



US009830082B1

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

(10) **Patent No.:** **US 9,830,082 B1**
(45) **Date of Patent:** **Nov. 28, 2017**

(54) **HYBRID HYPER-CONVERGED
INFRASTRUCTURE AND STORAGE
APPLIANCE**

(56) **References Cited**

U.S. PATENT DOCUMENTS

- (71) Applicant: **EMC IP Holding Company LLC**,
Hopkinton, MA (US)
- (72) Inventors: **Sudhir Srinivasan**, Acton, MA (US);
Devon Reed, Medway, MA (US);
Daniel Cummins, Hudson, NH (US)
- (73) Assignee: **EMC IP Holding Company LLC**,
Hopkinton, MA (US)
- (*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 0 days.

5,257,376	A	10/1993	Beck et al.
7,076,605	B1	7/2006	Son
7,076,690	B1	7/2006	Todd et al.
8,010,485	B1	8/2011	Chatterjee et al.
8,495,625	B1	7/2013	Sanders
8,930,333	B1	1/2015	Prince et al.
9,032,160	B1	5/2015	Natanzon et al.
9,063,994	B1	6/2015	Natanzon et al.
9,118,698	B1	8/2015	Radovanovic
9,158,540	B1	10/2015	Tzelnic et al.
9,286,007	B1	3/2016	Bono
9,286,261	B1	3/2016	Tzelnic et al.
2004/0078697	A1	4/2004	Duncan
2006/0179061	A1	8/2006	D'Souza et al.
2008/0005468	A1	1/2008	Faibish et al.
2009/0002380	A1	1/2009	Langyel et al.
2009/0144463	A1	6/2009	Cherian et al.
2009/0327798	A1	12/2009	D'Amato et al.
2010/0131636	A1	5/2010	Suri et al.
2011/0125951	A1	5/2011	Youngworth
2011/0225123	A1	9/2011	D'Souza et al.

(Continued)

(21) Appl. No.: **15/086,966**

(22) Filed: **Mar. 31, 2016**

Related U.S. Application Data

(63) Continuation-in-part of application No. 14/847,859,
filed on Sep. 8, 2015.

(51) **Int. Cl.**
G06F 12/12 (2016.01)
G06F 3/06 (2006.01)
G06F 9/455 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 3/0608** (2013.01); **G06F 3/065**
(2013.01); **G06F 3/0619** (2013.01); **G06F**
3/0641 (2013.01); **G06F 3/0665** (2013.01);
G06F 3/0685 (2013.01); **G06F 9/45558**
(2013.01); **G06F 2009/45579** (2013.01); **G06F**
2009/45595 (2013.01)

(58) **Field of Classification Search**
CPC G06F 1/1628; G06F 1/182; G06F 1/3234;
G06F 2009/45575; G06F 9/45558
See application file for complete search history.


Primary Examiner — Yong Choe

(74) *Attorney, Agent, or Firm* — BainwoodHuang

(57) **ABSTRACT**

A technique for operating a Hyper-Converged Infrastructure (HCI) system includes running an IO stack on a physical computing server of the HCI system. The IO stack exposes multiple protocol endpoints for providing host applications with access to data objects. Protocol endpoints are exposed both to host applications running within the HCI system and to host applications running on external computing devices, thus enabling the HCI system to double as a data storage system with respect to external computing devices.

20 Claims, 15 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

2012/0066543	A1	3/2012	Pafumi et al.
2012/0079165	A1	3/2012	Archer et al.
2012/0206463	A1	8/2012	Reid et al.
2012/0210066	A1	8/2012	Joshi et al.
2012/0262493	A1	10/2012	Tsai et al.
2013/0117286	A1	5/2013	Gallant et al.
2014/0047263	A1	2/2014	Coatney et al.
2014/0059196	A1	2/2014	Onfroy et al.
2014/0059292	A1	2/2014	Phelan et al.
2014/0192062	A1	7/2014	Kp
2014/0280489	A1	9/2014	Nandyalem et al.
2014/0333620	A1	11/2014	Park
2015/0058475	A1	2/2015	Earl et al.
2015/0058555	A1	2/2015	Karamanolis et al.
2015/0058577	A1	2/2015	Earl
2015/0186175	A1	7/2015	Van Der Walt et al.
2015/0331759	A1	11/2015	Brettell et al.
2016/0063760	A1	3/2016	Woo
2016/0212198	A1	7/2016	Krishnasamy et al.




FIG. 1
(Prior Art)





FIG. 2

**FIG. 3**





FIG. 4

**FIG. 5**




FIG. 6




FIG. 7




FIG. 8A




FIG. 8B





FIG. 9

**FIG. 10**




FIG. 11




FIG. 12




FIG. 13




FIG. 14





FIG. 15



1

HYBRID HYPER-CONVERGED INFRASTRUCTURE AND STORAGE APPLIANCE

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation-in-part of U.S. application Ser. No. 14/847,859, filed Sep. 8, 2015, the contents and teachings of which are incorporated herein by reference in their entirety.

BACKGROUND

Converged infrastructures are collections of compute, storage, network, and software resources that may be acquired, serviced, and supported as single assemblies. For example, a corporate IT (Information Technology) manager can purchase a single converged infrastructure system rather than having to purchase, assemble, and deploy separate computing servers, storage arrays, networking switches, associated software, and so on.

Hyper-converged infrastructures, or “HCI’s,” are converged infrastructures that provide additional integration at the sub-assembly level. For example, a hyper-converged system may include servers that perform multiple roles, such as any combination of compute, storage, and networking.

FIG. 1 shows an example of one prior HCI arrangement. Here, multiple client machines (e.g., user computers or devices) 110(1) to 110(N) access a system 122 over a network 120. The system 122 includes a local area network (LAN) 130 and multiple nodes. Five nodes 140(1) to 140(5) are shown. Each node is seen to include one or more processors (e.g., CPUs) and runs one or more applications, e.g., for servicing requests arriving over the network 120 from any of clients 110(1) to 110(N) and/or for performing other functions. Each of the nodes 110(1) to 110(N) includes its own local attached storage and employs a “share nothing” architecture, in which the processor and attached storage resources on each node are private to that node. To enable the nodes 110(1) to 110(N) to function together in a cluster, each node includes a file system component (FS Cmp). The file system components of the nodes together form a shared file system, with the file system component on each node communicating over the LAN 130 with the file system component on each of the other nodes to maintain a shared file system image and a single namespace across all of the nodes.

SUMMARY

The HCI arrangement of FIG. 1 involves the use of simple building-blocks and has become a popular choice among many IT specialists. Nodes can be standardized and manufactured as identical units. As resource demands increase over time, new nodes may be added.

Unfortunately, the HCI arrangement of FIG. 1 involves certain drawbacks. For example, service level agreements often require systems to maintain multiple copies of each file or other data object to prevent data loss and to promote data availability. A common reliability model requires three copies to be maintained of each data object, such that a system may suffer the simultaneous loss of two disk drives (e.g., magnetic disk drives, solid state drives, etc.) without losing data. As the disk drives of each node in the HCI arrangement of FIG. 1 are private to that node, however, making redundant copies involves copying data objects from one node to

2

another over the LAN 120. Such copies consume network bandwidth both in the nodes and in the LAN 120. They can also increase latency for other communications made over the LAN 120. Although increased network traffic can be avoided by storing multiple copies of data objects locally, i.e., within the attached storage of a single node, doing so may impair reliability, as the loss of that node can result in the loss of multiple copies of the data object. In addition, data services, such as de-duplication, run across multiple nodes and can also add considerably to network traffic.

Further, the HCI arrangement of FIG. 1 may not scale efficiently. Although new nodes may be added as demand increases, each new node must communicate with all other nodes to maintain the consistent file system image across the entire cluster. The amount of communication required to maintain the file system in a near real-time, consistent state increases dramatically with the number of nodes. Such additional communication further burdens networking resources and tends to starve out other communications over the LAN 120.

Further still, the arrangement of FIG. 1 does not normally allow the storage capacity of individual nodes to be changed as storage demands on such nodes grow. Rather, the need to maintain multiple copies of data objects across different nodes tends to force multiple nodes to require the same amount of storage. Thus, one cannot normally add storage to one node without also adding the same amount of storage to each of the other nodes. A common solution to a node running short on storage is to add a new node to the cluster and to move one or more applications from that node to the new node. This solution may lead to an over-provisioning of CPU resources, however.

