



(51) International Patent Classification:

G06F 9/455 (2018.01) G06F 9/44 (2018.01)
G06F 9/50 (2006.01)

(21) International Application Number:

PCT/EP2017/058445

(22) International Filing Date:

07 April 2017 (07.04.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(71) Applicant: NEC LABORATORIES EUROPE GMBH
[DE/DE]; Kurfürsten-Anlage 36, 69115 Heidelberg (DE).

(72) Inventors: MANCO, Filipe; Karlsruher Str. 122, 69126
Heidelberg (DE). KUENZER, Simon; Prinz-Eugen-Str. 5,
68735 Waghäusel (DE). SCHMIDT, Florian; Hauptstr. 2,

69221 Dossenheim (DE). HUICI, Felipe; Silbershohl 3 A,
69221 Dossenheim (DE).

(74) Agent: ULLRICH & NAUMANN; Schneidmühlstraße
21, 69115 Heidelberg (DE).

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR,
KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG,
MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,
PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC,
SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR,
TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,

(54) Title: METHOD FOR OPERATING VIRTUAL MACHINES ON A VIRTUALIZATION PLATFORM AND COR-
RESPONDING VIRTUALIZATION PLATFORM




Fig. 6b

(57) Abstract: A method for operating virtual machines on a virtualization platform, the method comprising: embedding, preferably by a toolstack of the virtualization platform, control information in a predetermined memory area of a front-end virtual machine where at least one virtual device is to be initialized, wherein said control information is required for initiating a communication with a back-end virtual machine where at least one back-end driver runs; retrieving, by the front-end virtual machine, said control information from said predetermined memory area of the front-end virtual machine; and performing the communication between the front-end virtual machine and the back-end virtual machine via a direct communication channel in order to exchange information for initializing said at least one virtual device of the front-end virtual machine, in particular by communicating with said at least one back-end driver via said direct communication channel, wherein said direct communication channel is established based on said control information embedded in said predetermined memory area of the front-end virtual machine. Furthermore, a corresponding virtualization platform is disclosed.

GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

METHOD FOR OPERATING VIRTUAL MACHINES ON A VIRTUALIZATION PLATFORM AND CORRESPONDING VIRTUALIZATION PLATFORM

5 The present invention relates to a method for operating virtual machines on a virtualization platform.

Furthermore, the present invention relates to a virtualization platform for operating virtual machines.

10

In recent years lightweight virtualization technologies such as Docker retrievable at “<https://www.docker.com/>” and LXC retrievable at “[https:// linuxcontainers.org](https://linuxcontainers.org)” are gaining enormous traction not only in the research field but also in terms of real-world deployment. Google, for instance, is reported to run all of its services in containers, retrievable at “http://www.theregister.co.uk/2014/05/23/google_containerization_two_billion/”, and Container as a Service (CaaS) products are available from a number of major players including Azure’s Container Service (retrievable at “<https://azure.microsoft.com/en-us/services/container-service/>”), Amazon’s EC2 Container Service and Lambda offerings (retrievable at “<https://aws.amazon.com/lambda>”), and Google’s Container Engine service (retrievable at “<https://cloud.google.com/container-engine>”).

20

Beyond these services, lightweight virtualization is crucial to a wide range of use cases, including just-in-time instantiation of services (e.g. described in the non-patent literature of MADHAVAPEDDY, A., LEONARD, T., SKJEGSTAD, M., GAZA-GNAIRE, T., SHEETS, D., SCOTT, D., MORTIER, R., CHAUDHRY, A., SINGH, B., LUDLAM, J., CROWCROFT, J., AND LESLIE, I. Jitsu: Just-in-time summoning of unikernels. In 12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15) (Oakland, CA, 2015), USENIX Association, pp. 559–573) – e.g., filters against Distributed Denial of Service (DDoS) attacks, TCP acceleration proxies, content caches, etc. – and networks functions virtualization (NFV), all the while providing significant cost reduction through consolidation and power minimization.

25

30

5 The reasons for containers to have taken the virtualization market by storm are clear. In contrast to heavyweight, hypervisor-based technologies – which represent virtualization platforms – such as VMWare, KVM or Xen, containers provide extremely fast instantiation times, small per-instance memory footprints, and high density on a single host, among other features.

10 However, no technology is perfect, and containers are no exception. Security, for one, has been and continues to be a thorn on their side. First, their large trusted computing base (TCB), at least compared to type-1 hypervisors, has resulted in a large number exploits. Second, a container that causes a kernel panic will bring down the entire host. Further, any container that can monopolize or exhaust system resources (e.g., memory, file descriptors, user IDs, forkbombs, etc.) will cause a Denial-of-Service (DOS) attack on all other containers on that host. Over the years, a significant amount of effort has resulted in the introduction of mechanisms such as user namespaces and Seccomp that harden or eliminate a large number of these attack vectors. However, a simple misconfiguration can still lead to an insecure system.

20 Beyond security, another downside of containers is that their sharing of the same kernel rules out the possibility to specialize the kernel and its network stack to provide better functionality and performance to specific applications. Finally, containers do not currently support live migration, although support for it is under development.

25 At least for multitenant deployments, this leaves with a difficult choice between

- (1) containers and the security issues surrounding them, and
- (2) the burden coming from heavyweight, VM-based platforms.

30 Clearly, it cannot easily, overnight, be fixed all of the security issues related to containers, nor prevent new ones from arising.

Thus, the ability to quickly boot virtual machines (VMs), destroy them, migrate them, and concurrently run as many of them on a single server is important to a vast number of applications in the field of Network Function Virtualization (NFV). For example, to run as many vCPEs (virtual customer premises equipments) as possible on a single server, to instantiate firewalls on a per-connection basis, to dynamically create filters to deal with Denial-of-Service attacks, to be able to quickly and dynamically boot monitoring services to oversee financial transactions, and to host services whose key performance indicators (KPIs) depend on boot times such as block chain and function-based services such as Amazon's Lambda, among many others.

A significant part of the overhead when booting a virtual machine or migrating it comes from the scalability of the back-end information store, for example the XenStore in the Xen hypervisor, which is used to keep control information about the instances currently running in the system.

Hence, known virtualization platforms use a back-end information store to keep track of control information about the virtual machines currently running on the system such as a unique machine identifier, a name, the amount of memory allocated, etc., along with information about the virtual devices they are using, for example network device addresses and device capabilities. While certainly useful, the back-end information store is often a source of bottlenecks that only get worse as the number of virtual machines increases. The reason for this is that an operation like virtual machine creation requires multiple interactions with such a back-end information store.

The back-end information store is crucial to the way a virtualization platform such as Xen functions, with many xl commands making heavy use of it. As way of illustration Fig. 5a shows the process when creating a virtual machine and its (virtual) network device. First, the toolstack writes an entry to the network back-end's directory, essentially announcing the existence of a new virtual machine in need of a network device. Previous to that, the back-end placed a watch on that directory; the toolstack writing to this directory triggers the back-end to assign an event channel and other information (e.g., grant references, a mechanism for

sharing memory between virtual machines and to write it back to the back-end information store such as the XenStore of Xen. Finally, when the virtual machine boots up, it contacts the back-end information store to retrieve the information previously written by the network back-end. As such, the above is a simplification:
5 in actuality, the virtual machine creation process alone can require interaction with over 30 back-end information store entries, a problem that is exacerbated with increasing number of virtual machines and devices. Worse, the back-end information represents a single point of failure.

