs to acquire a bundle of resources offered by a server as well as SAN switch bandwidth and storage space at a storage array. In such scenarios, deployment of the consumer can be supported by extending

step **502** to shop for multiple providers and then repeating the remaining steps for each of these providers. Additionally, for example, as explained below with respect to FIG. **6**, the example process **500** shown in FIG. **5** may be modified or extended to enable the adjustment of resource allocations to obtain desired service level agreements (SLAs).

According to various embodiments, the above-described supply chain economic principles may also be used to manage software licenses, such as temporary (time-limited) software licenses. For example, regardless of type (such as authorizations of software use per user, per CPU, per server, or per container), licenses may be modeled as resources to be purchased by an application manager **234**, much like other resources that it may purchase from the container **212**. License element managers (while not shown, may be included as part of the platform layer **230**) may be used to set the prices of the licenses based on costs and demands. In this manner, license management may be greatly simplified and unified with the allocation of other types of resources. For example, an application that is unable to acquire a needed license may suspend its operations and release its resources, as explained below, thus increasing the overall efficiency of the system. Additionally, licenses may be more efficiently used, since in situations where the licenses are highly utilized, they will be allocated to high priority tasks, while lower priority tasks may be suspended until they can afford the licenses. As soon as a license is no longer needed, it may be released and available for other tasks. Additionally, an administrator may consider the ROI of licenses, as with other resources, to plan the expansion, or contraction, of licenses capacity. For example, if a license's ROI is above a certain threshold, it may be desirable to acquire more licenses to increase the supply to meet demand.

FIG. **6** is a flow diagram of an example process **600** for delivering service level agreement targets through resource allocation in a container system, which includes many of the steps of process **500** shown in FIG. **5** and discussed above. Although not required, for the purpose of simplifying the following description, it is assumed that the target service level agreement relates to an application running on a container. However, the service level of other types of computer elements may be controlled in the following manner according to various embodiments.

Following the initial monitoring of resource utilization and optimizing of the container's budget (step **516**), it is determined whether the consumer service period has terminated (decision block **602**), in which case the provider element manager notifies the container element manager (step **518**) as described above. Otherwise, the container element manager monitors and obtains the value of the SLA parameter of interest, such as the average transaction rate of an application, the average transaction delay of an application, the average communications latency of the application, or the number of transactions performed within a predetermined prior time period by an application (step **604**). For example, an application element manager may monitor the value of the SLA parameter, through respective instrumentation, and inform the container element manager of the SLA parameter. The application may define its SLA goal as 100 transactions per second, in which case the SLA parameter of interest is transaction-rate. In general, because SLA parameters can be assumed to increase monotonically with the amount of resources allocated to an application, the management of SLAs may be accomplished as described herein by finding a budget and a respective resource allocation that will accomplish the target SLA value.

The container element manager determines whether the SLA parameter of interest is below a desired target (decision block **606**), in which case, for example, the application's payments to the container (e.g., of virtual currency units) are increased such that the container's budget is increased, and it is able to purchase more resources to increase the SLA parameter of the application (step **608**). After such an increase, the container's budget use is again monitored and optimized or improved as described above.

If the container manager determines that the SLA parameter is at or above the desired target, it is determined whether the SLA parameter exceeds the desired target by more than an acceptable threshold (decision block **610**), in which case the payments are reduced, thus reducing the container's budget and the resources it buys, saving on applications costs, and keeping the SLA performance within a desired tolerance range (step **612**). After such a reduction, the container's budget use is again monitored and optimized or improved as described above. If the SLA parameter is within the acceptable range, however, a reduction is not applied, and the process is repeated until it is determined that the consumer service period has been completed or terminated.

According to various embodiments, the process **600** for delivering service level agreement targets through resource allocation in a container system may be modified, adapted, and/or simplified for certain resources and SLA metrics. For example, in the case of allocation of I/O pathways to reduce or minimize latency, the process **600** may be modified as follows. The SLA parameter may be selected as the latency-hop-count, e.g., the number of physical switches traversed by an I/O pathway. For example, I/O pathways between elements located, or resident, at the same server (e.g., the containers **120** and **122** in FIG. 1) generally do not traverse any physical switch, and thus may be described as having a latency-hop-count of 0. Such I/O pathways may also be referred to as having Class-0 Latency SLA. On the other hand, I/O pathways between elements located or resident at different servers (e.g., the containers **120** and **124** in FIG. 1) and attached to a common switch (e.g., a common switch of the network **160**) may be described as having a latency-hop-count of 1, and may be referred to as having Class-1 Latency SLA. According to various embodiments, an I/O pathway may involve two or more physical switches, and may be described as having a latency-hop-count of 2 (or more) and referred to, for example, as having Class-2 Latency SLA.

According to various embodiments, the latency-hop-count associated SLA value may be described with respect to the ordinal preference {Class-0, Class-1, Class-2, . . . Class-n}, where Class-0 is preferred to Class-1, Class-1 is preferred to Class-2, and so on to the extent additional Classes are defined. With respect to the process **600**, a comparison can be made between a Target Latency Class and an Actual Latency Class (e.g., Target=Class-0, Actual=Class-1) at step **606**. If the Actual Latency Class does not meet the Target Latency Class, payments to the consumer (e.g., the container) may be increased at step **608**, and, following return to step **516**, an I/O pathway can be acquired that can deliver the Target SLA Value (e.g., Class-0). For example, the process **600** described with respect to FIG. 6 can be modified in a manner consistent with the above description so as to simplify the monitoring and control of SLA values to classification of the I/O pathway into latency class.

It will be understood that the SLA-delivery process **600** described above may be flexibly adapted to achieve various goals, such as improving its handling of stochastic fluctuations of an SLA parameter. For example, the steps of

increasing (step **608**) and decreasing (step **612**) payments by the application to the container may use standard mechanisms of Stochastic Approximation theory, including the Robbins-Monro or Kiefer-Wolfowitz algorithms, to regulate the changes in payments to assure convergence. Such a design may be implemented, for example, to achieve more desirable results in connection with non-monotonic SLA parameters. For example, an embodiment using a Robbins-Monro procedure may replace steps **606-612** with the following iteration:

$$R(n+1) \leftarrow R(n) + a(n)[SLATarget - SLAParameter(R(n))]$$

where n is a counter of the iterations, $R(n)$ is a vector describing the resource bundle allocated after n iterations, $SLATarget$ is the desired value of the $SLAParameter$, and $SLAParameter(R(n))$ is the observed value of the $SLAParameter$ after n iterations. The vector $a(n)$ represents the increase/decrease of resources through the n -th step of the iteration; typically $a(n) = a/n$, where a is a fixed bundle.

Although the SLA-delivery process **600** described above uses an economic model and virtual currency units to control SLA levels, other manners of controlling SLA levels may be used according to various embodiments. For example, the allocation of resources to a container, or to an application, may be independent of any economic budget or transfer of virtual currency units, and may instead be based on other measures of an application's or container's importance.