Further, the HCI arrangement of FIG. 1 does not normally allow the computing power of individual nodes to be changed. If a node runs short on computing power, no other node can normally assist because the storage for the over-worked node is captive (e.g., in the share-nothing arrangement). Although data objects can be copied from one node to another over the LAN 120, a CPU operating on one node cannot directly access a data object stored on another node. Thus, as before, a common solution is to add a new node and to redistribute the workload. Adding a new node is also the common solution when the cluster as a whole runs short on computing power. But more nodes may lead to an over-provisioning of storage. Also, as previously indicated, a greater number of nodes can result in reduced performance.

Moreover, conventional protocols for maintaining a quorum of nodes in a cluster, with awareness of which nodes are working and which nodes are down, require a minimum of five nodes to support double-fault scenarios. Thus, to meet the double-fault standard, the arrangement of FIG. 1 must typically include at least five nodes. This five-node minimum requirement imposes a high entry point for the cost of hyper-converged infrastructure systems, which can make such systems out of reach for many small and medium sized businesses.

In contrast with the prior hyper-converged infrastructure arrangement, which is based on individual nodes having a share-nothing architecture, an improved technique includes providing an HCI (Hyper-Converged Infrastructure) system with one or more HCI units, each HCI unit including a pair of physical computing servers coupled to shared, nonvolatile storage. The shared, nonvolatile storage of each HCI unit is dedicated to the pair of physical computing servers and provides redundant storage of application data. Each of the pair of physical computing servers runs a set of application instances and an IO (Input/Output) stack. The application

instances receive client requests over a network and generate IO requests specifying reads and writes of application data. The IO stack processes the IO requests to effect the specified reads and writes in the shared non-volatile storage.

Example embodiments of a hyper-converged infrastructure according to the improved technique hereof overcome many deficiencies of the prior HCI arrangement. For example, any copies of data objects needed to meet service level agreements, such as requirements concerning data availability in the face of drive failures, can be maintained locally within the shared, nonvolatile storage of a single HCI unit (e.g., using RAID and/or duplication technology), thus avoiding increases in traffic over the network. Further, as increases in network traffic are avoided, the capabilities of high-speed storage media can be more fully utilized, as network latency is less likely to limit storage performance. Application efficiency is therefore improved. Unlike in the prior arrangement, providing local storage of copies of data objects comes without a loss of reliability. If one physical computing server in an HCI unit should fail, the other physical computing server within the same HCI unit can continue to access the data from the shared, nonvolatile storage. Likewise, data services, such as deduplication, can be performed entirely locally, within the shared, nonvolatile storage, thus also avoiding increases in network traffic.

Example embodiments also scale more efficiently than does the prior HCI arrangement. The improved technique imposes no particular requirements for resource balancing between physical computing servers within HCI units or even between different HCI units. As it is not required to make copies of data objects across HCI units to meet reliability requirements, disk drives can be added or removed from the shared, nonvolatile storage of any HCI unit without affecting the others. Thus, storage resources can be applied where needed, without having to balance storage among HCI units. In some examples, the IO stack of any physical computing server can expose data objects to external servers. Thus, for example, if a physical computing server runs short on computing power, an external server may be connected to that physical computing server to share the workload, with the physical computing server allowing the external server to access the same data objects that can be accessed by the overworked server.

In some examples, the shared file system of the prior arrangement is replaced with an administrative server, which maintains a system-level view of data objects across the HCI system without providing a single namespace. The administrative server, which requires relatively infrequent, low-bandwidth communications, allows the example HCI systems to grow to large numbers of HCI units without significantly increasing network traffic.

In some examples, a HCI system includes only a single HCI unit. Because HCI units include only two physical computing servers, there is no need for a quorum to maintain understanding as to whether physical computing servers are working or down. Thus, configurations can start small, e.g., with a single HCI unit, and grow over time by the addition of more HCI units, as demands increase and budgets permit.

In further respects, an improved technique for operating an HCI system includes running an IO stack in a physical computing server of the HCI system. The IO stack exposes multiple protocol endpoints for providing host applications with access to data objects. Protocol endpoints are exposed both to host applications running within the HCI system and to host applications running on external computing devices, thus enabling the HCI system to operate as a data storage system with respect to external computing devices.

Certain embodiments are directed to a method of operating a hyper-converged infrastructure (HCI) system. The method includes running an IO (Input/Output) stack on a physical computing server of the HCI system, the IO stack exposing multiple protocol endpoints for enabling host application instances to access a set of data objects for reading and writing. The method further includes processing, by the IO stack, a first set of IO requests to effect reading and writing of a first data object of the set of data objects, the first set of IO requests received from a first host application instance by a first protocol endpoint of the IO stack, the first host application instance running on the physical computing server within the HCI system. The method still further includes processing, by the IO stack, a second set of IO requests to effect reading and writing of a second data object of the set of data objects, the second set of IO requests received from a second host application instance by a second protocol endpoint of the IO stack, the second host application instance running on a computing device external to the HCI system, the IO stack thereby serving both the first application instance running within the HCI system and the second application instance running external to the HCI system.

Other embodiments are directed to a hyper-converged infrastructure (HCI) system including multiple HCI units interconnected via a network, with at least one of the HCI units including a physical computing server constructed and arranged to perform a method of operating and HCI system, such as the method described above.

Still other embodiments are directed to a computer program product. The computer program product includes a set of non-transitory, computer-readable media storing instructions which, when executed on a physical computing server of a Hyper-Converged Infrastructure (HCI) system, cause the HCI system to perform a method, such as the method described above.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

The foregoing and other features and advantages will be apparent from the following description of particular embodiments of the invention, as illustrated in the accompanying drawings, in which like reference characters refer to the same or similar parts throughout the different views. In the accompanying drawings,

FIG. 1 is a block diagram of an example HCI (hyper-converged infrastructure) arrangement according to the prior art;

FIG. 2 is a block diagram of an example environment and HCI system in which embodiments of the improved technique hereof can be practiced;

FIG. 3 is a block diagram of an example IO (Input/Output) stack, which may be operated within each physical computing server of the HCI system of FIG. 2;

FIG. 4 is a block diagram showing an additional server used to augment the processing power of a physical computing server of the HCI system of FIG. 2;

FIG. 5 is a block diagram of an example physical computing server of the HCI system of FIG. 2 in additional detail;

FIG. 6 is a block diagram of an example set of processors in the example physical computing server of FIG. 5 in additional detail;

FIG. 7 is a block diagram of an example administrative client that may run on each physical computing server in the HCI system of FIG. 2;

FIGS. 8A and 8B are block diagrams that show an example movement of an application instance from one physical computing server in one HCI unit to another physical computing server in another PCI unit, with consequent redirection of IO requests;

FIG. 9 is an example screenshot produced by an administrative server operating in the HCI system of FIG. 2;

FIG. 10 is a flowchart showing an example process of managing host application instances in an HCI system, such as the one shown in FIG. 2;

FIG. 11 is a block diagram showing an external computing device coupled to a physical computing server of the HCI system of FIG. 2 to access one or more data objects in the HCI system;

FIG. 12 is a block diagram that shows multiple protocol endpoints and associated host-accessible data objects served by the IO stack of the physical computing server of FIG. 12;

FIG. 13 is a block diagram showing an example arrangement for running the IO stack and various host applications in respective containers;

FIG. 14 is a block diagram showing a variant of FIG. 13 in which the containers in which the host applications run are themselves run within virtual machines;

FIG. 15 is a block diagram showing another variant of FIG. 13 in which the containers in which the IO stack and the host applications run are themselves run within virtual machines; and

FIG. 16 is a flowchart showing an example method for operating an HCI system.

DETAILED DESCRIPTION OF THE INVENTION

Embodiments of the invention will now be described. It is understood that such embodiments are provided by way of example to illustrate various features and principles of the invention, and that the invention hereof is broader than the specific example embodiments disclosed.

Section I: Example HCI Arrangement:

This section presents an improved technique for managing applications, which includes providing an HCI (Hyper-Converged Infrastructure) system with one or more HCI units, each HCI unit including a pair of physical computing servers coupled to shared, nonvolatile storage.

FIG. 2 shows an example environment 200 in which embodiments of the improved technique hereof can be practiced. Here, multiple client computing devices (“clients”) 210(1) through 210(N) connect to an HCI system 222 over a network 220. The HCI system 222 includes HCI units 240. The HCI units 240 are connected to the network 220. It should be understood that the HCI system 222 may include any number of HCI units 240, including a single HCI unit.