10 As previously mentioned, the importance of small creation and boot times is at least partly demonstrated by the rise of containers and their typically faster-than-VMs boot times, although containers trade-off this performance against isolation, which is basic to a number of the application scenarios mentioned above. Known virtualization platforms, and the virtual machines that run on top of them, appears
15 to be inherently and fundamentally heavyweight and slow to boot.

In view of the above, it is therefore an object of the present invention to improve and further develop a method of the initially described type for operating virtual machines on a virtualization platform in such a way that virtual machine creation,
20 migration and/or destruction is improved, in particular with regard to virtual machine creation, migration and/or destruction times, preferably while retaining the isolation and security afforded by the underlying virtualization platform. Furthermore, a corresponding virtualization platform is to be developed.

25 In accordance with the invention, the aforementioned object is accomplished by a method for operating virtual machines on a virtualization platform, the method comprising:

embedding, preferably by a toolstack of the virtualization platform, control information in a predetermined memory area of a front-end virtual machine where
30 at least one virtual device is to be initialized, wherein said control information is required for initiating a communication with a back-end virtual machine where at least one back-end driver runs;

retrieving, by the front-end virtual machine, said control information from said predetermined memory area of the front-end virtual machine; and

performing the communication between the front-end virtual machine and the back-end virtual machine via a direct communication channel in order to exchange information for initializing said at least one virtual device of the front-end virtual machine, in particular by communicating with said at least one back-end driver via said direct communication channel, wherein said direct communication channel is established based on said control information embedded in said predetermined memory area of the front-end virtual machine.

Furthermore, the aforementioned object is accomplished by a virtualization platform for operating virtual machines, comprising a hypervisor that is configured to run one or more back-end virtual machines and one or more front-end virtual machines,

wherein an entity, in particular a toolstack, of said virtualization platform is configured to embed control information in a predetermined memory area of a front-end virtual machine where at least one virtual device is to be initialized, wherein said control information is required for initiating a communication with a back-end virtual machine where at least one back-end driver runs,

wherein the front-end virtual machine is configured to retrieve said control information from said predetermined memory area of the front-end virtual machine, and

wherein the front-end virtual machine and the back-end virtual machine are configured to perform the communication between the front-end virtual machine and the back-end virtual machine via a direct communication channel in order to exchange information for initializing said at least one virtual device of the front-end virtual machine, in particular by communicating with said at least one back-end driver via said direct communication channel, wherein said direct communication channel is established based on said control information embedded in said predetermined memory area of the front-end virtual machine.

According to the invention it has first been recognized that an enormous improvement with regard to virtual machine creation, migration and/or destruction times can be achieved by providing a virtualization platform that is able to function without a centralized back-end information store. Specifically, the single point of failure that a virtualization platform's back-end information store may represent is

removed by embedding, preferably by a toolstack of the virtualization platform, control information in a predetermined memory area of a front-end virtual machine where at least one virtual device is to be initialized. The control information is required for initiating a communication with a back-end virtual machine where at least one back-end driver runs. According to the invention, the front-end virtual machine retrieves the control information from the predetermined memory area of the front-end virtual machine, preferably when the front-end virtual machine boots up. Thus, a method or virtualization platform according to the invention may achieve its improvement by having an entity, in particular a modified toolstack, that embeds control information directly inside of a front-end virtual machine's memory and by having the front-end virtual machine retrieve this control information from the predetermined memory area, e.g. from its local memory when the front-end virtual machine boots up. Furthermore, according to the invention the communication between the front-end virtual machine and the back-end virtual machine is performed via a direct communication channel in order to exchange information – in particular in a decentralized way – for initializing the at least one virtual device of the front-end virtual machine. For example, this may be achieved by communicating with the at least one back-end driver via the direct communication channel, wherein the direct communication channel is established based on the control information embedded in the predetermined memory area of the front-end virtual machine.

Hence, a method and a virtualization platform according to the present invention improve virtual machine creation, virtual machine migration and/or virtual machine destruction, in particular with regard to the virtual machine creation, virtual machine migration and/or virtual machine destruction times, while retaining the isolation and security afforded by the underlying virtualization platform.

Aspects of the invention may extend to corresponding systems, methods, and computer program products such as computer readable storage media having instructions stored thereon which are operable to program a programmable computer processor to carry out a method as described in the aspects and embodiments that are set out below or recited in the claims and/or to program a

suitably adapted computer to provide the virtualization platform recited in any of the claims.

5 The term “virtualization platform” may refer in particular in the claims, preferably in the description, to an entity comprising software packages that can emulate a whole physical computer machine. Specifically, the entity may provide multiple virtual machines on one physical platform. The virtualization platform can be implemented and provided by a software entity/product such as an emulator and/or hypervisor.

10 The term “toolstack” may refer in particular in the claims, preferably in the description, to at least one interface for hypervisor management and/or configuration of virtual machines. In particular, the toolstack may comprise a collection of tools and/or libraries, which provide interfaces between the hypervisor and the user/administrator of the system to interact and control the hypervisor. Thus, the toolstack may comprise a number of tools and libraries that are built on top of each other to provide certain functionalities. The toolstack being a collection of tools and libraries may be employed and/or may be configured to handle virtual machine management requests such as requesting, by the toolstack, a builder entity to create a new guest virtual machine in response to a user request.

15 The term “hypervisor” may refer in particular in the claims, preferably in the description, to computer software, firmware and/or hardware, which creates and runs virtual machines. A hypervisor may also be designated as a virtual machine monitor (VMM).

20 According to embodiments of the invention the predetermined memory area for the front-end virtual machine may be created by a hypervisor of the virtualization platform. Thus, a virtualization platform can be provided, which employs decentralized interactions between virtual machines for initialization and deinitialization purposes for operations such as creating, destroying and/or migrating virtual machines. This significantly reduces the time it takes for such operations to complete, removes a single point of failure, and allows for highly-

scalable performance by removing the bottleneck of centralized solutions having a centralized information store.

5 According to embodiments of the invention the predetermined memory area dedicated to the control information may be created for each new front-end virtual machine that is created upon a virtual machine creation command. Thus, the provision of the predetermined memory area for each front-end virtual machine results in significant improvements for boot and migration times, among other metrics.

10 According to embodiments of the invention the predetermined memory area may be set aside for the control information within the memory that is set aside for use by the front-end virtual machine, preferably when the front-end virtual machine is created.

15 According to embodiments of the invention the predetermined memory area may be employed to keep track of the front-end virtual machine's control information about devices – preferably about any device – such as block devices and/or networking devices, which the front-end virtual machine may have. Thus, the control information of a virtual machine comprises information about devices of the virtual machine.