The process **500** described above may also be modified or extended according to various other embodiments. For example, since current container systems are not readily adaptable to handling the management of storage I/O through HBA or storage systems schedulers, as an alternative to an arbitrary first-come-first-serve process, the process **500** described above may be modified or extended as shown in FIG. 7 to facilitate the handling of storage I/O.

FIG. 7 is a flow diagram of an example process **700** for economic-based I/O scheduling in a container system, which includes many of the steps of the process **500** shown in FIG. 5 and discussed above. Although not required, for the purpose of simplifying the following description, it is assumed that the consumer is a container, the provider is a server, and the resource is storage I/O. It will be understood that, according to alternative embodiments, the resource being managed may be other types of I/O, such as network I/O.

Following the deployment of the container at a server (step **508**), the server element manager monitors storage or network I/O usage by one or more containers, such as by collecting data from one or more of the container system, the HBAs (step **702**), or the NIC. According to various embodiments, the server element manager may be configured to prevent congestion along storage I/O pathways, as might occur in cases of usage levels approaching the capacity limits. For example, the server element manager may prevent congestion by using pricing functions as described below that increase prices dramatically when utilization approaches 50% of the capacity.

The server element manager optimizes or improves the resources allocated to containers, as described above (step **516**), such that containers acquire a share of the storage I/O resources that is commensurate with and optimally reflects their budget. The server element manager then periodically estimates both the average storage I/O capacity used and the average available I/O capacity, and updates the respective attributes of the storage I/O objects in the above-described supply chain model databases **246** with this usage data (step **704**). It is noted that the usage data reported to the supply

chain model databases **246** will impact price computations, with excessive utilization of storage I/O capacity resulting in respective price increases, and higher prices in turn deflecting demand by new or existing containers to servers with lower utilization (and prices) of storage I/O. For example, price competition over using storage I/O resources may result in migration of low budget containers from overloaded servers to other servers where storage I/O resources are more highly available, and are thus priced lower. Higher priority containers, on the other hand, may use their higher budgets or credit to obtain a preferential share of storage I/O resources.

The server element manager also computes the actual (versus projected) costs expended by each container, and applies these prices to handle its current commitments to containers (step **706**). For example, higher usage of storage I/O results in higher prices and immediate costs assigned to containers, such that containers of lower priority and high storage use requirements may quickly exhaust their budget or credit and be suspended or terminated, as described below. In this manner, the low priority containers relinquish storage I/O capacity to containers having a higher priority and, thus, a higher budget.

Based on the computed costs, the server element manager evaluates whether the container's budget is sufficient to pay the cost (decision block **708**). If it is, the service period of the container continues until it ends, and the server element manager notifies the container element manager of the completion of the service period (step **518**).

Otherwise, if the container's budget is not sufficient, the server element manager evaluates whether the container's credit (costs minus budget) exceeds an acceptable credit threshold (decision block **710**). According to various embodiments, high priority containers may have higher budgets and credits and can thus afford to overpay the server element manager to guarantee that they do not run out of storage I/O resources. If it is determined that the container's credit exceeds the threshold, the container element manager initiates a credit check process to decide whether the container can increase its budget or sufficiently lower its service demands, and thus continue to run (decision block **712**). If possible, the container makes any necessary adjustments (such as a budget increase in the case of high priority containers, or reduced service demands) and continues to run (step **714**), until the service period has ended and the server element manager has notified the container manager of the termination of the service period as described above. Otherwise, the server element manager suspends or terminates the container execution and notifies the container element manager, which becomes responsible for addressing the suspension or termination (step **716**).

Upon termination of the service period and notification to the container element manager, the server element manager reports usage data to the container element manager and settles any credit, overpayments or underpayments with the container element manager (step **718**). The container element manager may then proceed to shop for a new server offering lowest cost services to meet the container's needs (step **520**), as explained above.

The economic-based scheduling process **700** described above may be used effectively to de-correlate peaks of competing, bursty I/O flows. For example, consider the scenario of four containers sharing a common server and a 4 Mbps Fiber Channel HBA, where the containers generate average storage I/O flows of 250 Mbps, 250 Mbps, 200 Mbps and 300 Mbps, respectively. The aggregate demand average of 1 Gbps consumes only 25% of the HBA capacity.

A resource scheduler may limit its consideration to only the average demand which, in this case, would be manageable by the HBA and SAN. However, consider an alternate scenario where the I/O traffic streams are bursty, with a peak/average ratio of five for each container. If the four I/O streams associated with the containers are uncorrelated, their peaks will be likely dispersed and the peak of the aggregate stream will generally be less than 2 Gbps, which can be handled by the HBA and SAN with negligible or relatively few queuing delays. However, if the I/O streams are correlated, their peaks may be compounded to generate, for example, up to 5 Gbps peaks, utilizing 125% of the capacity and generating sustainable congestion, delays, and losses. The scheduling process **700** described above reduces the likelihood of compounded peaks, since they result in peak prices and a corresponding depletion of budgets and credits of low budget containers, leading to suspension, termination, or migration of such containers to servers with lower storage I/O prices until they find servers where their peaks are sufficiently de-correlated from other containers.

Thus, the allocation of containers to common servers according to the scheduling process **700** may result in substantially de-correlated peaks and substantially reduce the peak/average ratio seen by servers. For example, consider the example of four containers above. If their peaks are uncorrelated, the peaks of the aggregate stream will generally require at most 1.5 Gbps (the peak of the largest component stream), while their average traffic is 1 Gbps. The burstiness ratio (peak/average) of the aggregate stream $1.5/1=1.5$ therefore represents only 30% of the burstiness of the individual streams (1.5 divided by 5). The economic-based scheduling process **700** described above substantially reduces interference not only between traffic averages, but it also reduces the interference between correlated traffic peaks. This results in smoother, less bursty, aggregate workloads, which may permit more efficient processing.

It will be understood that, according to various embodiments, the process **700** described above to manage storage I/O flows may be applied to other forms of I/O, such as network I/O. For example, the above description should be understood to include alternative processes whereby references to "storage" are replaced by references to "network." It will similarly be understood that storage I/O flows typically utilize network-I/O flows, such as Ethernet (e.g., Fibre Channel over Ethernet (FCoE)), Transmission Control Protocol/Internet Protocol (TCP/IP) (e.g., Network File System (NFS)), and SAN (e.g., Fibre Channel (FC), Internet Small Computer System Interface (iSCSI)), to transfer information such as storage access commands. The scheduling process **700** is therefore independent of the specific underlying network, and of the specific access commands carried by the described flows. Accordingly, the process **700** may be applied to schedule network I/O flows and thereby provide similar or identical benefits to those associated with storage I/O flows, such as smoothing the peaks of bursty traffic and/or supporting priority services.