The HCI system 222 is further seen to include an administrative server 242. The administrative server 242 may run as a software construct on any of the physical computing servers 250, on its own computer, or elsewhere. In some examples, the administrative server 242 is implemented in a shared arrangement across multiple physical computing servers 250 or other computers. The administrative server 242 may run in a virtual machine, in a container, or in some other fashion. Various arrangements are contemplated.

Each HCI unit 240 includes shared storage 270 and a pair of physical computing servers 250, i.e., a first physical computing server 250a and a second physical computing server 250b. In some examples, an HCI unit 240 may include greater than two physical computing servers 250,

each coupled to respective shared storage 270. Thus, at least two physical computing servers 250 are provided per HCI unit 240, and different HCI units 240 may include different numbers of physical computing servers 250. In an example, each of the physical computing servers 250 is a physical hardware server, such as a server blade, which plugs into a chassis that encloses and cools the physical computing server. It should be understood, however, that the physical computing servers 250 require no particular physical form. The physical computing servers 250 in each HCI unit 240 are connected together via interconnect 258, which in an example is a high-speed interconnect, such as PCI express, which forms a private network. Although FIG. 2 shows different HCI units 240 in respective blocks, it should be understood that HCI units 240 need not have separate enclosures. For example, the physical computing servers 250 of multiple HCI units 240 may be located together within the same chassis. Further, each HCI unit 240 may connect to the network 220 by way of individual network adapters in their respective physical computing servers 250.

The shared storage 270 includes a set of nonvolatile storage devices, such as magnetic disk drives 272, electronic flash drives 274, and/or other types of nonvolatile storage devices, in any suitable arrangement and number. The example shown is merely illustrative. In an example, the shared storage 270 within each HCI unit 240 is configured such that each physical computing server 250 within that HCI unit is capable of reading and writing each nonvolatile storage device contained within the shared storage 270. In an example, connections 268a and 268b between physical computing servers 250 and storage devices in the shared storage 270 may be achieved using SAS (Serial-Attached-SCSI), e.g., with each physical computing server 250 attached to each storage device in the shared storage 270; however, this is not required. Although the shared storage 270 is shown as being enclosed within an HCI unit 240, it should be understood that no particular physical arrangement is required. For example, shared storage 270 of multiple HCI units 240 may be co-located in the same physical chassis or enclosure, or in multiple enclosures or storage shelves, with the relationships between physical computing servers 250 and respective shared storage 270 of each HCI unit 240 established via physical connections (e.g., 268a and 268b) and configuration settings.

Each physical computing server 250 “includes” (i.e., realizes by operation of executable instructions) a respective set of host application instances 252 (i.e., one or more instances of one or more application programs), an IO (Input/Output) stack 260, and an administrative client 262. The host application instances 252 are instances of one or more software programs. In an example, these programs are server-side components of transactional applications, such as backend database programs, web applications, server applications, and so forth. This is not required, however, as the programs may be of any type. Multiple instances may be run of each program, and multiple programs may be run.

The IO stack 260 on each physical computing server 250 provides an execution path for IO requests 254 generated by the application instances 252 running on that physical computing server 250. As will be described, the IO stack 260 exposes host-application-accessible data objects built from storage elements (e.g., slices) within the shared storage 270. Host application instances 252 can access these host-application-accessible data objects for reading and/or writing via the IO stack 260. In some examples, as will also be described, host applications 252 running on other servers in

the HCI system 222 may also access the host-application-accessible data objects exposed by the IO stack 260.

The administrative client 262 running on each physical computing server 250 collects information about the host-application-accessible data objects exposed by the IO stack 260 on that physical computing server 250, such as their names, size, etc., and provides the information to the administrative server 242, which maintains a system-level view of all host-application-accessible data objects served across the HCI system 222. In some examples, the administrative client 262 on each physical computing server 250 also receives instructions from the administrative server 242 to perform administrative operations on host-application-accessible data objects. In some examples, the administrative server 242 is realized by one or more administrative clients 262, which may be elected as leader by other administrative clients 262.

It should be understood that the client computing devices 210(1) through 210(N) may be any type of network-connectable computing devices, such as desktop computers, laptop computers, smart phones, tablets, PDAs (personal data assistants), game consoles, set-top boxes, and so forth, for example. The environment 1200 may include any number of clients, including a single client. The network 220 may be any type of network or combination of networks, such as a storage area network (SAN), a local area network (LAN), a wide area network (WAN), the Internet, and/or some other type of network or combination of networks, for example.

In example operation, clients 210(1) through 210(N) issue client requests to the HCI system 222 to perform transactions via host application instances running on the HCI system 222. For example, client 210(1) runs a client program (e.g., a browser or application frontend) that directs a client request 212 to a host application instance 252 running on the physical computing server 250a in HCI unit 240(1). The physical computing server 250a receives the client request 212, and the host application instance processes the client request 212. Such processing may include generating IO requests 254 specifying reads and/or writes to one or more host-application-accessible data objects built from storage elements within the shared storage 270 of HCI unit 240(1) (and/or, in some cases, to one or more host-application-accessible data objects built from storage elements within the shared storage 270 of other HCI units 240). The host-application-accessible data objects may be LUNs (Logical Units), file systems, VVOLs (virtual volumes, available from VMware, Inc. of Palo Alto, Calif.), or other types of data objects, for example. The IO stack 260 receives the IO requests 254 and processes them to effect the specified reads and/or writes on the data objects served from the shared storage 270. When each IO request 254 is completed, the IO stack 260 returns an acknowledgement 256 to the host application instance 252. In some examples, the host application instance 252 then returns an acknowledgement 214 to the client 210(1). Operation can proceed in this fashion for any number of client requests 212, directed to any number of host application instances 252 on any number of physical computing servers 250 in the HCI system 222.

It should be understood that there is no necessity for there to be a one-to-one relationship between client requests 212 and IO requests 254. Rather, it may often be the case that a single client request 212 gives rise, via operation of an application instance 252, to multiple IO requests 254, or that multiple client requests 212 give rise to a single IO request 254. Also, in some examples, some application instances 252 running within the HCI system 222 may act as clients

or peers of other application instances 252 running within the HCI system 222, which act as servers or peers. Thus, client requests 212 need not arise externally to the HCI system 222, but may arise internally. In some examples, a single client request 212 received by one physical computing server 250 may result in IO requests 254 being issued by multiple physical computing servers 250. Further, it should be understood that some application instances 152 may generate IO requests 254 in the absence of any client requests 212. For example, an application instance 252 may generate IO requests 254 at its own initiative, e.g., in response to the occurrence of predetermined events, on some schedule, and/or for any reason.

It can be seen from the foregoing that the HCI system 222 provides any desired level of data protection without inducing excessive network traffic. The shared storage 270 in any HCI unit 240 may provide redundant storage at any desired level. For instance, the nonvolatile storage devices within the shared storage 270 may be arranged in RAID (Redundant Array of Independent Disks) groups, which prevent loss of data or data availability even in the face of multiple drive failures. For example, RAID 6 provides double parity and may be used to support two simultaneous drive failures without data loss. Other arrangements may provide triple parity to support three simultaneous drive failures. Alternatively, or in addition, the shared storage 270 may support mirrored copies (RAID 1) of data objects, with as many redundant copies made as needed to support service level agreements. As the shared storage 270 maintains redundant copies locally, within a single HCI unit 240, little to no network traffic is needed to support such copies. Likewise, data services, such as deduplication and compression, may be performed locally to each HCI unit 240, again without entailing any increase in network traffic.

Such savings in network traffic comes at virtually no cost in terms of reliability. If a first physical computing server within an HCI unit 240 should fail, the second physical computing server in the same HCI unit 240 can continue to access the data of the first physical computing server from the shared storage 270. Any data transfer required to perform failover may be conducted over the high-speed interconnect 258.

FIG. 3 shows an example IO stack 260 of the physical computing servers 250 in greater detail. Beginning at the bottom of FIG. 3 and working up, HDD/EFD support 334 includes drivers that perform actual reading from and writing of magnetic disk drives, electronic flash drives, etc., in the shared storage 270. RAID manager 332 arranges the storage devices into RAID groups and provides access to the RAID groups using RAID protocols.

Storage pool 330 expresses RAID groups in the form of internal LUNs, which the storage pool 330 carves into slices. A "slice" is an increment of storage space, such as 256 MB or 1 GB in size, for example, which is composed from a portion of an internal LUN. In an example, the storage pool 330 provisions slices to internal file systems to support the storage of host-application-accessible data objects. The storage pool 330 may also de-provision slices from internal file systems if their storage is no longer needed. In some example, each host-application-accessible data object has an owner, which is designated to be a particular physical computing server 250. In addition, each slice may have an owner, also designated to be a particular physical computing server 250. In an example, only slices owned by a physical computing server may be provisioned to host-application-accessible data objects also owned by that physical computing server.