20 According to embodiments of the invention the control information may comprise virtual machine identifiers, virtual device information, information about communication channels for devices, back-end identifiers, event channel identifiers and/or grant references. Further, the virtual device information may comprise a MAC (Media Access Control) address and/or an IP (Internet Protocol) address. Thus, a virtualization platform can be provided that employs decentralized interactions between virtual machines for initialization and
25 deinitialization purposes for operations such as creating, destroying and/or migrating virtual machines. This significantly reduces the time it takes for such operations to complete, removes a single point of failure, and allows for highly-scalable performance by removing the bottleneck of centralized solutions having a centralized information store.

According to embodiments of the invention, it may be provided that the toolstack of the virtualization platform keeps track of which back-end devices are available and is responsible for assigning control information that is needed for a communication between the available back-end devices and the front-end virtual machine. The back-end device represents the physical device being controlled by a back-end driver of the back-end virtual machine. Thus, a virtualization platform can be provided which employs decentralized interactions between virtual machines for initialization and deinitialization purposes for operations such as creating, destroying and/or migrating virtual machines. This significantly reduces the time it takes for such operations to complete, removes a single point of failure, and allows for highly-scalable performance by removing the bottleneck of centralized solutions having a centralized information store.

The use of front-end driver and back-end driver represents an established concept in virtual machine drivers in which a physical device is controlled via a driver (the back-end, e.g. "netback" for network devices) that allows other entities (the front-end drivers, e.g. "netfront" for network devices) to also access the device. The front-end drivers create a virtual device inside virtual machines, which behaves like a physical device. The front-ends and back-ends communicate over a channel so that the virtual machines can use the virtual device to access the physical device's capabilities. For example, a network back-end driver (netback) controls a physical network adapter, and provides a communication interface for network front-end drivers (netfront), which provide virtual network interfaces in other virtual machines, so that they can access the network the physical machine is connected to.

According to embodiments of the invention a hypercall may be employed by a virtual machine in order to write to and/or read from the predetermined memory area. A hypercall is a software trap from a virtual machine to the hypervisor, just as a syscall being a software trap from an application to the kernel. Virtual machines, i.e. domains, will use hypercalls to request privileged operations like, e.g., updating page tables. Like a syscall, the hypercall is synchronous, but the return path from the hypervisor to the virtual machine uses event channels. An event channel is a queue of asynchronous notifications, and notify of the same

sorts of events that interrupts notify on native hardware. When a virtual machine with pending events in its queue is scheduled, the operating system's event-callback handler is called to take appropriate action. Thus, the hypercall is a request from a virtual machine to the hypervisor, for information, or for running a certain action on the virtual machine's behalf. Conceptually similar to a system call, in which an application asks the operating system to, e.g., write data to a hard disk: since the application must not access the hard disk itself for reasons of security and abstraction, it asks the operating system to do so on its behalf. Some operations, such as creating or destroying virtual machines, must not be done by a virtual machine itself; it needs to ask the hypervisor to do so on its behalf (and only virtual machine dom0 may typically ask for virtual machine creation, for example; such requests by other virtual machine, i.e. unprivileged domains, would be rejected.

According to embodiments of the invention the predetermined memory area may be configured to be shared read-only with unprivileged front-end virtual machines. Thus, inadmissible write access and manipulations on the predetermined memory area by unprivileged virtual machines can be prevented.

According to embodiments of the invention the predetermined memory area may be configured to be shared writable with a privileged virtual machine that comprises the toolstack. Thus, the privileged virtual machine where the toolstack is located can control and embed the information stored in the predetermined data area. Hence, data security and data integrity is ensured at the predetermined data area.

According to embodiments of the invention the predetermined memory area may be implemented as a memory page. Thus, a fast and efficient memory access is provided. Furthermore, by using memory pages the memory can be shared and protected from unprivileged accesses at page level.

According to embodiments of the invention a front-end driver of the at least one virtual device may be configured to read the control information from the predetermined memory area and to initiate a communication with a back-end

driver of the back-end virtual machine as provided by the control information. The control information can be employed to directly contact a back-end device of the back-end virtual machine via the direct communication channel. Thus, a virtualization platform can be provided which employs decentralized interactions between virtual machines for initialization and deinitialization purposes for operations such as creating, destroying and/or migrating virtual machines. This significantly reduces the time it takes for such operations to complete, removes a single point of failure, and allows for highly-scalable performance by removing the bottleneck of centralized solutions having a centralized information store.

According to embodiments of the invention an operating system of the front-end virtual machine may be configured to retrieve the control information from the predetermined memory area when the front-end virtual machine boots up. Thus, a significant improvement is provided for boot and migration times, among other metrics.

According to embodiments of the invention the embedding of the control information into the predetermined memory area, the retrieving of the control information from the predetermined memory area by the front-end virtual machine and the performing of the communication between the front-end virtual machine and the back-end virtual machine may be employed for creating, destroying and/or migrating the front-end virtual machine. Thus, the time it takes for these operations to complete can be significantly reduced. Further, a single point of failure is removed, and highly-scalable performance is provided by removing the bottleneck of centralized solutions having a centralized information store.

According to embodiments of the invention the toolstack may embed the control information in the predetermined memory area by means of a notification mechanism. Thus, an efficient data exchange of the control information can be provided.

According to embodiments of the invention a toolstack of the virtualization platform may inform the back-end virtual machine that is in control of a physical back-end device about a new front-end virtual machine having a virtual device being

created. The toolstack may receive new control information on how the new front-end virtual machine can access the back-end virtual machine's back-end driver that controls the physical back-end device. The toolstack may write the new control information into the predetermined memory area inside the new front-end virtual machine before the new front-end virtual machine starts its booting process. Further, the new front-end virtual machine's operating system may access the new control information to create a direct communication channel with the back-end virtual machine, wherein further communication including finishing initialization of the front-end virtual machine is done via the direct communication channel. Thus, a single point of failure is removed, and highly-scalable performance is provided by removing the bottleneck of centralized solutions having a centralized information store.

An embodiment of the present invention may provide a modification of a virtualization platform's toolstack such that it is able to embed control information – commonly placed in a back-end store or central information store – directly in a virtual machine, and more specifically in a special, read-only memory page. Furthermore, it may be provided that back-ends (e.g. for storage, networking, etc.) are modified to no longer rely on a back-end store or central information store.

There are several ways how to design and further develop the teaching of the present invention in an advantageous way. To this end it is to be referred to the patent claims subordinate to patent claim 1 on the one hand and to the following explanation of further embodiments of the invention by way of example, illustrated by the figure on the other hand. In connection with the explanation of the further embodiments of the invention by the aid of the figure, generally further embodiments and further developments of the teaching will be explained.

In the drawings

Fig. 1 is a schematic view illustrating an overview of the Xen architecture including toolstack, the XenStore, software switch and split virtual drivers between the driver domain dom0 and the guests,

Fig. 2 is a schematic view illustrating an architecture overview for a method or a system providing an improved and adapted implementation of Xen's standard toolstack,

5 Fig. 3 is a schematic view illustrating an architecture overview for a method or a virtualization platform according to an embodiment of the present invention,

10 Fig. 4 is a schematic view illustrating an architecture overview for a method or a virtualization platform according to an embodiment of the present invention,

15 Fig. 5a is a schematic view illustrating a standard virtual machine creation process in Xen using the XenStore,

Fig. 5b is a schematic view illustrating a virtual machine creation process using a method or a virtualization platform according to an embodiment of the present invention,

20 Fig. 6a is a schematic view illustrating a standard virtual machine creation process using a standard virtualization platform having a central information store,

25 Fig. 6b is a schematic view illustrating a virtual machine creation process using a method or a virtualization platform according to an embodiment of the present invention, and

30 Fig. 7 is a diagram illustrating the quantification of faster virtual machine creation, destruction and migration times of an embodiment according to the present invention compared to the standard virtual machine platform Xen having the XenStore as back-end store.