The order of steps described above with respect to scheduling process **700** is illustrative only, and can be done in different orders. Moreover, the aforementioned beneficial effects are true not only for I/O streams, but for workloads sharing other resources as well.

The contracting of services between a consumer and a provider, as described in the example processes above, may include the use of a standard request-response protocol (such as SOAP) to submit a purchase order to the provider and

transfer a respective payment. In response, the provider may deploy the service requested by the consumer and respond with a service confirmation.

FIG. 8A is an example purchase order data structure **800** issued by a consumer element manager for use in purchasing services from a provider element manager. The first two fields of the data structure **800**, source-ID field **802** and provider-ID field **804**, respectively identify the source consumer and destination provider. The third field, transaction-ID field **806**, identifies the particular purchase order. The fourth field of the data structure **800**, service field **808**, identifies the service and provides parameters to quantify the purchase. The fifth field of the data structure **800**, payment field **810**, provides payment data including payment amount and authentication data to establish the validity of the payment. Finally, the sixth field of the data structure **800**, authentication field **812**, provides data to authenticate the validity of the purchase order transaction.

FIG. 8B is an example service confirmation data structure **850** issued by the provider element manager for use in confirming or rejecting the purchase of services by the consumer element manager. The first three fields of the data structure **850**, source-ID field **852**, provider-ID field **854** and transaction-ID field **856**, correspond to the first three fields of the data structure **800** described above. The fourth field of the data structure **850**, service confirmation field **858**, includes data to confirm the service and enable the source to access it. Alternatively, assuming the provider has rejected the transaction, the service confirmation field **858** would include data with the reason or reasons for rejection, such as insufficient resources or a price change. Finally, the fifth field of the data structure **850**, authentication field **860**, provides data to authenticate the validity of the service confirmation.

As described below, various embodiments may also be used to address the problems of container sprawling and energy consumption in container systems using supply chain economics. Regarding sprawling, as explained in greater detail below, these embodiments may be used to suspend or terminate containers that are no longer needed or productive. These embodiments may also be used to terminate containers, or to disallow their re-activation if in a standby state, that are determined to be inconsistent with the current versions of their container system and applications. Regarding energy consumption, these embodiments may be used to consolidate and shift containers into fewer servers, for example, while still providing desired SLA performance, and switching other unused or non-productive servers OFF or into standby mode to reduce energy use. The supply chain software model and processes described above provide mechanisms and metrics to quantify how productive or non-productive a service element is.

The following description details an example process **900**, shown in FIG. 9, for managing the states of container system elements, which as explained further below, may be used to address sprawling and energy consumption issues. For simplicity, the following description assumes that the system element is a container, although the general principles that follow may be readily adapted for any type of system element.

A container is first initialized, for example, through the use of an initialize signal generated by a management station (step **902**) or an automated action of a container manager. Similarly, for example, an application element may interpret events generated by a launch as an initialize signal.

After being initialized, the container attempts to obtain an initial budget to acquire resources for its operations (step

904). It is next determined whether the container was successful in obtaining an initial budget (decision block **906**), in which case the container tries to acquire the resources needed to launch a respective service component (step **908**). Otherwise, it begins the termination procedure by releasing any resources allocated to it (step **910**).

If the container is successful at acquiring resources (decision block **912**), it is provisioned, deployed, and remains in an active state (step **914**) until it receives a signal to switch the service element OFF to an idle or standby state (step **916**). After the terminate signal has been received, the container begins the termination procedure by releasing resources allocated to it, as described above.

On the other hand, if the container is not successful at acquiring resources, the container will wait an amount of time for sufficient resources to become available before attempting to acquire resources again (step **918**). For example, during this waiting period, the container may use an exponential “backoff” mechanism, whereby the container repeats its attempts to acquire resources, but doubles the waiting period between repetitions with every failure. If it is determined that the container should continue to try to acquire resources (decision block **920**), it will do so as described above in step **908**. Otherwise, for example, if failures persist beyond some timeout period, the container abandons attempts to launch and begins to terminate.

Once resources have been released, it is determined whether the container should remain in a standby state (decision block **922**), in which case the execution of the container stops, but it remains in a suspended or standby state and retains sufficient state data, for example, by using storage services to retain state data in image form, and for which the container may be required to pay (step **924**). Otherwise, the container terminates execution and may be deleted (step **926**).

According to various embodiments, the applications being executed by the container are first terminated, and then the container is terminated. Such a graceful termination may be pursued through a recursive termination of the supply chain elements supported by the container. For example, a container element manager may issue a terminate signal to a corresponding operating system manager, which propagates the signal to an application manager, which in turn signals termination to its application. The application may then begin the termination steps as described above with respect to the process **900**, after which a termination complete signal to the application manager, and is forwarded to the operating system manager, which in turn sends a terminate signal and receives a termination complete signal back from the operating system. Finally, the operating system’s termination complete signal may be forwarded to the container manager, which can signal the container to terminate. It will be understood that terminating (or even suspending) a container operations may result in damages if conducted improperly or at an inappropriate time. Thus, according to various embodiments, a notification procedure may be invoked to notify administrators of pending terminations or suspensions, such that termination or suspension may only be completed once administrator permission has been received.

For a container in standby state, it is determined whether termination should follow (such as by receipt of a terminate signal) (decision block **928**), in which case the container terminates execution as described above. Otherwise, for example, if it is determined that the container should reactivate, the container seeks to obtain a budget to acquire resources for its operations as described above, for example,

35

upon receiving an initialize signal. It will be understood that the specific actions described above in connection with process 900 may be modified for non-container system elements, and that the order of steps in process 900 are also illustrative only.

According to various embodiments, a process such as process 900 described above may be used to control container sprawling by suspending or terminating non-productive system elements, such as containers. For example, consider the ROI of a container, which measures the relationship between the payments it collects from applications and the prices it pays for underlying server and I/O resources. If the container's ROI is greater than one, the container is earning more than it expends, and the container may be classified as being productive in creating applications value that exceeds the costs of the infrastructures it uses. However, if the container's ROI is less than one, this means that the container produces less value than the cost of resources it consumes, and the container may thus be classified as non-productive. In this manner, ROI is one example of a metric of productivity that may be used in determining whether a system element should be suspended or terminated, or whether it should remain active.

A process such as process 900 described above may be used to assure, for example, that applications' budgets are sufficient to keep one or more containers' ROI greater than one, and to notify applications' administrators (element managers) as needed when budgets are low. If the ROI of one or more containers remains less than one for more than a threshold period, for example, it may indicate that an application's budget is too low to sustain productive operation, and that the corresponding, non-productive container should be suspended or terminated. For example, a container may receive a terminate signal to switch it OFF to an idle or standby state (per step 916 of process 900 described above) as soon as the container's productivity level or score (for example, measured by its ROI) has been determined to be less than one for a predetermined time period. Additionally, for example, the length of time that the container's ROI has been less than one may be a factor in deciding whether the container should be terminated, or only suspended for the time being.