Internal file system manager **328** builds and manages internal file systems upon slices served by the storage pool **330**. Internal file systems can realize a variety of host-application-accessible data objects, such as LUNs, host file systems, and VVOLs in the form of files. VVOLs are similar to VMDKs (Virtual Machine Disks) and provide storage realizations of virtual machine disks. The shared storage **270** may support any number of internal file systems, and each internal file system may include any number of files.

Object mapping **326** expresses files stored in internal file systems as respective, host-application-accessible data objects **324**. For example, the object mapping **326** maps each file representing a host-application-accessible data object **324** to a respective internal volume. Such mapping may express the logical address range of the file as a physical address range of the internal volume. Higher levels of the IO stack **260** can then access the internal volume using block-based semantics. For hosting LUNs and VVOLs, the object mapping **326** may simply provide the internal volumes directly. For hosting file systems, the object mapping **326** may build a host file system on top of an internal volume and express the resulting file system.

Cache manager **322** manages a memory cache on the respective physical computing server **250**. As will be described, cache management may include storing data, specified in IO requests **254**, to a local memory cache, mirroring the data to a memory cache on the other physical computing server **250** in the pair (i.e., in the same HCI unit **240** over the interconnect **258**), and providing an acknowledgement **256** to the IO request **254** only when the data are stored in both memory caches. Typically, such acknowledgements **256** may be provided long before the data are fixed in the nonvolatile devices of the shared storage **270**. Thus, the cache manager **322** enables the IO stack **260** to respond to IO requests **254** with very low latency.

Redirector **320** selectively redirects IO requests **254** to another physical computing server **250**, which may be in the same HCI unit **240** or in a different HCI unit **240**. For example, the redirector **320** checks the identity of the host-application-accessible data object **324** specified in each IO request **254**. If the identified host-application-accessible data object **324** is served from the same physical computing server **250** on which that IO stack **260** is running, the redirector **320** allows the IO request **254** to pass through to lower levels of the IO stack **260**. However, if the identified host-application-accessible data object **324** is served from another physical computing server **250**, the redirector **320** forwards the IO request **254** to the other physical computing server **250**.

The protocol end points **318** expose the host-application-accessible data objects **324** to host application instances **252** in accordance with respective protocols for accessing those data objects. Thus, the protocol end points **318** may expose LUNs and block-based VVols using a block-based protocol (e.g., Fiber Channel or iSCSI, for example) and may expose host file systems and file-based VVols using a file-based protocol (e.g., NFS or CIFS, for example).

It should be understood that the particular order of layers in the IO stack **260** may be varied. For example, the cache manager **322** may, in some variants, be disposed directly above the storage pool **330** or even below the storage pool **330**. Similarly, the redirector **320** may be disposed below the host-application-accessible data objects **324** or at any other suitable location. Further, the cache manager **322** and/or the redirector **320** may be provided in multiple portions, which

operate at multiple layers within the IO stack **260**. The particular arrangement shown is intended to be merely illustrative.

FIG. 4 shows an example arrangement for adding computing power to the HCI system **222**. Here, an additional server **450**, which may run one or more host application instances **252** and an administrative client **262**, is connected to physical computing server **250a** of HCI unit **240**. Additional server **450** may be another physical computing server **250**, e.g., of the same type and in the same chassis as those already described, or it may be an external server located elsewhere. Connection between the additional server **450** and the physical computing server **250a** may be established via network **220** or in any other suitable fashion (e.g., using iSCSI, Fibre Channel, etc.). The connection allows host application instances **252** running on the additional server **450** to send IO requests **254** to the IO stack **260** of physical computing server **250a**, and to receive responses **256** therefrom. Thus, host application instances **252** running on additional server **450** can access the IO stack **260** of physical computing server **250a** in parallel with the host application instances **252** running locally on physical computing server **250a**, to access the same host-application-accessible data objects **324** that are accessible to physical computing server **250a**. Any number of additional servers, like the server **450**, may be added in this fashion, with such additional servers connected to any physical computing server **250** in the HCI system **220**. It should be noted that the above arrangement allows computing power to be added without having to add corresponding amounts of storage. Thus, the above arrangement avoids over-provisioning of storage when adding computing power.

Although FIG. 4 shows the additional server **450** connecting to a single physical computing server **250a**, it should be understood that the additional server **450** may connect, over the network **220**, to any physical computing server **250** within the HCI system **222**. For example, the additional server **450** may issue IO requests to one physical computing server **250** at one time and to another physical computing server **250** at another time. In this manner, the additional server **450** may be regarded as supporting the entire HCI system **222**, including any of its physical computing servers **250**.

FIG. 4 further shows additional storage devices **474** added to shared storage **270** (e.g., four new EFDs, three new magnetic disk drives, and a cloud drive). These additional storage devices **474** may be of any type. They may be added to available slots within an existing shelf or chassis, or in one or more expansion shelves, for example. In some examples, additional storage may be provided in the form of cloud-based storage, e.g., on an external Ethernet drive and/or on a drive or array available over a WAN or over the Internet. It should be noted that storage devices **474** can be added without having to add additional computing power. Although shown in the same figure, the additions of computing power and storage devices are independent. Thus, the above arrangement allows storage to be added to an HCI unit **240** without having to add corresponding amounts of computing power and thereby avoids over-provisioning of computing power when adding storage. Further, adding storage devices **474** to the shared storage in one HCI unit **240** does not impose any requirement that storage also be added to shared storage **270** in any other HCI unit **240**.

FIG. 5 shows an example physical computing server **250** in additional detail. Here, the physical computing server **250** is seen to include one or more communication interfaces **510**, a set of processors **512**, and memory **520**. The com-

11

munication interfaces **510** include, for example, SCSI target adapters and TCP/IP network interface adapters for converting electronic and/or optical signals received over the network **220** to electronic form for use by the physical computing server **250**. The set of processors **512** includes one or more processing chips and/or assemblies. In a particular example, the set of processors **512** includes multiple CPU (Central Processing Unit) cores, as well as associated processing chip sets (e.g., coprocessors and the like). The memory **520** includes both volatile memory (e.g., RAM), and non-volatile memory, such as one or more ROMs, disk drives, solid state drives, and the like. The set of processors **512** and the memory **520** together form control circuitry, which is constructed and arranged to carry out various methods and functions as described herein. Also, the memory **520** includes a variety of software constructs realized in the form of executable instructions. When the executable instructions are run by the set of processors **512**, the set of processors **512** are caused to carry out the operations of the software constructs. Although certain software constructs are specifically shown and described, it is understood that the memory **520** typically includes many other software constructs, which are not shown, such as various applications, processes, and daemons.

As further shown in FIG. 5, the memory **520** includes a cache **522**. In an example, the cache **522** is composed of a set of DRAM (Dynamic Random Access Memory) devices, which are backed up by battery **524**. In the event of the power loss, power from battery **524** continues to energize the DRAM devices and allows them to retain their contents. Also, as described in connection with FIG. 3 above, cache manager **322** acts to mirror the contents of cache **522** to a similar cache (which is also battery-backed) on another physical computing server **250** (i.e., on the other physical computing server in the same HCI unit). Mirroring of cached data may be conducted at high speed over the interconnect **258** and helps to provide low-latency responses **256** to IO requests **254** (FIG. 2). It is observed that data stored in battery-backed DRAM on two different servers is, for all relevant intents and purposes, securely persisted.

In some examples, the physical computing server **250** runs an operating system (OS) virtualization framework **530** for managing virtual operating environments (VOEs). For example, OS virtualization framework **530** may manage one VOE (e.g., VOE **540(1)**) for the IO stack **260**, other data services **542**, a virtualization client **544**, and the administrative client **262**. The OS virtualization framework **530** may also manage a respective VOE (e.g., VOE **540(2)** and VOE **540(3)**), for each host application instance, e.g., **550(1)** and **550(2)**. The data services **542** may provide a wide range of functions, such as deduplication, compression, replication, backup, and so forth. Alternatively, the data services **542** (or any subset of them), the virtualization client **544**, and/or the administrative client **262** may be operated in another, or in separate, VOEs.