Fig. 1 shows a Xen architecture including toolstack, the XenStore, software switch and split virtual drivers between the driver domain (dom0) and the guest domains.

The term “domain” is primarily used when using the Xen hypervisor and is synonymous with the term “virtual machine”.

Fig. 1 shows several building blocks which may be described as follows:

5

Building blocks “Dom0” (reference sign 1) and respectively “DomU 1” (reference sign 2) represents a domain, which is a virtual machine in the context of the Xen terminology.

10

In particular “Dom0” is the control domain in context of the Xen terminology, from the fact that the first started domain (having number 0) is awarded special privileges, such as being allowed to instruct the hypervisor to start and stop additional virtual machines. Hence, Dom0 is the administrative domain of Xen.

15

In particular “DomU” is an unprivileged domain in context of the Xen terminology, i.e. every domain other than “Dom0”.

20

Building block “Xen Hypervisor” (reference sign 3) represents a system running on a physical machine and which is in charge of, e.g., creating, destroying, scheduling virtual machines. To put it more simply, the Xen Hypervisor is an operating system for operating systems. The Xen Hypervisor may also be designated as virtual machine monitor.

25

Building block “xl” (reference sign 4) represents a command-line tool that a system administrator can use to run commands such as “create a domain”, “destroy a domain”, etc.

30

Building block “libxl” (reference sign 5) represents a library that the xl command-line tool interacts with, and which in turn interacts with libraries libxc and libxs to facilitate xen-related tasks such as domain creation or destruction.

Building block “libxc” (reference sign 6) represents a Xen control library. Specifically, libxc is a library that comprises, among other things, the code to interact with the hypervisor, by providing an interface to the various hypercalls

required to control the system, in particular with regard to domain created, destruction, etc. A hypercall is a request from a domain to the hypervisor, for information, or for running a certain action on the domain's behalf. Conceptually similar to a system call, in which an application asks the operating system to, e.g., write data to a hard disk: since – in the context of the Xen architecture – the application must not access the hard disk itself for reasons of security and abstraction, it asks the operating to do so on its behalf. Some operations, such as creating or destroying domains, must not be done by a domain itself; it needs to ask the hypervisor to do so on its behalf (and only dom0 may typically ask for domain creation, for example; such requests by other domains would be rejected).

Building block “libxs” (reference sign 7) represents a XenStore library, in particular a library that provides an application programming interface (API) to interact with the XenStore.

Building block “SW switch” (reference sign 8) represents a software switch, which is a software implementation realizing network switching between its different ports. The physical network interface card NIC is connected to the virtual network devices via the software switch to allow control over the data flow between physical and virtual devices and for example also between different virtual devices of virtual machines.

Building block “XenStore” (reference sign 9) represents an implementation of an information storage. While the XenStore is primarily an information store, the fact that entities such as a netback driver can register watches to be informed of new information being put into the XenStore means that the XenStore, in a way, can also trigger events happening. So in the case of the XenStore, the netback registers a watch on the part that deals with network devices, so that whenever the toolstack wants to create a new virtual network device, it can simply write the information (e.g., which MAC address the device should have) into the XenStore, and the watch will trigger an event that tells the netback driver of this, which can then create a new virtual network device.

Building block “NIC” (reference sign 10) represents a network interface card (NIC), which is a network adapter used for network communication. The NIC driver is the driver in charge of controlling the NIC.

- 5 Building block “block” (reference sign 11) represents a block device, which is a term for a device that allows reading and writing to it. A typical example can be a hard disk, on which blocks of data are written and read from.

10 Building blocks “netback” (reference sign 12) and “netfront” (reference sign 13) represent split virtual drivers. A virtual driver may consist of two split virtual drivers, which can be designated as netback driver and netfront driver, respectively. The netback driver is the back-end of the driver for virtual network devices. This portion of the driver exports a unified network-device interface that can be accessed by any operating system, e.g. Linux (reference sign 14), that implements a compatible front-end. For example, when spawning a virtual machine that needs network access, a virtual network device is created. The virtual machine in charge of the physical network interface card (NIC) uses the back-end to create such a virtual network device. In this regard, it is noted that the virtual machine in charge of the physical network interface is generally the privileged control virtual machine, 15 though this does not necessarily have to be the case. The virtual machine that is being created then uses the netfront driver to interact with that virtual device, much like a normal NIC driver would do with a physical NIC. This way, the data from the newly created virtual machine flows via its netfront driver to the control domain's netback, and from there via the NIC driver out of the physical machine. Similar driver concepts exist for block devices, where a blockfront and blockback model can be created to create a virtual hard disk. 25

30 Building block “xenbus” (reference sign 15) represents a driver that provides an interface between the XenStore and the drivers for virtual devices, e.g. netback and netfront, which are illustrated in the Fig. 1, Fig. 2 and Fig. 3 as examples. Thus, the information from the drivers such as how to access the virtual network device can be stored in the XenStore and be read from there.

Building block “toolstack” (reference sign 16) represents a collection of tools and/or libraries, which provide interfaces between the hypervisor and the user/administrator of the system to interact and control the hypervisor. Thus, the toolstack may be a number of tools and libraries that are built on top of each other to provide certain functionalities. The toolstack being a collection of tools and libraries may be employed and/or may be configured to handle virtual machine management requests such as requesting, by the toolstack, a builder entity to create a new guest virtual machine in response to a user request. In the context of the embodiment illustrated by Fig. 1 the toolstack relates to the part of the control flow required to, e.g., create and/or destroy domains, which resides within the user space of the control domain.

In the context of the Xen architecture, e.g. as illustrated in Fig. 1, the toolstack is a component of Dom0, which is responsible for creating, destroying, and managing the resources and privileges of virtual machines. To create a new virtual machine a user provides a configuration file describing memory and CPU allocations as well as device configurations. Then the toolstack parses this file and writes this information in the XenStore. The toolstack takes advantage of Dom0 privileges to map guest memory, to load a kernel and virtual BIOS, and to set up initial communications channels with the XenStore and with the virtual console when a new virtual machine is created.

In the example of Fig. 1, the toolstack comprises building blocks xl, libxl, libxc and libxs.

The Xen hypervisor as illustrated by Fig. 1 is a type-1 hypervisor in charge of managing basic resources such as CPUs and memory. When it finishes booting, it automatically creates a special virtual machine called dom0. Dom0 typically runs Linux and hosts the toolstack, which includes the xl command and the libxl and libxc libraries needed to carry out commands such as virtual machine creation, migration and shutdown.