Similarly to dealing with the sprawling issue, the process 900 described above and similar processes may also be used for energy management. For example, such processes may be used to suspend or terminate (switch OFF) servers that are classified as being non-productive, as in the case where a server's ROI is less than one for a sufficiently long period of time. In this case, the server element manager, much like the case of the container manager described above, can monitor the ROI and detect termination or suspension conditions. The server manager may then pursue a termination process, similar to the recursive termination process described above, where all containers on the server are first terminated, or moved to another server, before the server manager suspends the server into Standby state (so as to consume less energy and cooling resources, for example) or switches the server OFF.

According to various embodiments, process 900 and similar processes may also be used to assure consistency of a suspended container with changes in applications. For example, the container manager may prevent such inconsistencies by sending a terminate signal, as described above, to all containers whenever their respective operating system or applications software has changed, thus causing the applicable containers to transition from standby to terminate state, at which point it may be deleted.

36

Although the above descriptions consider a single-domain container environment, it will be understood that the principles described herein may also be applied to multi-domain environments, e.g., a multi-cloud environment. For example, FIG. 10 is a block diagram of an example multi-domain software system environment 1000 for managing virtualized resources in "multi-cloud" systems. According to various embodiments, as shown in FIG. 10, container environment 1000 includes two example software systems 1002 and 1004, each of which is similar to the more detailed example software system 200 shown in FIG. 2, and which operate in a first and second domain, respectively.

As shown, the software system 1002 operating in the first domain includes a user interface subsystem 1006 and one or more functional managers 1008 and 1010. Together, these elements make up a functional management layer 1012 of software system 1002, and provide specific management applications as described above in connection with FIG. 2.

Software system 1002 also includes one or more element managers 1014 and 1016, which monitor and control one or more respective container stack elements 1018 and 1020. The software system 1002 also includes one or more databases 1022 (such as the supply chain databases 246 and operations databases 248 described with reference to FIG. 2), as well as a platform manager 1024. These elements are included in a platform layer 1026 of the software system 1002 to provide the infrastructures for monitoring the container stack elements 1018 and 1020, modeling these container stack elements as part of a supply chain economy, and controlling the operations of the container stack elements, as described above.

The software system 1004 operates in the second domain, includes similar elements as the software system 1002, and also includes a proxy manager 1030. According to various embodiments, the domain software system 1004 exports one or more resources or services to the domain software system 1002 by using the proxy manager 1030. The proxy manager 1030 exports instrumentation to monitor and control these provided resources to one or more of the element managers 1014 and 1016, such as container element managers, of the first domain software system 1002. The first domain software system 1002 may view the second domain software system 1004 as a service element integral with its supply chain model.

According to various embodiments, the second domain software system 1004 is in complete control of the resources (or services) and capabilities exported to the first domain software system 1002. For example, the software system 1004 may be an external cloud provider exporting raw server services to the software system 1002. In this case, the software system 1002 can access these services, using its local element managers 1014 and 1016, to allocate, for example, CPU, memory, and storage resources at the second domain software system 1004 and then monitor and control their use and operations.

Moreover, according to various embodiments, software systems 1002 and 1004 are separately owned and/or managed. For example, software system 1002 may be owned and operated by a small business that experiences steady computing needs except for two hours in each day, during which time its computing needs are consistently elevated. In this case, rather than purchasing permanent computing resources to handle the two hours of elevated needs per day, for example, software system 1002 may lease or purchase additional computing resources from software system 1004 (e.g., owned by Amazon.com, Inc.) on an as-needed basis and transfer excess workloads to software system 1004

(“bursting”). For example, computing resources from software system **1004** may be leased or purchased to facilitate the execution of a multi-tier web service by a cluster of containers (or applications). In that example, the software system **1002** may lease or sell resources from software system **1004** to execute this cluster of containers (or applications) and then migrate the container cluster (or application cluster). For example, the migration may take place from a private cloud of a small business to the public cloud of another business (e.g., of Amazon, Inc.). It is noted that, according to various embodiments, even if needed computing resources are available from within software system **1002**, such resources may be purchased from software system **1004** based on relative price offerings.

The asymmetric relationship between software systems **1002** and **1004** shown in FIG. **10** and described above may be extended to provide full symmetry. In that case, the first domain software system **1002** would incorporate its own proxy manager (not shown) to export services to the second domain software system **1004**, which would integrate it within its supply chain through one or more of its respective element managers.

The supply-chain principles discussed herein may be used to scale containers up/down, by adding or removing resources to existing container components, or to scale containers out, by adding more container components or suspending containers. The decisions to scale up/down and out are based on a supply chain **1100** outlined in FIG. **11**, and the revenues and expenses of each of the entities involved in that Figure. The supply chain **1100** may also be used to determine the sizing and placement of containers, when a selected container is deployed, and to determine future sizing and placement requirements based on anticipated changes in container load.

Turning to FIG. **11**, according to various embodiments, the supply chain **1100** may include two types of entities, namely, Service Entities (SEs), such as a Virtual Machine **1110** or a Container **1120**, and Resources, such as CPU Allocation (CPUAllocation) **1102** and Memory Allocation (MemAllocation) **1101**.

In some embodiments, the market may suggest an increase of (scaling up) the Memory Allocation **1101** of the Container **1120**, or it may suggest the creation of another instance (scaling out) of the Container **1120**. According to various embodiments, decisions to scale up/down will apply to Resources only, while decisions to scale out will apply to SEs only.

For example, in FIG. **11**, the MemAllocation **1101** of the Container **1120** may reduce as a result of congestion for resources at the Virtual Machine level. Increased utilization of MemAllocation **1101** of the Virtual Machine **1110** will lead to increased MemAllocation price. In turn, the increased MemAllocation price increases expenses of MemAllocation for the Container **1120**, leading to a decision to reduce the size of MemAllocation of the Container **1120**.

With reference now to a supply chain **1200** shown in FIG. **12**, the Container **1120** consumes directly from a Physical Machine **1210**. The MemAllocation size may also reduce as a result of congestion for resources at the Physical Machine level. Increased utilization of Physical Machine MemAllocation will lead to increased MemAllocation price, which in turn increases expenses for MemAllocation on the Container **1120**, leading to a decision to reduce the size of MemAllocation of the Container **1120**.

Container MemAllocation size may increase as a result of over provisioned resources at the Virtual Machine level. Decreased utilization of Virtual Machine CPUAllocation

due to a high capacity will lead to decreased CPUAllocation price, which in turn decreases expenses for CPUAllocation on the Container **1120**. If the Container **1120** has high revenues for CPUAllocation this would lead to a decision to increase the capacity of CPUAllocation on the Container **1120**.