In some examples, the OS virtualization framework **530** includes a hypervisor, and the VOEs are respective virtual machines managed by the hypervisor. According to this example, the IO stack **260**, data services **542**, virtualization client **544**, and administrative client **262** may run together in one virtual machine, and each host application instance may run in its own, separate virtual machine. The physical computing server **250** may employ any suitable virtual machine technology, such as vSphere (available from VMware, Inc. of Palo Alto, Calif.), KVM (Kernel-based Virtual Machine), Hyper-V (available from Microsoft Corporation of Redmond, Wash.), and so forth. If the physical

12

computing servers **250** in the HCI system **222** operate virtual machines (VMs), then, in some examples, a separate VM manager (e.g., vCenter for VMware deployments) may be provided, e.g., on any of the physical computing servers **250** or elsewhere, and each virtualization client **544** may include a client component for communicating with the VM manager (e.g., a VASA provider). In some examples, the VM manager is collocated with the administrative server **242**, although this is by no means required.

In other examples, the OS virtualization framework **530** includes a container-hosting component, and the VOEs are respective containers managed by the container-hosting component. Unlike virtual machines, which virtualize entire operating systems, containers virtualize only the userspace component of an operating system, but not the kernel. Thus, where the OS virtualization framework **530** includes a container-hosting component, the physical computing server **250** operates an OS kernel **532**. Each container is a software process that provides an isolated userspace execution context built upon the kernel **532**. Examples of container-hosting components include Docker, available from Docker, Inc. of San Francisco, Calif., and LXC (Linux Containers).

FIG. 6 shows aspects of the set of processors **512** of FIG. 5 in additional detail. Here, it is seen that the set of processors **512** in physical computing server **250** includes numerous CPU (central processing unit) cores **610**. Although 24 cores **610** are shown, it is understood that any number of cores **610** may be provided. In the example shown, the CPU cores **610** are divided into two groups: a first group **620** and a second group **630**. In an example, the first group **620** of CPU cores **610** is dedicated to data storage and overhead operations, such as running the IO stack **260**, data services **542**, virtualization client **544**, and administrative client **262**. In an example, no host application instances **252** run on the first group of cores **620**. The second group **630** of CPU cores **610** is dedicated to host application instances **252**; however this second group **630** does not operate the IO stack **260** or participate in performing data services **542** or operating the administrative client **262**.

This segregation of cores into different groups for performing different functions may be fixed or it may be adjustable. In one example, a fixed set of cores may be dedicated to a controller virtual machine (e.g., VOE **540(1)** of FIG. 5), with other cores made available to applications. In another example, there is no fixed assignment of cores. A system administrator may establish a setting for assigning cores **610** to one role or another, and the system administrator may change the setting over time. In some examples, the physical computing server **250** monitors its own CPU utilization and dynamically assigns cores **610** to group **620** or to group **630** based on monitored performance. Thus, for example, if some number of cores in group **620** is idle or underutilized while the cores in group **630** are operating near capacity, the physical computing server **250** may dynamically reallocate cores to move some cores **610** from group **620** to group **630**.

FIG. 7 shows an example administrative client **262** in further detail. Here, it is seen that the administrative client **262** includes a client adapter **710** and a system object map **720**. In an example, the client adapter **710** gathers information about host-application-accessible data objects **324** served from the respective physical computing server **250** and reports the gathered information to the administrative server **242**. The client adapter **720** may also receive instructions from the administrative server **242** and implement them locally.

13

The system object map 720 provides a map of host-application-accessible data objects 324 across the entire HCI system 222. For example, the system object map 720 includes entries that associate identifiers of host-application-accessible data objects 324 with identifiers of the respective physical computing servers 250 that serve those objects. In an example, the administrative client 262 on each physical computing server 250 maintains its own copy of the system object map 720, to enable fast access to the system object map 720 at runtime. This is not strictly required, however, as multiple physical computing servers 250 may share one or more copies of the system object map 720.

FIGS. 8A and 8B show an example scenario in which a system object map 720 may be used. As shown in FIG. 8A, a host application instance 252a runs on physical computing server 250b of HCI unit 240(1). During operation, host application instance 252a issues IO requests 254 to the respective IO stack 260a to effect reads and/or writes of host-application-accessible data objects 324 (e.g., LUNs, VVOLS, file systems, etc.) realized in the shared storage 270 of HCI unit 240(1) and accessible via the IO stack 260a.

At some point, e.g., when physical computing server 250b in HCI unit 240(1) becomes very busy, and administrator may wish to move host application instance 252a to a less busy physical computing server, such as to physical computing server 250a in HCI unit 240(2). If application instance 252a runs within a virtual machine, then moving the application instance to a different server may be effected by executing simple commands (e.g., using vMotion for VMware implementations).

FIG. 8B shows the state of the system after the host application instance 252a has been moved to server 250a in HCI unit 240(2). Here, although the application instance 252a has been successfully moved, the host-application-accessible data objects 324 accessed by instance 252a have not. Therefore, when application instance 252a issues IO requests 254 to IO stack 260b, on physical computing server 250a of HCI unit 240(2), the IO stack 260b will not be able to service the IO requests locally, as the addressed data objects do not reside locally. In this circumstance, the IO stack 260b accesses the system object map 720 to identify the physical computing server that serves the data objects specified in the IO requests 254. The IO stack 260 then forwards the IO requests (e.g., via the redirector 320; see FIG. 3) to the identified server. In this example, the system object map 720 identifies physical computing server 250b on HCI unit 240(1) as the host of the addressed object, and the redirector 320 forwards the IO requests 254 there.

At some point, it may be desired to move the host-application-accessible data objects 324 accessed by application instance 252a from the shared storage 270 in HCI unit 240(1) to the shared storage 270 in HCI unit 240(2). However, the decision to move the data is left to the system administrator and/or may be policy-driven. For example, movement of data may not be desirable if movement of the application 252a is intended to be temporary.

FIG. 9 shows an example system-level view 900 of host-application-accessible data objects 324 across the HCI system 222, which may be provided, for example, by the administrative server 242, e.g., via a GUI (graphical user interface) application. Each HCI unit may be shown (e.g., HCI Unit 1, HCI Unit 2, etc.). A user can click the text for each HCI unit to reveal details, including names and properties of constituent physical computing servers (e.g., PCS 1, PCS 2). Clicking the text for each PCS reveals further information, such as names and properties of virtual machines running on each PCS. Drilling down further

14

reveals names 910 of individual host-application-accessible data objects 324 (VVOL 22, LUN 30, etc.). The system-level view 900 thus provides information of all host-application-accessible data objects 324 across the entire HCI system 222, as well as additional information. It should be understood that the particular format of the system level view 900 is intended to be merely illustrative, as presentation details and modes of operation may vary. It should be understood that the information presented in system-level view 900 has relatively low information content and tends to change slowly. Thus, communications required to maintain the system-level view 900 typically place only small demands on the network 220.

FIG. 10 shows an example method 1000 for managing host application instances in a hyper-converged infrastructure system. The method 1000 may be carried out, for example, by the software constructs, described in connection with FIGS. 2-5 and 7, which reside, at least in part, in the memories 520 of the physical computing servers 250 and are run respective sets of processors 512. The various acts of the method 1000 may be ordered in any suitable way. Accordingly, embodiments may be constructed in which acts are performed in orders different from that illustrated, which may include performing some acts simultaneously, even though the acts are shown as sequential in the illustrated embodiments.

At 1010, a pair of physical computing servers coupled to a network are provided. Each of the pair of physical computing servers runs a respective set of host application instances and runs an IO (Input/Output) stack. For example, as shown in FIG. 2, a pair of physical computing servers 250a and 250b are provided within HCI unit 240(1). Each of the physical computing servers 250a and 250b runs a respective set of host application instances 252 and runs a respective IO stack 260.

At 1012, a set of nonvolatile storage devices are shared between the pair of physical computing servers. For example, as further shown in FIG. 2, each of the physical computing servers 250a and 250b in HCI unit 240(1) is coupled to shared, nonvolatile storage 270, with storage devices 272 and 274 being accessible in parallel by both physical computing servers 250a and 250b. The HCI system 222 may include any number of HCI units 240, with each HCI unit 240 including its own pair of physical computing servers 250a and 250b and its own shared nonvolatile storage 270.

At 1014, the set of nonvolatile storage devices redundantly store data of the host applications instances running on the pair of physical computing servers. For example, the shared, nonvolatile storage 270 of HCI unit 240(1) stores redundant copies, e.g., in fault-tolerant RAID groups and/or as mirrored copies, of host-application-accessible data objects 324, which may be built from storage elements, such as slices, which are derived from the shared, nonvolatile storage 270, and which are served to host application instances 252 via the IO stack(s) 260 of one or both physical computing servers 250a and 250b.