In addition, dom0 hosts the XenStore, a proc-like central registry that keeps track of management information such as which virtual machines are running in the

system and device information, along with the libxs library containing code to interact with it. The XenStore provides watches that can be associated with particular directories of the store and that will trigger callbacks whenever those directories are read or written to.

5

Typically, dom0 also hosts a software switch (Open vSwitch is the default) to mux/demux packets between NICs and the VMs, as well as the (Linux) drivers for the physical devices. Strictly speaking, this functionality can be put in a separate virtual machine. For communication between dom0 and the other guests, Xen implements a split-driver model: a virtual back-end driver running in dom0 (e.g., the netback driver in the case of networking) communicates over shared memory with a front-end driver running in the guests (the netfront driver). So-called event channels, essentially software interrupts, are used to notify drivers about the availability of data.

15

Fig. 2 shows an architecture overview for a method or a system that provides an improved and adapted implementation of Xen's standard toolstack.

20

Some of the building blocks illustrated in Fig. 2 correspond to the building blocks of Fig. 1. However, according to the embodiment of Fig. 2, the toolstack is modified compared to the standard platform architecture of Fig. 1. In particular the embodiment of Fig. 2 shows other building blocks as follows:

25

Building block "chaos" (reference sign 17) represents a command-line tool that replaces the building block xl of Fig. 1. Building block chaos provides the same functionality as xl, however, chaos is just more efficient with regards to, e.g., domain creation time (i.e., its use results in smaller times).

30

Building block "chaos daemon" (reference sign 18) represents a daemon – i.e. a program running in the background – whose job it is to run the prepare phase and thus prepare a number of virtual machine shells as well as keeping them ready to be used to create new domains by running the execute phase on them.

Building block “libchaos” (reference sign 19) represents a library, which in the case of the split toolstack the command line tool chaos and the chaos daemon interact with, and which in turn interacts with libxc and libxs in order to facilitate xen-related tasks such as domain creation or destruction.

5

Thus, in the embodiment of Fig. 2 the toolstack comprises building blocks chaos daemon, chaos, libchaos, libxc and libxs for providing the concept of a split toolstack (reference sign 20).

10

The embodiment of Fig. 2 considers and improves shortcomings of conventional virtual machine platforms where a significant portion of the overheads related to virtual machine creation and other operations comes from the toolstack itself. More specifically, it turns out that a significant portion of the code that executes when e.g. a virtual machine creation command is issued does not actually need to run at virtual machine creation time. This is because this code is common to all virtual machines, and so can be pre-executed and thus off-loaded from the creation process in accordance with the implementation of the embodiment illustrated by Fig. 2.

15

20

Hence, according to the embodiment of Fig. 2, the standard Xen toolstack is replaced with the libchaos library and is split into two phases, namely into a prepare phase and a execute phase. The prepare phase is responsible for functionality common to all virtual machines such as having the hypervisor generate an ID and other management information and allocation CPU resources to the virtual machine. This functionality is offloaded to the chaos daemon, which generates a number of virtual machine shells and places them in a pool. The daemon may ensure that there are always a certain number of configurable virtual machine shells available in the system.

25

30

The execute phase then begins when a virtual machine creation command is issued. First, chaos contacts the daemon and asks for one of these virtual machine shells to be removed from the pool. It then completes this shell with virtual machine specific operations such as parsing its configuration file, building initialization its devices and booting the virtual machine.

The term “virtual machine shell” may refer to an empty virtual machine or rather to a partially initialized virtual machine for which all basic and generic steps have been done, and which only needs to be filled with an image and some device-specific configuration, e.g. MAC addresses for network devices, to create a fully functional virtual machine. The virtual machine shell can be pre-created before a virtual machine has to be started. Then, the virtual machine shell can be used when a virtual machine creation command is received, to speed up the creation process by only having to execute the execute phase of the toolstack. Hence, the virtual machine shell is a pre-created virtual machine that can be filled with content. The virtual machine shell is a partially-created virtual machine that waits to be filled with execution logic – that is at a minimum an operating system – in order to form a fully functional virtual machine.

The embodiment of Fig. 2 provides a process for virtual machine creation. As already described with regard to the architecture of Fig. 1, one of the key elements of a virtualization platform such as Xen, KVM, VMWare, etc., is the toolstack, which may comprise a command line tool (e.g., “xl” in Xen or “qemu-kvm” in KVM) and a set of libraries in charge of executing operations such as creating virtual machines, destroying them, migrating them, and accessing their console, etc.

The embodiment illustrated by Fig. 2 optimizes virtual machine creation times by means of splitting the toolstack’s functionality into a set of functional processes each containing a number of code functions $fn_1 \dots fn_N$, and further classifying those into

- (1) a prepare phase that may be executed periodically and offline, and
- (2) an execute phase that is executed when the command to create a virtual machine is received.

The prepare phase of the embodiment illustrated by Fig. 2 comprises five functional processes that are common to virtual machines:

- 5
- (1) Process of hypervisor reservation: Hypervisor reservation comprises such jobs as creating the internal control structures, allocating information such as a unique ID to the virtual machine to be started, and filling that information into those control structures.
- (2) Process of compute allocation: Compute allocation comprises e.g. creating virtual CPUs and allocating them to the virtual machine.
- 10 (3) Process of memory reservation: Memory reservation comprises e.g. setting aside memory for the virtual machine in which it can then run.
- (4) Process of memory preparation: Memory preparation comprises e.g. setting up the page table to allow paged memory access and scrubbing memory if needed (if memory has been used by other virtual machines before, to prevent information leakage between virtual machines).
- 15
- (5) Process of device pre-creation: Device pre-creation comprises e.g. creating a console and a virtual network card for eventual interfacing with real network hardware. This creates the virtual machine shell which is a partially initialized virtual machine for which all basic and generic steps have been done, and which only need to be filled with an image and some device-specific configuration (e.g., MAC addresses for network devices) to create a fully functional virtual machine.
- 20
- 25

The virtual machine shells generated by the process of device pre-creation can be pre-created before a virtual machine has to be started, and are then used when a virtual machine creation request is received, to speed up the virtual machine creation process by only having to execute the execute phase.

30

The execute phase of the embodiment illustrated by Fig. 2 comprises the functional processes as follows:

(6) Process of configuration parsing: Configuration parsing comprises e.g. that the configuration file describing the virtual machine (e.g., image, required memory, required device and their configuration parameters) is parsed.

5

(7) Process of device initialization: Device initialization comprises e.g. that the devices are initialized according to the information acquired in the process of configuration parsing.

10

(8) Process of image build: Image build comprises e.g. that the image, which contains the operating system and potentially a virtualized disk containing further tools, is opened, that its format is parsed, and that its data is loaded into the virtual machine memory according to the image's specification.

15

(9) Process of virtual machine boot: Virtual machine boot comprises e.g. that the virtual machine finishes its creation, and starts its boot process. At this point, the control is handed over to the virtual machine and the toolstack's creation step is terminated.

20

According to the embodiment of Fig. 2, the pre-created/prepared virtual machine shells are kept track of, so that at any point, it is clear which virtual machine shells have been pre-created/prepared and which can thus be used for creating a virtual machine in them. This information can be kept in numerous ways, from a dedicated, fixed location in the hypervisor's memory over a file on the management domains file system.