Decisions for both resources and SEs are based on revenues and expenses of these resources. Similarly, expenses and revenues can be set to a predetermined value as desired. For example, the price of MemAllocation can be set to a minimum value to force higher expenses if attempting to maintain the size of the MemAllocation of the Container at or below some value. This advantageously avoids unnecessary resizing only for the purpose of having additional MemAllocation. Accordingly to other embodiments, the price of MemAllocation can be set to a maximum value.

FIG. **13** shows an example process **1300**, which illustrates how a decision is made to scale a resource allocation up or down. Turning to FIG. **13**, the process **1300** first determines if the revenue/expense of a commodity is greater than a predetermined value X (decision block **1301**). If so, then the capacity of the resource is scaled up until the revenues are equal to the expenses (step **1302**). If the revenue/expense of the resource is less than a predetermined value Y (decision block **1303**), then the resource allocation is scaled down until the revenues are equal to the expenses (step **1305**). Otherwise, if the revenues/expense of the resource is within the range defined by the values X and Y (decision blocks **1301** and **1303**), then the resource allocation is not scaled (step **1304**).

Advantageously, the values of X and Y provide a mechanism to tune the responsiveness of the system to increases or decreases in demand. The value of revenues/expenses captures the profitability of the resource allocation (or the SE). If the ratio is >1 , the resource is profitable. If it is <1 , it is losing money. In process **1300**, X is typically (but not necessarily) ≥ 1 and Y is typically (but not necessarily) < 1 . Stated in another way, an increase in capacity typically is suggested when the resource is profitable, and a decrease when it is operating at a loss.

As an additional advantage, decisions capture the entire state of the system, and can optimize the system as a whole. Increased utilization of a resource will lead to increased price for the resource, which in turn increases expenses for the resource. In some embodiments, the ideal price for scaling the resources provides 70% utilization.

In some embodiments, revenues and expenses can refer to the accumulated revenues and expenses over a period of time. Different periods of time can be used to adjust the decision-making behavior (e.g., aggressive versus conservative behavior). Short time frames lead to aggressive decisions, where the system responds very quickly to changes in the supply and demand anywhere along the supply chain. This can be used, for example, to respond quickly to congestion for resources and guarantee the quality of service offered to the entities in the system. Long time frames dampen the effects of short-term changes in the supply and demand, and reflect accurately the longer-term trends of the demand and supply.

A similar decision tree to the one shown in FIG. **13** is depicted in FIG. **14**, which illustrates an exemplary process **1400** for scaling SEs. Instead of resizing resources as shown in FIG. **13**, the process **1400** concerns creating a new instance of a SE, or suspending the operation of an existing SE, depending on the expenses and revenues of the SE. Turning to FIG. **14**, the process **1400** first determines whether the revenue/expense of a SE is greater than a

predetermined value X (decision block **1401**). If so, then a new instance of the SE is created (step **1402**). If the revenue/expense of the SE is less than a predetermined value Y (decision block **1403**), then the operation of the SE is suspended (step **1405**). Otherwise, if the revenues/expense of the SE is within the range defined by the values X and Y (decision blocks **1401** and **1403**), then the SE is unchanged (step **1404**).

As discussed above, in addition to managing container resources, the supply-chain principles discussed herein also may be used to manage application performance in other virtualization systems. For example, an application server requires a certain amount of memory and CPU resources. A database will also require a certain amount of storage. In order for the application to perform adequately, the application must be allocated a sufficient amount of resource. In order for the infrastructure to be utilized efficiently, the application should only consume what it requires at any given point in time.

Accordingly, with respect to application performance, the supply-chain principles discussed in FIGS. **13** and **14** can be used to scale up/down, by adding or removing resources allocated to the application, or to scale out, by adding more application components, or suspend application components. Some examples of application resources include, without limitation, java heap, thread pools, and connection pools in an application server or data space and log space in a relational database. These decisions are based on a supply chain **1500** outlined in FIG. **15**, and the revenues and expenses of each of the entities involved in that Figure.

Turning to FIG. **15**, the supply chain **1500** includes the two types of entities discussed with reference to FIG. **11**. Specifically, the supply chain **1500** illustrates the SEs, such as the Physical Machine **1210**, the Virtual Machine **1120**, or an Application Server **1530**, and the Resources, such as Memory (Mem) **1501**, Virtual Memory (VMem) **1502**, and Heap **1503**.

As discussed above, the resources and SEs have expenses and revenues. For example, the revenues of a virtual central processing unit (VCPU) **1504** sold by the Virtual Machine **1120** are generated from the Application Server **1530** buying this resource. Expenses of the VCPU **1504** come from paying to acquire a necessary resource, such as CPU **1505**, from the underlying Physical Machine **1210** hosting the Virtual Machine **1120**.

Similarly, a SE has revenues which can be the sum of the revenues of the resources it sells, while its expenses can be the sum of the expenses of the resources it buys. As another example, the revenues of the Virtual Machine **1120** can be the sum of the revenues of the VCPU **1504** and the VMem **1502** that it sells to the Application Server **1530** in FIG. **15**, while its expenses are the sum of the expenses to acquire the CPU **1505** and Mem **1501** from the Physical Machine **1210**.

Revenues and expenses can depend on the prices of resources, which in turn can be a function of supply, e.g., attributes of the resource such as its capacity, as well as the demand—how much of the capacity is currently utilized by resources or SEs consuming this resource. In one embodiment, price is a function of the utilization (U) of the resource, and depends on it through the formula:

$$\frac{1}{(1-U)^2}$$

For example, an application server requires java heap in order to process transactions. This java heap is allocated from the underlying virtual machine's virtual memory allocation. In the event that the demand for java heap is very high (e.g., generating revenue for the application server), and the price of virtual memory from the virtual server (e.g., determined by the combination of supply and demand) is sufficiently low, then the application server will be able to buy more virtual memory from the virtual server and allocate additional java heap. In the event that the demand for java heap is low and the price of virtual memory is high then the application server will decrease its allocation of java heap and return virtual memory to the virtual machine to be used by other applications.

In some embodiments, the buyer can be assigned a budget for purchasing the resources.

Decisions for both resources and SEs are based on the revenues and expenses of these resources. Similarly, expenses and revenues can be set to a predetermined value as desired. For example, the price of VMem can be set to a minimum value to force higher expenses if attempting to maintain the size of the Heap at or below some value. This advantageously avoids unnecessary resizing only for the purpose of having additional VMem. Accordingly to other embodiments, the price of VMem can be set to a maximum value.

In some embodiments, the market may suggest to increase (scale up) the Heap size of an Application Server, or it may suggest to create another instance (scale out) of the Application Server. These decisions can be based on the process **1300** for resizing resources and process **1400** for scaling SEs as discussed above.