At 1016, each of the pair of physical computing servers (i) receives client requests over the network, the client requests directed to one of the set of host application instances running on that physical computing server, (ii) generates, via operation of that host application instance, IO requests specifying reads and writes of application data, and (ii) processes the IO requests, via the respective IO stack, to effect the specified reads and writes of application data in the set of nonvolatile storage devices. For example, a host application instance 252 running on physical computing

server **250a** receives a client request **212** over the network **220**. The host application instance **252** processes the client request **212** and generates, in response to the client request **212** or on its own initiative, IO requests **254** specifying reads and writes of a host-application-accessible data object **324** served from the shared, nonvolatile storage **270**. The IO stack **260** running on physical computing server **250(a)** processes the IO requests **254** to effect the reads and writes specified by the IO requests **254** on the host-application-accessible data objects **324** built from storage elements, such as slices, in the shared, non-volatile storage **270**.

This section has described an improved technique for providing an HCI (Hyper-Converged Infrastructure) system **222**. The system includes one or more HCI units **240**, each HCI unit **240** including a pair of physical computing servers **250a** and **250b** coupled to shared, nonvolatile storage **270**. The shared, nonvolatile storage **270** of each HCI unit **240** is dedicated to the pair of physical computing servers **250a** and **250b** and provides redundant storage of application data. As compared with prior HCI implementations, embodiments of the invention hereof provide reduced network congestion, improved resource allocation, and better scaling.

Section II: Example Hybrid HCI/Data Storage System Arrangement:

An improved technique for operating a Hyper-Converged Infrastructure (HCI) system includes running an IO stack on a physical computing server of the HCI system. The IO stack exposes multiple protocol endpoints for providing host applications with access to data objects. Protocol endpoints are exposed both to host applications running within the HCI system and to host applications running on external computing devices, thus enabling the HCI system to double as a data storage system with respect to external computing devices. In a non-limiting example, the HCI system is constructed in accordance with the embodiments described in Section I; however, this is not required.

FIG. **11** shows an example arrangement in which the HCI system **222** (FIG. **2**) acts as a data storage system with respect to external computing devices. As is known, a data storage system is an arrangement of hardware and software that includes one or more storage processors coupled to arrays of non-volatile storage devices, such as magnetic disk drives, electronic flash drives, and/or optical drives, for example. The storage processors service storage requests, arriving from host machines ("hosts"), which specify files or other data elements to be written, read, created, or deleted, for example. Software running on the storage processors manages incoming storage requests and performs various data processing tasks to organize and secure the data elements stored on the non-volatile storage devices.

In the example shown, the HCI system **222** performs a dual role, both as an HCI with respect to locally-running applications and as a data storage system with respect to applications running on external hosts. In its data storage system role, the physical computing servers **250** act as storage processors, i.e., by receiving IO requests arriving from external hosts and processing the IO requests to effect read and write operations on data objects stored in the HCI system **222**. The data objects accessed by external hosts may include block-based objects, such as LUNs and block-based VVols (virtual machine disks). They may also include file-based objects, such as file systems and file-based VVols. Host applications may access block-based objects using iSCSI (Internet Small Computer Systems Interface), Fibre Channel, and/or some other block-based protocol, for example. Also, host applications may access file-based objects using NFS (Network File System), CIFS (Common

Internet File System), and/or some other file-based protocol, for example. In some examples, the data objects accessed by external hosts are distinct from the ones accessed by local applications running within the HCI system **222**. However, this is not required, as the same data objects may be accessed by both local applications and remote hosts.

As shown in the example of FIG. **11**, an external computing device **1150** (a host) connects to the HCI system **222** over the network **220** (FIG. **2**), e.g., using iSCSI, Fibre Channel, NFS, and/or CIFS. Although only one external computing device **1150** is shown, any number of external computing devices may connect to the HCI system **222**. The external computing device **1150** (or each such device) may be any type of computing device, such as a desktop computer, a server, a virtual machine, a workstation, a laptop computer, a smartphone, a tablet computer, a personal reader, a set top box, a video streaming device, and so on, i.e., any device capable of running a host application **1152** and issuing IO requests to the HCI system **222**. In an example, the external computing device **1150** is not managed by the HCI system **222**, e.g., it does not include an administrative client **262**.

In example operation, the physical computing server **250a** runs a set of local application instances **252**, which generate IO requests **254a**. The IO requests **254a** specify reads and writes of data objects (e.g., LUNs, file systems, VVols, etc.), which are composed from storage elements in the shared storage **270**. The IO stack **260** running within the physical computing server **250a** receives the IO requests **254a** and processes them to effect the requested reads and writes. For example, and referring briefly back to FIG. **3**, the IO stack **260** exposes protocol endpoints **318**, i.e., one protocol endpoint **318** for each data object to be served. The IO stack **260** receives the IO requests **254a** at the protocol endpoints **318** and processes the IO requests using the layers described in connection with FIG. **3**. For illustrative purposes, it is assumed that IO requests **254a** are all directed to a first data object, and that the IO stack **260** in the physical computing server **250a** manages access to the first data object via a first protocol endpoint **318**. It should be understood, however, that the IO requests **254a** may be directed to any number of data objects, access to which is managed by any number of corresponding protocol endpoints. In response to each IO requests **254a**, the IO stack **260** provides a response **256a**, e.g., which acknowledges completion of the respective IO request **254a**.

In parallel with the set of application instances **252** on physical computing server **250a** issuing IO requests **254a**, the host application instance **1152** on external computing device **1150** issues IO requests **254b**. The IO stack **260** on the physical computing server **250a** receives the IO requests **254b**, which may arrive at any time relative to the IO requests **254a**. The IO stack **260** processes the IO requests **254b** as it does for the IO requests **254a**, e.g., using the processing layers described in connection with FIG. **3**. For illustrative purposes, it is assumed that IO requests **254b** are all directed to a second data object, and that the IO stack **260** in the physical computing server **250a** manages access to the second data object via a second protocol endpoint **318**. This is merely an example, however, as IO requests **254b** may be directed to any number of data objects via respective protocol endpoints. In response to each IO request **254b**, the IO stack **260** issues a response **256b**, e.g., which acknowledges completion of the respective IO request **254b**.

As FIG. **11** further shows, the IO stack **260** running on physical computing server **250a** may further receive IO requests **254c**. For example, the IO stack running in physical

computing server **250b** may redirect IO requests **254c** to physical computing server **254a**. The IO stack **260** on physical computing server **250a** receives the IO requests **254c**, which may arrive at any time relative to the IO requests **254a** and **254b**. The IO stack **260** processes the IO requests **254c** as it does for the IO requests **254a** and **254b**, e.g., using the processing layers described in connection with FIG. 3. For illustrative purposes, it is assumed that IO requests **254c** are all directed to a third data object, and that the IO stack **260** in physical computing server **250a** manages access to the third data object via a third protocol endpoint **318**. This is merely an example, however, as IO requests **254c** may be directed to any number of data objects via respective protocol endpoints. In response to each IO request **254c**, the IO stack **260** provides a response **256c**, e.g., which acknowledges completion of the respective IO request **254c**.

In an example, the HCI system **222** includes a storage administrative program (not shown) to administer host-application-accessible data objects, such as to manage their creation, destruction, replication, migration, storage tiering, and so forth. In some examples, the storage administrative program is integrated with the administrative server **242** (FIG. 2), such that the administrative server **242** manages both HCI-related settings and data-storage-related settings.

With the arrangement described, the IO stack **260** operates in a similar fashion regardless of whether the IO requests originate from a local application instance (e.g., App(s) **252**) or a remote application instance (e.g., App(s) **1152**). Thus, the physical computing server **250a** enables the HCI system **222** to perform a dual role as both an HCI system and a data storage system. The HCI system **222** may perform both roles simultaneously. Although data storage operations are described herein with reference to physical computing server **250a**, it should be appreciated that physical computing server **250b** may also perform these same functions, as may any other physical computing server **250** in the HCI system **222**.

Enabling the HCI system **222** to operate as a data storage system promotes customer convenience, cost savings, energy savings, and ease of administration. Rather than having to purchase, power, administer, cool, and maintain two different systems, e.g., an HCI system and a separate data storage system, customers may instead be able to meet their needs with a single HCI system **222**.