25

According to the embodiment of Fig. 2, the use of a pre-create daemon is provided. The daemon is a background server process running inside the management domain, with which the toolstack can interact.

30

The daemon will aim to pre-create a number of virtual machine shells defined by the operator, keeping them available for future use. Whenever a virtual machine is to be created, the toolstack requests such a virtual machine shell from the daemon

to fill it with the operating system and start the virtual machine. When the daemon runs low on virtual machine shells, it will aim to fill up its virtual machine shell backlog by pre-creating/preparing new virtual machine shells without impeding the performance and behavior of the already running virtual machines, e.g., by pre-creating virtual machine shells during times of low system load.

Fig. 3 shows an architecture overview for a method or a virtualization platform according to an embodiment of the present invention.

- Some of the building blocks illustrated in Fig. 3 correspond to the building blocks of Fig. 1 and Fig. 2. However, according to the embodiment of Fig. 1, the toolstack is modified compared to the standard platform architecture of Fig. 1. In particular, the embodiment of Fig. 3 shows other building blocks as follows:
- Building block “chaos” (reference sign 17) represents a command-line tool that replaces the building block xl of Fig. 1. Building block chaos provides the same functionality as xl, however, chaos is just more efficient with regards to, e.g., domain creation time (i.e., its use results in smaller times).
- Building block “chaos daemon” (reference sign 18) represents a daemon – i.e. a program running in the background – whose job it is to run the prepare phase and thus prepare a number of virtual machine shells as well as keeping them ready to be used to create new domains by running the execute phase on them.
- Building block “libchaos” (reference sign 19') represents a library, which in the case of the split toolstack the command line tool chaos and the chaos daemon interact with, and which in turn interacts with libxc and NOXS (reference sign 21) in order to facilitate tasks such as domain creation or destruction.
- Thus, in the embodiment of Fig. 3 the toolstack comprises building blocks chaos daemon, chaos, libchaos and libxc for providing the concept of a split toolstack (reference sign 20').

Building block “NOXS” (reference sign 21) represents a lean driver for replacing the XenStore (reference sign 9). While the XenStore is primarily an information store, the fact that entities such as the netback driver can register watches to be informed of new information being put into the XenStore means that the XenStore, in a way, can also trigger events happening. So in the case of the XenStore, the netback registers a watch on the part that deals with network devices, so that whenever the toolstack wants to create a new virtual network device, it can simply write the information (e.g., which MAC address the device should have) into the XenStore. The watch will trigger an event that tells the netback driver of this, which can then create a new virtual network device. With store-less operation, as provided by the embodiment of Fig. 3, this functionality naturally goes away, so that a way for the toolstack to directly interact with drivers such as the netback driver is needed. The driver NOXS as illustrated in Fig. 3 is simply an interface that allows the toolstack to interact with the different virtual drivers in a convenient way whenever it, for example, wants to create a new virtual network device.

Fig. 4 shows an architecture overview for a method or a virtualization platform according to a further embodiment of the present invention.

Some of the building blocks illustrated in Fig. 4 correspond to the building blocks of Fig 1, Fig. 2 and Fig. 3. However, according to the embodiment of Fig. 1, the toolstack is modified compared to the standard platform architecture of Fig. 1. In particular, the embodiment of Fig. 4 shows other building blocks as follows:

Building block “chaos” (reference sign 17) represents a command-line tool that replaces the building block xl of Fig. 1. Building block chaos provides the same functionality as xl, however, chaos is just more efficient with regards to, e.g., domain creation time (i.e., its use results in smaller times).

Building block “libchaos” (reference sign 19”) represents a library, which interacts with the command line tool chaos, and which in turn interacts with libxc and NOXS (reference sign 21) in order to facilitate tasks such as domain creation or destruction.

In the embodiment of Fig. 3 the toolstack comprises building blocks chaos daemon, chaos, libchaos and libxc for providing the concept of a split toolstack (reference sign 20”).

- 5 Building block “NOXS” (reference sign 21) corresponds to the NOXS driver of the embodiment of Fig. 3.

10 Thus, according to the embodiment of Fig. 4, libxl and the corresponding xl command is replaced with a streamlined, thin library and command called libchaos and chaos, respectively; these no longer make use of the XenStore and its accompanying libxs library. Further, the toolstack keeps track of which back-ends are available (e.g., network, block) and is responsible for assigning any information needed for communication between them and the guests (e.g., event channels and grant references).

15 Fig. 5a and Fig. 5b shows a standard virtual machine creation process in Xen using the XenStore in Fig. 5a versus a virtual machine creation process using a method or a virtualization platform according to an embodiment of the present invention in Fig. 5b.

20 Having identified the major bottlenecks in the standard system of Xen as illustrated by Fig. 5a, an implementation according to the embodiment of Fig. 5b introduces a re-design of some of the basic Xen mechanisms in order to provide lightweight virtualization over that virtualization platform.

25 There is a main obstacle to implementing lightweight virtualization on Xen, namely the XenStore, since many operations require multiple interactions with it, a problem that worsens with an increasing number of virtual machines.

30 To address this, according to the embodiment illustrated by Fig. 5b significant portions of the Xen ecosystem is redesigned. The XenStore is removed. Specifically, the XenStore is replaced with a lean driver called noxs and, additionally, a split toolstack that separates functionality between a prepare phase

and an execute phase may be provided. Furthermore, a much leaner chaos command-line tool and libchaos library than the standard xl/libxl are used.

5 The XenStore is crucial to the way Xen functions, with many xl commands making heavy use of it. As way of illustration, Fig. 5a shows the process when creating a virtual machine and its (virtual) network device. First, the toolstack writes an entry to the network back-end's directory, essentially announcing the existence of a new virtual machine in need of a network device (step 1 in Fig. 5). Previous to that, the back-end placed a watch on that directory; the toolstack writing to this directory
10 triggers the back-end to assign an event channel and other information (e.g., grant references, a mechanism for sharing memory between guests) and to write it back to the XenStore (step 2 in Fig. 5a). Finally, when the virtual machine boots up, it contacts the XenStore to retrieve the information previously written by the network back-end (step 3 in Fig. 5a).

15 The above is a simplification: in actuality, the virtual machine creation process alone can require interaction with over 30 XenStore entries, a problem that is exacerbated with increasing number of virtual machines and devices. Worse, the XenStore represents a single point of failure.

20 As it turns out, most of the necessary information about a virtual machine is already kept by the hypervisor (e.g., the virtual machine's id, but not the name, which is kept in the XenStore but is functionally not strictly needed). This insight here is that the hypervisor already acts as a sort of centralized store, so its functionality is extended to implement the noxs (no XenStore) mechanism. Thus,
25 according to the embodiment of Fig. 5b, it is possible to forego the use of the XenStore for operations such as creation, pause/unpause and migration.

30 Specifically, libxl and the corresponding xl command is replaced with a streamlined, thin library and command called libchaos and chaos, respectively (cf. Fig. 3 and Fig. 4); these no longer make use of the XenStore and its accompanying libxs library. Further, the toolstack keeps track of which back-ends are available (e.g., network, block) and is responsible for assigning any

information needed for communication between them and the guests (e.g., event channels and grant references).