As discussed above, revenues and expenses can refer to the accumulated revenues and expenses over a period of time and different periods of time can be used to adjust the decision-making behavior (e.g., aggressive versus conservative behavior). For example, longer periods of time can be used to anticipate future needs for extra application servers based on steadily increasing revenues that reflect an increase in demand. Conversely, a longer term decrease in revenues indicates that the steady state operation of a system may not require a particular SE.

The use of supply chain economic principles and other principles explained above serve several purposes and provide several potential benefits, both expressly enumerated and otherwise. For example, these principles can be used to provide a common software framework and abstractions to unify and automate the management of container systems. More specifically, they can be used to optimize or improve the allocation of IT resources (such as I/O resources or software licenses) to best process applications workloads according to their business value. The principles of supply chain economics can also be used to balance workloads to minimize disruptive operating conditions, such as I/O congestion, and to reduce resource waste by terminating or switching-off underutilized resources. These principles can also be used to empower business units to monitor and control the delivery of SLAs to their applications, as well as the ROI of individual elements and the overall container system. In addition, for example, these principles can be used to handle the management of virtual resources in a multi-cloud (or multi-domain) system.

Additionally and/or alternatively, the management of resources in container systems and conventional virtualization systems can include not only supply-chain based methods, but also access regulation to the resources. FIG. **16** illustrates an exemplary system **1600** for regulating access

of consumers 1610 (e.g., electronic applications) to resources and services (e.g., storage). In one embodiment, this regulation occurs through the use of access permits (not shown) that the consumer 1610 acquires from an intermediate entity—an Action Manager (AM) 1620—prior to accessing the resource or service. As shown in FIG. 16, the AM 1620 regulates access to a provider 1630 of the resource or service. For example, regulating access includes controlling the number of concurrent accesses, and/or the rate at which consumers 1610 access the resource, as desired.

In some embodiments, there is one type of permit per provider 1630. According to various embodiments, the AM 1620 can sell multiple types of action permits, regulating access to a number of resources. Each permit can be associated with a predetermined price. Additionally and alternatively, this price can be dynamically adjusted taking into consideration the availability of permits possessed by the AM 1620.

Permits sold by the AM 1620 can create both revenues and expenses for the AM 1620. The revenues come from the price the consumer 1610 has to pay to the AM 1620 to buy the permit. The expenses come from the price the AM 1620 has to pay to the resource provider 1630 for the right to sell these permits. For example, the AM 1620 may need to pay for Input/output Operations Per Second (IOPS) offered by a storage controller in order to allow access to the consumer 1610.

In some embodiments, the price that the AM 1620 pays for the right to sell these permits is determined by the provider 1630 based on one or more of the following parameters: the capacity and the percentage the provider 1630 wishes to make available to the consumers 1610; the current load of the provider 1630; and the rate at which the provider 1630 wishes its resources to be accessed.

The AM 1620 dynamically can adjust the number of permits it possesses at any time, depending on its revenues and its expenses. For example, if the AM 1620 is profitable (e.g., the charges based on price it is selling the permits to the consumer 1610 is higher than the charges based on price it pays to the provider 1630 for the right to sell these permits), the AM 1620 can consider increasing the number of permits it sells. Alternatively, if the AM 1620 is losing money, the AM 1620 can consider decreasing the number of permits it is selling.

Advantageously, the AM 1620 can be used to avoid I/O congestion in storage controllers when several VMs request to execute heavy-storage applications (e.g., VM Reboots, Antivirus database updates, OS Updates, and so on) at the same time. In one embodiment, the AM 1620 limits the number of concurrent consumers that can access the provider 1630. It may limit access across types of applications or within each type of application. For example, permits can be priced and provided for all anti-virus, OS updates, etc. separately, or all of them may be constrained by the same permits. In this example, the provider 1630 is the storage controller, while the consumer 1610 is the application performing the heavy-storage task. For instance, the application can be performing an anti-virus update on the virtual machine.

Turning to FIG. 17, the consumer 1610 (e.g., an application) sends the AM 1620 a request 1601 to acquire the appropriate number of permits (e.g., 5) for the provider 1630 (e.g., a storage controller) of the storage associated with the VM. It will be understood that, although reference is made to a storage controller with respect to FIG. 17, according to various embodiments, other types of providers and resources are managed using similar principles and permits. After a

request 1601 has been received, the AM 1620 subsequently determines 1602 if the request includes a sufficient budget, and if the AM 1620 has enough permits to satisfy the request 1601. If so, the AM 1620 replies to the consumer 1610 with the appropriate permits and charges. After buying the permits, the consumer 1610 accesses 1602 the storage through the provider 1630 and performs the update. After completing the update, the consumer 1610 releases 1604 the permits such that the AM 1620 can re-sell them. The AM pays 1605 the provider 1630 for the use of the permits it is selling. According to various embodiments, payment for the use of permits can occur before, after, or simultaneously with storage access.

In an alternative embodiment, the number of concurrent accesses to a resource may vary. For example, the AM 1620 adjusts the number of permits it is selling, to reflect the ability of the provider 1630 to satisfy concurrent requests by consumers 1610. For example, when the AM 1620 pays the provider 1630 for the use of the permit, the AM 1620 adjusts the number of the permits it sells based on how profitable it is. If demand for permits for a specific provider 1630 is high, the AM 1620 raises the prices for this permit, advantageously increasing revenues.

To become even more profitable, the AM 1620 can request the right to sell more permits from the provider 1630. If the provider 1630 agrees, the provider 1630 raises the price the AM 1620 has to pay for these rights. As the demand increases, the provider 1630 continues to increase the price it charges the AM 1620. At a threshold price, the AM 1620 can no longer make a profit, and the AM 1620 does not request any further increase in the number of rights it can sell. Similarly, the number of permits sold by the AM 1620 can decrease as a result of reduced demand by consumers 1610, or increased prices by the provider 1630.

In yet another embodiment, the AM 1620 controls rate of concurrent accesses to a particular resource. For example, the AM 1620 limits the rate at which the applications are accessing the storage controller to perform the heavy-storage tasks. In this case, once the application releases the permit, and until the predetermined period of time has elapsed, the AM 1620 cannot resell this permit. The storage controller can charge the AM 1620 a very small amount for the right to sell a first predetermined number of permits within a period of time, and then increase the price to infinity for permits beyond the first predetermined number in this period.

In yet another embodiment, the consumer request to access one or more permits is made directly to the resource or service provider.

In yet another embodiment, the AM 1620 controls the total number and/or the rate at which a group of consumers accesses a group of resources.

Another aspect discussed above formulates and evaluates the option to move the consumer to a new provider. “Formulating” includes the attributes taken into account when considering the option to move to the new provider. The cost of moving can be part of the comparison between two different alternatives (e.g., keeping a VM in an existing infrastructure or moving the VM to an external cloud provider). Cost can be expressed in actual currency or any unit suitable for the comparison. For example, moving time can be expressed in a real value that quantifies the cost of the VM downtime. In contrast, if there is a strict limit on acceptable downtime, the cost of moving the VM can be expressed in terms of time.