FIG. 12 shows examples of protocol endpoints **318** and associated host-application-accessible data objects **324** in additional detail. Here, each host-application-accessible data object **324** is associated with a respective protocol endpoint **318**. For example,

- LUN **324a** is associated with iSCSI protocol endpoint **318a**, which may include, for example, a SCSI target;
- Block-based VVol **324b** is associated with iSCSI protocol endpoint **318b**, which may include, for example, a SCSI target. Note that block-based VVols may also be associated with Fibre Channel protocol endpoints;
- LUN **324c** is associated with Fibre Channel protocol endpoint **318c**, which may include, for example, a SCSI target;
- File-based VVol **324d** is associated with NFS protocol endpoint **318d**, which may include an NFS mount point;
- Host file system **324e** is associated with NFS protocol endpoint **318e**, which may include an NFS mount point; and
- Host file system **324f** is associated with CIFS protocol endpoint **318f**, which may include a CIFS share.

For each host-application-accessible data object **324**, the respective protocol endpoint **318** provides an access point to which host applications can connect for reading, writing, and/or performing other activities on that data object. The first, second, and third data objects, described above, may be any of the data objects **324**, and the first, second, and third protocol endpoints may be those provided for accessing the first, second, and third data objects, respectively. It should be appreciated that the particular protocol endpoints, data objects, and combinations of protocol endpoints with data objects as shown are merely examples, and what is depicted is not intended to represent all possibilities.

In an example, host application instances bind to protocol endpoints **318** to enable them to access respective host-application-accessible data objects **324**. Any number of host application instances may bind to each protocol endpoint. However, application instances that are not bound to a protocol endpoint cannot access the data object associated with that protocol endpoint. Protocol endpoints **318** are available for binding to both local application instances **252** and external application instances **1152**.

FIG. 13 shows example software constructs within the memory **252** of physical computing server **250**. In an example, the arrangement of FIG. 13 is a variant of that shown in FIG. 5 and is common to all physical computing servers **250** in the HCI system **222**.

In the example shown, the IO stack **260** and each of application instances **1320(1)**, **1320(2)**, etc., run within respective containers, i.e., **1310(1)** for the IO stack and **1310(2)**, **1310(3)**, etc. for the application instances (there may be any number of application instances). A container-hosting component **1350** provides a set of tools for creating, managing, and communicating between and among containers **1310**. Each “container” provides an isolated userspace execution environment. As compared with a virtual machine, which virtualizes an entire operating system as well as hardware, a container virtualizes only the userspace portion of an operating system. A container does not virtualize the operating system kernel or the hardware. Containers are thus similar to virtual machines but are more lightweight. In an example, the container-hosting component **1350** is commercially available from Docker, Inc., of San Francisco, Calif. Also, in an example, the physical computing server **250** runs Linux as its host operation system.

FIG. 14 shows a variant of the arrangement of FIG. 13. Here, the IO stack **260** runs in a container **1410(1)**, which is managed by container-hosting component **1350** (e.g., Docker), as in FIG. 13. Here, however, application instances **1420(1)**, **1420(2)**, etc. are run within respective virtual machines **1412(1)**, **1412(2)**, etc. A hypervisor **1450** manages the virtual machines **1412(1)**, **1412(2)**, etc. In an example, the virtual machines **1412(1)**, **1412(2)**, etc. are kernel-based virtual machines (KVM) and the hypervisor **1450** is a KVM hypervisor. In some variants, the KVM hypervisor is merged with the Linux kernel, such that both are provided together as a single component.

According to some variants, the application instances **1420(1)**, **1420(2)**, etc., also run within respective containers **1410(2)**, **1410(3)**, etc., which may be Docker containers, for example. Here, each of the virtual machines **1412(1)**, **1412(2)**, etc., virtualizes an entire guest operating system, including its kernel. For example, virtual machine **1412(1)** virtualizes kernel **1430(1)** and virtual machine **1412(2)** virtualizes kernel **1430(2)**. Although not shown, each of the virtual machines **1412(1)**, **1412(2)**, etc., also runs its own container-hosting component **1350** (e.g., Docker), for managing the container within the respective virtual machine.

FIG. 15 shows another variant of the arrangement of FIG. 13. Here, both the IO stack 260 and application instances 1520(1), 1520(2), etc., run within respective virtual machines, 1512(1), 1512(2), 1512(3), etc., which are managed by a hypervisor 1550. Each of the virtual machines 1512(1), 1512(2), 1512(3), etc., runs a guest operating system, which includes a kernel (e.g., see kernels 1530(1), 1530(2), 1530(3), etc.). Although not shown, the memory 520 may also run a host operating system, such as Linux or Windows, for example. The arrangement of FIG. 15 may be provided, for example, in VMWare implementations.

In some examples, the IO stack 260 and the application instances 1520(1), 1520(2), etc., each run within respective containers 1510(1), 1510(2), 1510(3), etc. For example, each of the virtual machines 1512(1), 1512(2), 1512(3), etc., runs its own container-hosting component 1350 (e.g., Docker, not shown), for managing the container within the respective virtual machine.

FIG. 16 shows an example method 1600 for operating an HCI system. The method 1600 may be carried out, for example, by the software constructs, described in connection with FIGS. 2-5, 7, and 11-15, which reside in the memory 520 of one or more physical computing servers 250 and are run by respective sets of processors 512. The various acts of the method 1600 may be ordered in any suitable way. Accordingly, embodiments may be constructed in which acts are performed in orders different from that illustrated, which may include performing some acts simultaneously, even though the acts are shown as sequential in the illustrated embodiments.

At 1610, an IO stack 260 is run on a physical computing server 250 of the HCI system 222. The IO stack 260 exposes multiple protocol endpoints 318 for enabling host application instances 252, 1152 to access a set of data objects 324 for reading and writing.

At 1620, a first set of IO requests 254a are processed by the IO stack 260 to effect reading and writing of a first data object (e.g., any of 324a to 324f) of the set of data objects. The first set of IO requests 254a are received from a first host application instance 252 by a first protocol endpoint (e.g., any of 318a to 3180 of the IO stack 260. The first host application instance 252 runs on the physical computing server 250 within the HCI system 222.

At 1630, a second set of IO requests 254b are processed by the IO stack 260 to effect reading and writing of a second data object (e.g., any other one of 324a to 3240 of the set of data objects. The second set of IO requests 254b are received from a second host application instance 1152 by a second protocol endpoint (e.g., any other one of 318a to 3180 of the IO stack 260. The second host application instance 1152 runs on a computing device 1150 external to the HCI system 222. The IO stack 260 thereby serves both the first application instance 252 running within the HCI system 222 and the second application instance 1152 running external to the HCI system 222.

An improved technique has been described for operating a Hyper-Converged Infrastructure (HCI) system 222. The technique includes running an IO stack 260 on a physical computing server 250 of the HCI system 222. The IO stack 260 exposes multiple protocol endpoints 318 for providing host applications 252, 1152 with access to data objects 324. Protocol endpoints 318 are exposed both to host applications 252 running within the HCI system 222 and to host applications 1152 running on external computing devices 1150, thus enabling the HCI system 222 to double as a data storage system with respect to external computing devices 1150.

Having described certain embodiments, numerous alternative embodiments or variations can be made. For example, although embodiments have been described in connection with the HCI system 222 disclosed in Section I, this is merely one example. Alternatively, embodiments may be constructed using other HCI system designs, including designs that do not employ HCI units 240 each including pairs of physical computing servers 250 coupled to shared storage 270.

Further, although various container and virtual machine configurations have been specifically shown and described, these are merely examples, as other embodiments may be constructed using other container and/or virtual machine technologies and configurations different from those shown.

Further, although features are shown and described with reference to particular embodiments hereof, such features may be included and hereby are included in any of the disclosed embodiments and their variants. Thus, it is understood that features disclosed in connection with any embodiment are included as variants of any other embodiment.

Further still, the improvement or portions thereof may be embodied as a computer program product including one or more non-transient, computer-readable storage media, such as a magnetic disk, magnetic tape, compact disk, DVD, optical disk, flash drive, SD (Secure Digital) chip or device, Application Specific Integrated Circuit (ASIC), Field Programmable Gate Array (FPGA), and/or the like (shown by way of example as medium 1050 in FIG. 10 and as medium 1650 of FIG. 16). Any number of computer-readable media may be used. The media may be encoded with instructions which, when executed on one or more computers or other processors, perform the process or processes described herein. Such media may be considered articles of manufacture or machines, and may be transportable from one machine to another.