5 In addition, Xen's hypervisor is modified to create a new, special device memory page for each new virtual machine that is used to keep track of a virtual machine's information about any devices, such as block and networking, which it may have.

10 Furthermore, a hypercall is included to write to and read from this memory page, and make sure that the page is shared read-only with guests (but read/write with dom0).

15 When a chaos create command is issued, the tool-stack first requests the creation of devices from the back-end(s) through a system call (an ioctl) handled by the NOXS Linux kernel module (step 1 in Fig. 5b). This mechanism may work if the back-ends run in dom0. Furthermore, this mechanism may also work if back-ends run on a different virtual machine. The back-end then returns the details about the communication channel for the front-end. Second, the toolstack calls the new hypercall asking the hypervisor to add these details to the device page (step 2 in Fig. 5b).

20 When the virtual machine boots, instead of contacting the XenStore, it will ask the hypervisor for the address of the device page and will map the page into its address space using hypercalls (step 3 in Fig. 5b); this requires modifications to the guest's operating system. The guest will then use the information in the page to initiate communication with the back-end(s) by mapping the grant and bind to the event channel (step 4 in Fig. 5b). At this stage, the front and back-ends setup the device by exchanging information such as its state and its MAC address (for networking); this information was previously kept in the XenStore and is now stored in a device control page pointed to by the grant reference. Finally, front and back-ends notify each other of events through the event channel, which replaces the use of XenStore watches.

30 To support migration without a XenStore a new pseudo-device to handle power-related operations can be created. This device can be implemented, following

Xen's split driver model, with a back-end driver (power-back) and a front-end (powerfront) one. These two devices can share a device page through which communication happens.

5 With this in place, migration may begin by chaos opening a TCP connection to a migration daemon running on the remote host and by sending the guest's configuration so that the daemon pre-creates the domain and creates the devices. Next, to suspend the guest, chaos issues an ioctl to the back-end which will set a field in the shared page to denote that the shutdown reason is suspend. The front-
10 end will receive the request to shutdown, upon which the guest will save its internal state and unbind noxs-related event channels and device pages. Once the guest is suspended it may be relied on libxc code to send the guest data to the remote host.

15 Fig. 6a shows a standard virtual machine creation process using a standard virtualization platform having a central information store. The standard virtual machine boot process of Fig 6a involves an information store during the creation of the virtual machine dom2.

20 Step (1): The toolstack writes information about the new front-end virtual machine dom2 into the information store.

Step (2): The back-end virtual machine dom1 that has control of the physical network device receives a notification and puts information about the
25 network device into the information store.

Step (3): The new front-end virtual machine dom2 receives information from the information store to initialize its virtual network device, and puts its own initialization information in there.

30

Steps (2) and (3) iterate until the back-end virtual machine (dom1) and front-end virtual machine (dom2) have finished their initialization with each other.

Step (4): During initialization, a direct communication channel was set up that is now used for the data exchange during the virtual machine lifetime.

Fig. 6b shows a virtual machine creation process using a method or a virtualization platform according to an embodiment of the present invention, wherein a central information store is removed.

Step (1): The toolstack informs the back-end virtual machine dom1 in control of the physical network device about a new front-end virtual machine dom2 with a virtual network device being created. The toolstack receives information on how the new front-end virtual machine dom2 can access the network driver of the back-end virtual machine dom1.

Step (2): The toolstack writes this information into a memory area inside / dedicated to the new front-end virtual machine dom2 before it starts the booting process.

Step (3): The front-end virtual machine's operating system accesses this information to create a direct communication channel with the back-end virtual machine dom1.

Step (4): All further communication including finishing the initialization is done via a direct communication channel.

Thus, the embodiment of Fig. 6b provides a method whereby a virtualization platform is able to function without a centralized information store such as the XenStore of Xen. This is implemented by embedding control information directly into the new or running front-end virtual machine dom2, in a memory area set aside for this task within the memory set aside for use by the front-end virtual machine dom2, as the front-end virtual machine is created (cf. Fig. 6b). The control information should include, at a minimum, information about the front-end virtual machine's devices (e.g., for network or storage I/O). To accomplish this, the toolstack is modified to embed the above-mentioned information into the front-end virtual machine's memory. Further, the front-end virtual machine operating system

is modified such that when the front-end virtual machine starts, it no longer tries to contact a central information store: instead, the front-end virtual machine retrieves the control information from the predetermined memory area and uses that to directly contact any relevant communication back-end devices or rather its back-end drivers located at back-end virtual machines.

Thus, the information store is completely removed from the process, reducing virtual machine creation, destruction and migration times. As an extension, the process of embedding control information directly in the virtual machine page might be off-loaded from the local toolstack to an external entity such as an orchestrator framework (e.g., OpenStack). As a further extension, it would be possible for the toolstack to modify the special memory page by means of using a notification mechanism to let the virtual machine know that changes have been made and that the memory page, i.e. the memory area, needs re-reading; e.g., in case the physical devices change because of hot-plugging or migration of the virtual machine to a new physical machine, requiring re-initialization

Fig. 7 shows a diagram illustrating the quantification of faster virtual machine creation, destruction and migration times of an embodiment according to the present invention compared to the standard virtual machine platform Xen having the XenStore as back-end store.

The diagram of Fig. 7 plots a cumulative probability against a domain creation time in milliseconds.

To quantify faster virtual machine creation, destruction and migration times that are provided by an embodiment according to the present invention, Cumulative Distribution Function (CDF) is included when creating 1000 virtual machines as fast as possible on a single server using the Xen virtualization platform and a minimalistic virtual machine such as a unikernel. The standard platform (labeled "Back-end Store" in Fig. 4; reference sign 22) is compared against the method according to an embodiment of the present invention (labeled "No Back-end Store" in Fig. 4; reference sign 23). As illustrated by the diagram of Fig. 7, by using an embodiment according to the present invention practically all of the virtual

machines boot in as little as 10 milliseconds; without it, the CDF shows a much longer tail (up to 300 milliseconds) and boot creation times that take at least three times as long as with a method according to an embodiment of the invention.

- 5 Many modifications and other embodiments of the invention set forth herein will come to mind to the one skilled in the art to which the invention pertains having the benefit of the teachings presented in the foregoing description and the associated drawings. Therefore, it is to be understood that the invention is not to be limited to the specific embodiments disclosed and that modifications and other embodiments
- 10 are intended to be included within the scope of the appended claims. Although specific terms are employed herein, they are used in a generic and descriptive sense only and not for purposes of limitation.