“Evaluating” includes making the decision (e.g., initiating an action based on the decision) and determining the right

time to take the action. Compared to other economics-based decision-making systems, one embodiment described herein postpones the decision for the future, advantageously waiting for a sufficient amount of time until the decision-maker is convinced that the decision is the right one.

For example, a virtualization system is considering taking an action A with the cost of taking this action represented as $C(A)$. If the action is taken, the savings over time is $S(t)$. The decision to take the action at the time t_A when the savings would have exceeded the cost of the action is represented by the following Equation:

$$S(t_A) \geq C(A)$$

In one embodiment, with reference to FIG. 18, a virtualization system 1800 controls moves of VMs 1810 between different storage (or resource) providers 1820 to avoid frequent moves of VMs 1810 between different storage providers 1820 in a datacenter (DC) 1830 or across different datacenters.

For example, the VM 1810 is evaluating a move to one or more service providers 1820, such as storage providers SP_1, SP_2, \dots, SP_N . Although storage providers 1820 are used herein as an example, it will be understood that the concepts disclosed herein can be applied to other types of service or resource providers.

In some embodiments, the cost $C(A_i)$ of moving to provider i is set to a value that is proportional to the size of the data to be moved from the current SP to SP_i , multiplied by a factor P_i that captures the ‘proximity’ of the current SP to SP_i . For example, if the current and the future SPs are in the same datacenter 1830, P_i could be set to 1, whereas if they are in different datacenters 1830, it could be set to 10, to capture that it is more expensive to move across datacenters 1830 as opposed to moving within the same datacenter 1830.

The consumer periodically checks the prices at the current and each provider i , calculates the saving for this period and adds them to the savings from the previous periods. The price of the new provider for the current period may be higher than that of the current provider, and as a result the savings for this period will be negative and will decrease the total savings from previous rounds. The moment the savings up to now exceed the cost $C(A_i)$ the VM 1810 decides to move SP_i .

In an alternative embodiment, when the consumer considers moving to a new provider, the current provider gives the consumer some credit (e.g., equal to $C(A)$) to convince the consumer to stay. The consumer accepts the credit, and periodically checks the price of the new provider. If the price is cheaper, the consumer can use this credit to subsidize any loss of not having moved there. If it is more expensive, the consumer adds her gain to the credit. If the consumer runs out of credit, then the consumer can decide to move.

Advantageously, the system accounts for the fact that a decision that looks good now may not be good in the future. For example, a consumer that buys bandwidth from a network provider may see a cheaper price offered right now by a new provider. However, the new provider may change the price an hour later, and this new price may be higher than the price of the current provider an hour later.

Additionally, the system accounts for the actual behavior of other users. Assume a VM is interested in the latency of accessing data stored on a disk, and a decision is made to move its data from the current to a new disk that currently has lower latency. For large amounts of data, the move could take hours to complete. While the move takes place, other consumers who also see a slightly reduced latency move to

the same new provider—effectively increasing the latency for everyone, and making it a bad decision.

Furthermore, the amount of time it takes to determine that the decision may be good is related to the cost of performing the action. Therefore, expensive decisions are carefully validated over longer periods than cheaper decisions, ensuring that undertaking the cost of the action will pay off in the future.

Advantageously, the systems and methods above minimize bad decisions and decisions that would frequently alternate between the current and the new provider.

Embodiments of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer program products, i.e., one or more modules of computer program instructions encoded on a tangible program carrier for execution by, or to control the operation of, data processing apparatus. The tangible program carrier can be computer readable medium, such as a machine-readable storage device, a machine-readable storage substrate, a memory device, or a combination of one or more of them.

The terms “data processing apparatus” “data processor”, or “processing device” encompasses all apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

A computer program (also known as a program, software, software application, script, or code) can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program does not necessarily correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data (e.g., one or more scripts stored in a markup language document), in a single file dedicated to the program in question, or in multiple coordinated files (e.g., files that store one or more modules, sub programs, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

The processes and logic flows described in this specification can be performed by one or more programmable processors executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit).

Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will

receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a processor for performing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, to name just a few.

Computer readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto optical disks; and CD ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input.

Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a LAN and a wide area network ("WAN"), e.g., the Internet.

While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or of what may be claimed, but rather as descriptions of features that may be specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

Particular embodiments of the subject matter described in this specification have been described. Other embodiments are within the scope of the following claims. For example, the actions recited in the claims can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

What is claimed is:

1. A computer-implemented method for managing containers in a container-based computer system, comprising:
 - initializing a container, by a container manager running on a data processor in the container system, wherein the container is a software implementation of a computer that runs on top of a shared operating system, is allocated exclusive access to compute resources using a separate operating system kernel name space, and executes one or more applications on the shared operating system;
 - determining, by the container manager, a computer resource bundle to be purchased for the container using virtual currency units;
 - identifying multiple computer servers available to offer the computer resource bundle;
 - determining a purchase price for the computer resource bundle, in virtual currency units, for each of the identified computer servers, wherein at least one determined purchase price is based at least in part on either an application-specific usage charge according to an allocated priority level for a containerized application or associated class of applications executing on the container, or a determined application-specific performance improvement that would result from allocating at least a part of the computer resource bundle to the containerized application, or both;
 - automatically selecting, by the container manager, a first one of the identified computer servers based at least in part on the purchase price for the computer resource bundle for each of the computer servers; and
 - allocating at least a part of the computer resource bundle from the first one of the identified computer servers to the containerized application executing on the container.
2. The computer-implemented method of claim 1, further comprising:
 - initiating, following an increase in the purchase price for the allocated computer resource bundle, a credit check process by the container manager, wherein the credit check process results in an increase in budget of virtual currency units available for usage by the container; and
 - determining, following the initiation of the credit check process, whether the container is to be moved from the

47

first one of the multiple computer servers to a second one of the multiple computer servers, wherein the moving determination is based at least in part on the purchase price for the computer resource bundle offered by the second one of the multiple computer servers.

3. The computer-implemented method of claim 1, further comprising:

maintaining a budget, in virtual currency units, for the container;

determining whether the budget of virtual currency units for the container is sufficient to pay the purchase price for the computer resource bundle for the selected first one of the multiple computer servers;

allocating the computer resource bundle from the first one of the multiple computer servers to the container only if the budget of the virtual machine has been determined to be sufficient to pay the purchase price for the computer resource bundle for the first one of the multiple computer servers; and

transferring virtual currency units from the budget for the container to a budget for the first one of the multiple computer servers based on the purchase price for the computer resource bundle determined for the first one of the multiple computer servers.