As used throughout this document, the words “comprising,” “including,” “containing,” and “having” are intended to set forth certain items, steps, elements, or aspects of something in an open-ended fashion. Also, as used herein and unless a specific statement is made to the contrary, the word “set” means one or more of something. This is the case regardless of whether the phrase “set of” is followed by a singular or plural object and regardless of whether it is conjugated with a singular or plural verb. Further, although ordinal expressions, such as “first,” “second,” “third,” and so on, may be used as adjectives herein, such ordinal expressions are used for identification purposes and, unless specifically indicated, are not intended to imply any ordering or sequence. Thus, for example, a second event may take place before or after a first event, or even if no first event ever occurs. In addition, an identification herein of a particular element, feature, or act as being a “first” such element, feature, or act should not be construed as requiring that there must also be a “second” or other such element, feature or act. Rather, the “first” item may be the only one. Although certain embodiments are disclosed herein, it is understood that these are provided by way of example only and that the invention is not limited to these particular embodiments.

Those skilled in the art will therefore understand that various changes in form and detail may be made to the embodiments disclosed herein without departing from the scope of the invention.

What is claimed is:

1. A method of operating a hyper-converged infrastructure (HCI) system, the method comprising:
running an I/O (Input/Output) stack on a physical computing server of the HCI system, the I/O stack exposing

21

multiple protocol endpoints for enabling host application instances to access a set of data objects for reading and writing;

processing, by the I/O stack, a first set of I/O requests to effect reading and writing of a first data object of the set of data objects, the first set of I/O requests received from a first host application instance by a first protocol endpoint of the I/O stack, the first host application instance running on the physical computing server within the HCI system; and

processing, by the I/O stack, a second set of I/O requests to effect reading and writing of a second data object of the set of data objects, the second set of I/O requests received from a second host application instance by a second protocol endpoint of the I/O stack, the second host application instance running on a computing device external to the HCI system, the I/O stack thereby serving both the first application instance running within the HCI system and the second application instance running external to the HCI system.

2. The method of claim 1, further comprising:

processing, by the I/O stack, a third set of I/O requests to effect reading and writing of a third data object of the set of data objects, the third set of I/O requests received from a third host application instance by a third protocol endpoint of the I/O stack, the third host application instance running on a second physical computing server within the HCI system.

3. The method of claim 2,

wherein at least one of the first protocol endpoint, the second protocol endpoint, and the third protocol endpoint supports a block-based protocol.

4. The method of claim 3, wherein at least one other of the first protocol endpoint, the second protocol endpoint, and the third protocol endpoint supports a file-based protocol.

5. The method of claim 4, wherein the block-based protocol is one of iSCSI (Internet Small Computer Systems Interface) and Fibre Channel, and wherein the file-based protocol is one of NFS (Network File System) and CIFS (Common Internet File System).

6. The method of claim 3,

wherein the first application instance runs within a first container on the physical computing server, the first container providing a first isolated userspace execution environment on the physical computing server,

wherein the I/O stack runs within a second container on the physical computing server, the second container providing a second isolated userspace execution environment on the physical computing server, and

wherein the first set of I/O requests originate from the first application instance running within the first container and propagate within the physical computing server to the I/O stack running within the second container.

7. The method of claim 6,

wherein the physical computing server runs a hypervisor for managing virtual machines;

wherein the first container runs within a virtual machine, and

wherein the first set of I/O requests originate from the first application instance running in the first container within the virtual machine and propagate within the physical computing server to the I/O stack running within the second container.

8. The method of claim 7,

wherein the second container runs within a second virtual machine, and

22

wherein the first set of I/O requests originate from the first application instance running in the first container within the virtual machine and propagate, via the hypervisor, to the I/O stack running in the second container within the second virtual machine.

9. The method of claim 7,

wherein the first data object is a first VVol (Virtual Volume) that provides a virtual machine disk for supporting the first application instance, and

wherein the second data object is one of a LUN (Logical UNit), a file system, and a second VVol that provides a virtual machine disk for supporting the second application instance.

10. The method of claim 7,

wherein the first data object is a first VVol (Virtual Volume) that provides a virtual machine disk for supporting the first application instance, and

wherein the second data object is a second VVol that provides a virtual machine disk for supporting the second application instance within a virtual machine running on the external computing device.

11. A hyper-converged infrastructure (HCI) system, the HCI system comprising:

multiple HCI units; and

a network interconnecting the HCI units,

wherein at least one of the HCI units includes a physical computing server constructed and arranged to:

run an I/O (Input/Output) stack on a physical computing server of the HCI system, the I/O stack exposing multiple protocol endpoints for enabling host application instances to access a set of data objects for reading and writing;

process, by the I/O stack, a first set of I/O requests to effect reading and writing of a first data object of the set of data objects, the first set of I/O requests received from a first host application instance by a first protocol endpoint of the I/O stack, the first host application instance running on the physical computing server within the HCI system; and

process, by the I/O stack, a second set of I/O requests to effect reading and writing of a second data object of the set of data objects, the second set of I/O requests received from a second host application instance by a second protocol endpoint of the I/O stack, the second host application instance running on a computing device external to the HCI system, the I/O stack thereby serving both the first application instance running within the HCI system and the second application instance running external to the HCI system.

12. The HCI system of claim 11, wherein each HCI unit includes a pair of physical computing servers coupled to the network and a set of shared, nonvolatile storage devices coupled to each of the pair of physical computing servers and shared between the pair of physical computing servers, the set of shared, nonvolatile storage devices providing redundant storage of data of host application instances running on the respective pair of physical computing servers.

13. A computer program product including a set of non-transitory, computer-readable media storing instructions which, when executed on a physical computing server of a Hyper-Converged Infrastructure (HCI) system, cause the HCI system to perform a method, the method comprising:

running an I/O (Input/Output) stack on a physical computing server of the HCI system, the I/O stack exposing

23

multiple protocol endpoints for enabling host application instances to access a set of data objects for reading and writing;

processing, by the I/O stack, a first set of I/O requests to effect reading and writing of a first data object of the set of data objects, the first set of I/O requests received from a first host application instance by a first protocol endpoint of the I/O stack, the first host application instance running on the physical computing server within the HCI system; and

processing, by the I/O stack, a second set of I/O requests to effect reading and writing of a second data object of the set of data objects, the second set of I/O requests received from a second host application instance by a second protocol endpoint of the I/O stack, the second host application instance running on a computing device external to the HCI system, the I/O stack thereby serving both the first application instance running within the HCI system and the second application instance running external to the HCI system.

14. The computer program product of claim 13, wherein the method further comprises:

processing, by the I/O stack, a third set of I/O requests to effect reading and writing of a third data object of the set of data objects, the third set of I/O requests received from a third host application instance by a third protocol endpoint of the I/O stack, the third host application instance running on a second physical computing server within the HCI system.

15. The computer program product of claim 14, wherein at least one of the first protocol endpoint, the second protocol endpoint, and the third protocol endpoint supports a block-based protocol,

wherein at least one other of the first protocol endpoint, the second protocol endpoint, and the third protocol endpoint supports a file-based protocol,

wherein the block-based protocol is one of iSCSI (Internet Small Computer Systems Interface) and Fibre Channel, and

wherein the file-based protocol is one of NFS (Network File System) and CIFS (Common Internet File System).

16. The computer program product of claim 15, wherein the first application instance runs within a first container on the physical computing server, the first

24

container providing a first isolated userspace execution environment on the physical computing server, wherein the I/O stack runs within a second container on the physical computing server, the second container providing a second isolated userspace execution environment on the physical computing server, and

wherein the first set of I/O requests originate from the first application instance running within the first container and propagate within the physical computing server to the I/O stack running within the second container.

17. The computer program product of claim 16, wherein the physical computing server runs a hypervisor for managing virtual machines;

wherein the first container runs within a virtual machine, and

wherein the first set of I/O requests originate from the first application instance running in the first container within the virtual machine and propagate within the physical computing server to the I/O stack running within the second container.

18. The computer program product of claim 17, wherein the second container runs within a second virtual machine, and

wherein the first set of I/O requests originate from the first application instance running in the first container within the virtual machine and propagate, via the hypervisor, to the I/O stack running in the second container within the second virtual machine.

19. The computer program product of claim 17, wherein the first data object is a first VVol (Virtual Volume) that provides a virtual machine disk for supporting the first application instance, and

wherein the second data object is one of a LUN (Logical UNit), a file system, and a second VVol that provides a virtual machine disk for supporting the second application instance.

20. The computer program product of claim 17, wherein the first data object is a first VVol (Virtual Volume) that provides a virtual machine disk for supporting the first application instance, and

wherein the second data object is a second VVol that provides a virtual machine disk for supporting the second application instance within a virtual machine running on the external computing device.

* * * * *