List of reference signs

1	Dom0
2	DomU 1
3	Xen Hypervisor
4	xl
5	libxl
6	libxc
7	libxs
8	SW switch
9	XenStore
10	NIC
11	block
12	netback
13	netfront
14	Operating system Linux
15	xenbus
16	toolstack
17	chaos
18	chaos daemon
19	libchaos
19'	libchaos
19''	libchaos
20	split toolstack
20'	split toolstack
20''	split toolstack
21	NOXS driver
22	Evaluation of Back-end Store
23	Evaluation of No Back-end Store

Claims

1. Method for operating virtual machines on a virtualization platform, the method comprising:

5 embedding, preferably by a toolstack of the virtualization platform, control information in a predetermined memory area of a front-end virtual machine where at least one virtual device is to be initialized, wherein said control information is required for initiating a communication with a back-end virtual machine where at least one back-end driver runs;

10 retrieving, by the front-end virtual machine, said control information from said predetermined memory area of the front-end virtual machine; and

performing the communication between the front-end virtual machine and the back-end virtual machine via a direct communication channel in order to exchange information for initializing said at least one virtual device of the front-end virtual machine, in particular by communicating with said at least one back-end driver via said direct communication channel, wherein said direct communication channel is established based on said control information embedded in said
15 predetermined memory area of the front-end virtual machine.

20 2. Method according to claim 1, wherein said predetermined memory area for the front-end virtual machine is created by a hypervisor of the virtualization platform.

25 3. Method according to claim 1 or 2, wherein said predetermined memory area dedicated to control information is created for each front-end virtual machine that is created upon a virtual machine creation command.

30 4. Method according to any of claims 1 to 3, wherein said predetermined memory area is set aside for said control information within the memory that is set aside for use by the front-end virtual machine.

5. Method according to any of claims 1 to 4, wherein said predetermined memory area is employed to keep track of the front-end virtual machine's control information about devices that the front-end virtual machine has.

- 5 6. Method according to any of claims 1 to 5, wherein the toolstack of the virtualization platform keeps track of which back-end devices are available and is in charge of assigning control information that is needed for a communication between the available back-end devices and the front-end virtual machine.
7. Method according to any of claims 1 to 6, wherein a hypercall is employed in order to write to and/or read from said predetermined memory area.
- 10 8. Method according to any of claims 1 to 7, wherein said predetermined memory area is configured to be shared read-only with unprivileged front-end virtual machines, and/or
wherein said predetermined memory area is configured to be shared writable with a privileged virtual machine that comprises the toolstack.
- 15 9. Method according to any of claims 1 to 8 , wherein a front-end driver of said at least one virtual device is configured to read said control information from said predetermined memory area and to initiate a communication with a back-end driver of the back-end virtual machine as provided by said control information.
- 20 10. Method according to any of claims 1 to 9, wherein an operating system of the front-end virtual machine is configured to retrieve said control information from said predetermined memory area when the front-end virtual machine boots up.
- 25 11. Method according to any of claims 1 to 10, wherein said embedding, said retrieving and said performing is employed for creating, destroying and/or migrating the front-end virtual machine.
- 30 12. Method according to any of claims 1 to 11, wherein the toolstack embeds said control information in said predetermined memory area by means of a notification mechanism.
13. Method according to any of claims 1 to 12, wherein a toolstack of the virtualization platform informs the back-end virtual machine that is in control of a

physical back-end device about a new front-end virtual machine having a virtual device being created,

wherein the toolstack receives new control information on how the new front-end virtual machine can access the back-end virtual machine's back-end driver that controls the physical back-end device,

wherein the toolstack writes the new control information into the predetermined memory area inside the new front-end virtual machine before the new front-end virtual machine starts its booting process,

wherein the new front-end virtual machine's operating system accesses the new control information to create a direct communication channel with the back-end virtual machine,

wherein further communication including finishing initialization of the front-end virtual machine is done via the direct communication channel.

14. Virtualization platform for operating virtual machines, in particular for executing a method according to any of claims 1 to 13, comprising a hypervisor that is configured to run one or more back-end virtual machines and one or more front-end virtual machines,

wherein an entity, in particular a toolstack, of said virtualization platform is configured to embed control information in a predetermined memory area of a front-end virtual machine where at least one virtual device is to be initialized, wherein said control information is required for initiating a communication with a back-end virtual machine where at least one back-end driver runs,

wherein the front-end virtual machine is configured to retrieve said control information from said predetermined memory area of the front-end virtual machine, and

wherein the front-end virtual machine and the back-end virtual machine are configured to perform the communication between the front-end virtual machine and the back-end virtual machine via a direct communication channel in order to exchange information for initializing said at least one virtual device of the front-end virtual machine, in particular by communicating with said at least one back-end driver via said direct communication channel, wherein said direct communication channel is established based on said control information embedded in said predetermined memory area of the front-end virtual machine.

15. Computer program comprising instructions which, when the program is executed by a computer, cause and/or provide the virtualization platform of claim 14 to execute a method according to any of claims 1 to 13.

1/7




Fig. 1




Fig. 2




Fig. 3




Fig. 4




Fig. 5a




Fig. 5b




Fig. 6a




Fig. 6b




Fig. 7

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2017/058445

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/455 G06F9/50 G06F9/44
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	YUN CHAN CHO ET AL: "Sharing data between processes running on different domains in para-virtualized xen", CONTROL, AUTOMATION AND SYSTEMS, 2007. ICCAS '07. INTERNATIONAL CONFERENCE ON, ANONYMOUS, SEOUL, KOREA, 17 October 2007 (2007-10-17), pages 1255-1260, XP031223752, ISBN: 978-89-950038-6-2 page 1255, right-hand column, line 18 - page 1260, right-hand column, line 15; figures 2-10 ----- -/--	1-15



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

2 January 2018

Date of mailing of the international search report

11/01/2018

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

de Junca, Irène

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2017/058445

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	YUEBIN BAI ET AL: "A High Performance Inter-VM Network Communication Mechanism", 21 May 2010 (2010-05-21), ALGORITHMS AND ARCHITECTURES FOR PARALLEL PROCESSING, SPRINGER BERLIN HEIDELBERG, BERLIN, HEIDELBERG, PAGE(S) 363 - 380, XP019142366, ISBN: 978-3-642-13118-9 page 366, line 24 - page 372, line 7; figures 2-6 -----	1-15
A	US 9 507 617 B1 (WANG HSIAO-PENG [TW] ET AL) 29 November 2016 (2016-11-29) abstract; claims 1,6 column 3, line 19 - column 4, line 48; figures 3,4 column 5, line 23 - line 63; figure 6 -----	1-15
A	US 2017/031593 A1 (TSIRKIN MICHAEL [IL] ET AL) 2 February 2017 (2017-02-02) abstract paragraphs [0022] - [0029]; figure 2 paragraph [0031]; figure 3(310) -----	1-15
T	FILIPE MANCO ET AL: "My VM is Lighter (and Safer) than your Container", PROCEEDINGS OF THE 26TH SYMPOSIUM ON OPERATING SYSTEMS PRINCIPLES, SOSOP '17, 28 October 2017 (2017-10-28), pages 218-233, XP055436727, New York, New York, USA DOI: 10.1145/3132747.3132763 ISBN: 978-1-4503-5085-3 page 221, right-hand column, line 7 - page 222, left-hand column, line 2; figure 3 page 223, left-hand column, line 29 - page 224, right-hand column, line 26; figures 6,7 -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2017/058445

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 9507617	B1	29-11-2016	NONE

US 2017031593	A1	02-02-2017	NONE