4. The computer-implemented method of claim 1, further comprising:

maintaining a budget, in virtual currency units, for the containerized application executing on the container;

determining a purchase price, in virtual currency units, for continued allocation of the at least a part of the computer resource bundle to the containerized application; determining whether the budget of virtual currency units for the containerized application is sufficient to pay the purchase price for the continued allocation;

continuing to allocate, to the containerized application executing on the container, the at least a part of the computer resource bundle if the budget for the containerized application is sufficient to pay the determined purchase price for continued allocation; and

deducting virtual currency units from the budget for the containerized application based on the determined purchase price for continued allocation.

5. The computer-implemented method of claim 1, wherein the computer resource bundle includes one or more computer resources, the method further comprising:

determining a performance metric level of the containerized application executing on the container; and

adjusting the allocation of at least one of the one or more computer resources to the containerized application based on a comparison of the determined performance metric level to a predetermined threshold performance metric level.

6. The computer-implemented method of claim 5, wherein the performance metric level of the application executing on the container is correlated to at least one of the transaction rate for the application, the number of transactions performed by the application, and the completion time or delay for the application.

7. The computer-implemented method of claim 1, wherein the computer resource bundle includes one or more computer resources, the method further comprising:

determining a performance metric level of the containerized application executing on the container; and

adjusting a budget of virtual currency units for the containerized application based on a comparison of the

48

determined performance metric level to a predetermined threshold performance metric level.

8. The computer-implemented method of claim 1, wherein the computer resource bundle includes at least one of an allotment of computer memory, an allotment of application execution scheduling for one or more central processing units, an allotment of storage interface bandwidth, an allotment of network interface bandwidth, and an allotment of a storage array.

9. The computer-implemented method of claim 1, wherein the allocating the computer resource bundle from the selected one of the multiple computer servers to the containerized application comprises moving the container to the selected one of the multiple computer servers from a different computer server.

10. The computer-implemented method of claim 1, wherein the computer resource bundle includes a plurality of computer resources, and wherein the automatically selecting one of the multiple computer servers includes selecting the computer server having the lowest combined purchase price for the plurality of computer resources.

11. The computer-implemented method of claim 1, wherein the allocated computer resource bundle includes at least a first computer resource and a second computer resource, the method further comprising de-allocating the first computer resource based at least in part on a comparison of a determined usage of the first computer resource to a determined usage of the second computer resource.

12. The computer-implemented method of claim 1, wherein the computer resource bundle includes multiple computer resources, and wherein the determining a computer resource bundle to be purchased is based at least in part on a computed utilization of one of the multiple computer resources.

13. The computer-implemented method of claim 12, further comprising:

increasing the amount of the computer resource bundle to be purchased if it is determined that the computed utilization of the one of the multiple computer resources is below a predetermined threshold.

14. The computer-implemented method of claim 12, further comprising:

decreasing the amount of the computer resource bundle to be purchased if it is determined that the computed utilization of the one of the multiple computer resources is above a predetermined threshold.

15. The computer-implemented method of claim 1, wherein the determining a purchase price for the computer resource bundle includes receiving one or more auction bids for purchase of the computer resource bundle.

16. The computer-implemented method of claim 1, further comprising:

maintaining a budget, in virtual currency units, for the first one of the multiple computer servers;

determining a purchase price, in virtual currency units, for a second computer resource bundle for a provider of the second computer resource bundle;

determining whether the budget of virtual currency units for the first one of the multiple computer servers is sufficient to pay the purchase price for the second computer resource bundle;

allocating, to the first one of the multiple computer servers, the second computer resource bundle if the budget for the first one of the multiple computer servers is sufficient to pay the determined price; and

49

deducting virtual currency units from the budget for the first one of the multiple computer servers based on the determined purchase price.

17. The computer-implemented method of claim 16, wherein the second computer resource bundle includes at least one of an allotment of electrical power, an allotment of air temperature cooling resources, an allotment of physical space, an allotment of operational expense, and an allotment of capital expenditure.

18. The computer-implemented method of claim 1, further comprising converting virtual currency units to government issued monetary currency based on a predetermined exchange rate.

19. A computer-implemented method for managing containers in a computer system, comprising:

determining, by a container manager running on a data processor in a first container system that includes at least one container, a computer resource bundle to be purchased for a first container in the first container system using virtual currency units, wherein the first container in the first container system is a software implementation of a computer that runs on top of a shared operating system, is allocated exclusive access to compute resources using a separate operating system kernel name space, and executes one or more applications on the shared operating system;

allocating the computer resource bundle from a first computer server in the first container system to a containerized application executing on the first container;

receiving, from a proxy manager of a second container system offering the computer resource bundle, a purchase price for the determined computer resource bundle in virtual currency units, wherein the second container system also includes at least one container that runs on top of a second shared operating system, is allocated exclusive access to compute resources using a separate operating system kernel name space, and executes one or more applications on the second shared operating system, and wherein the purchase price is based at least in part on either an application-specific usage charge according to an allocated priority level for the containerized application or associated class of applications, or a determined application-specific performance improvement that would result from allocating the computer resource bundle, or both;

determining whether the determined computer resource bundle is to be purchased from a second computer server in the second container system based at least in part on the purchase price received from the proxy manager of the second container system;

allocating the computer resource bundle from the second computer server in the second container system to the containerized application executing on the first container; and

50

dispatching the first container from the first container system to execute at the second container system.

20. The computer-implemented method of claim 19, further comprising initiating, following an increase in the purchase price for the allocated computer resource bundle, a credit check process by the container manager, wherein the credit check process results in an increase in budget of virtual currency units available for usage by the first container.

21. A computer-implemented method for managing containers in a computer system, comprising:

allocating one or more computer resources to a plurality of containers in a container system, wherein each container is a software implementation of a computer that runs on top of a shared operating system, is allocated exclusive access to computer resources using a separate operating system kernel name space, and executes one or more applications on the shared operating system, and wherein the plurality of containers operate as components of a cluster for running related applications;

determining, by a container manager running on a data processor of the container system, a target latency for the plurality of containers;

determining, by the container manager, an actual latency for the plurality of containers;

comparing the target and actual latency for the plurality of containers;

determining a purchase price for the one or more computer resources, in virtual currency units, for a second computer server, wherein the purchase price is based at least in part on either an application-specific usage charge according to an allocated priority level for a containerized application or associated class of applications, or a determined application-specific performance improvement that would result from allocating the one or more computer resources by the second computer server, or both; and

determining that the plurality of containers is to be moved from a first computer server to the second computer server based at least in part on the comparison of the target and actual latency and the determined purchase price.

22. The computer-implemented method of claim 21, further comprising initiating a migration of the plurality of containers from the first computer server to the second computer server.

23. The computer-implemented method of claim 22, wherein said initiating comprises an atomic migration of the plurality of containers.

* * * * *