

EXTENDED CONFERENCE PAPER

Flexible Service Consolidation with Nested Virtualization and Library Operating Systems

Kenichi Kourai* | Kouta Sannomiya

¹Department of Creative Informatics,
Kyushu Institute of Technology, Fukuoka,
Japan

Correspondence

*Kenichi Kourai, 680-4 Kawazu, Iizuka,
Fukuoka 820-8502, Japan. Email:
kourai@ksl.ci.kyutech.ac.jp

Summary

In Infrastructure-as-a-Service clouds, users can reduce costs by scale-in or -down when running services are under-utilized. Since these optimizations of instance deployment require at least one minimum instance even for running an under-utilized service, cost reduction is limited. For further optimization, multiple services can be consolidated into one instance. However, services have to be stopped temporarily at the consolidation time and isolation between services becomes weaker after the consolidation. To solve these problems, this paper proposes FlexCapsule, which enables seamless and secure service consolidation in existing IaaS clouds. FlexCapsule runs each service in a lightweight virtual machine (VM) called an app VM, using a library operating system. An app VM runs inside an instance using a technique called nested virtualization. FlexCapsule can optimize instance deployment with negligible downtime by flexibly migrating app VMs. Thanks to strong isolation provided by app VMs, it can guarantee security between consolidated services. In addition, FlexCapsule provides multi-process support using app VMs such as process fork and process pools. We have implemented FlexCapsule in Xen using both fully virtualized OS^v and para-virtualized MiniOS. Then we confirmed its effectiveness using several applications.

KEYWORDS:

IaaS clouds, service consolidation, nested virtualization, library operating system, process pool, virtual private network

1 | INTRODUCTION

Infrastructure-as-a-Service (IaaS) clouds provide users with instances, which are usually virtual machines (VMs), and users run their services in instances. Since users can change instance deployment flexibly in IaaS clouds, they can respond to load changes rapidly. For example, users can use minimum instance deployment at start up and increase the number or the resource amount of instances when their services become over-utilized. In contrast, they can decrease the number or the resource amount to reduce costs when their services become under-utilized. Thus it is necessary to optimize instance deployment so that used instances are always sufficient but minimum.

However, it is not easy to perform such optimization in current IaaS clouds. If users adjust the number of instances by scale-out and -in, they need at least one instance even for an under-utilized service and cannot further reduce costs. As an optimization for one instance, users can adjust the amount of resources assigned to an instance by scale-up and -down. Unfortunately, most of the

existing clouds achieve this optimization by switching instance types offline because they do often not provide the function for dynamically changing resource allocation to an instance. Therefore, users need at least one minimum instance even for a mostly idle service. Although the cost of each instance may be low, the total cost could become high if users run many under-utilized services.

For further optimization, users can consolidate services running in multiple instances into one instance. When there are several under-utilized services, the user can run them in one instance and reduce costs. Later, when some of the services become over-utilized, the user can de-consolidate them to other instances. However, this service consolidation and de-consolidation cause service downtime when users move services between instances. This problem can be solved by using process migration¹, but a security issue arises due to consolidating services. Since multiple services run in the same instance, isolation among them becomes weaker than traditional instance-level isolation.

In this paper, we propose FlexCapsule, which achieves seamless and secure service consolidation for optimizing instance deployment. FlexCapsule runs each service in a lightweight virtual machine (VM) called an app VM, using a library operating system (LibOS)^{2,3,4}. To enable an app VM to flexibly run with appropriate resources in existing IaaS clouds, FlexCapsule runs an app VM inside an instance using nested virtualization^{5,6}. Since it constructs a virtual private network (VPN) for all app VMs across multiple instances, app VMs can be seamlessly migrated between instances. Thus FlexCapsule can optimize instance deployment without disrupting services enclosed in app VMs. In addition, it can guarantee security between services consolidated into one instance using strong isolation provided by app VMs.

We have implemented FlexCapsule in Xen 4.2.4⁷. FlexCapsule provides two types of LibOSes running in app VMs. One is based on OS^v 0.21⁴, while the other is based on MiniOS in Xen. We have added migration support to these LibOSes because these LibOSes use para-virtualization techniques to reduce virtualization overhead. As a helper of the LibOS, FlexCapsule provides an OS server running inside each instance. For example, when the fork function in the LibOS is invoked, the OS server clones the entire app VM. When the listen function is invoked, the OS server registers a rule for port forwarding to the app VM. Combining these mechanisms, the OS server enables app VMs to create process pools. According to our experiments, it was shown that FlexCapsule could optimize instance deployment according to performance requirements of services. Migration performance of app VMs was better than the normal VM running Linux, thanks to smaller memory footprints.

This paper is an extended version of our previous conference paper⁸. In this paper, we have developed the FlexCapsule OS based on MiniOS in addition to the OS^v-based one. Whereas OS^v is a fully virtualized LibOS with para-virtual device drivers, MiniOS is a completely para-virtualized LibOS. Therefore, MiniOS can further reduce the overhead and run in smaller memory footprint. Using the MiniOS-based FlexCapsule OS, we conducted several experiments and compared the results with those in the OS^v-based one.

This paper is organized as follows. Section 2 describes issues in optimizing instance deployment in existing IaaS clouds and related work. Section 3 proposes FlexCapsule for enabling seamless and secure service consolidation using app VMs. Section 4 describes its implementation and Section 5 shows experimental results. Section 6 concludes this paper.

2 | BACKGROUND

2.1 | Optimizing Instance Deployment

In IaaS clouds, the optimization of instance deployment is performed according to resource utilizations of instances, as illustrated in Fig. 1. The most popular optimization is scale-out and -in, which adjust the number of instances. When a service becomes over-utilized, the user can increase the number of instances by scale-out and distribute the load to more instances. In contrast, when a service becomes under-utilized, the user can decrease the number of instances by scale-in to reduce costs. However, if only one instance is deployed for an under-utilized service, the user cannot further reduce the number of instances. For example, consider intra servers that are rarely accessed during weekends and personal servers that are sometimes accessed. When there is almost no request to such a server, the system load becomes almost zero, but one instance is required if the server cannot be stopped. Thus the effectiveness of this optimization is limited when services are almost not running.

The optimization for one instance is scale-up and -down, which adjust the amount of resources assigned to each instance. When a service becomes over-utilized, the user can increase the number of virtual CPUs (vCPUs), the performance of vCPUs, and/or the amount of memory of the instance by scale-up. In contrast, when a service becomes under-utilized, the user can decrease the amount of such resources by scale-down to reduce costs. However, most of the existing clouds like Amazon EC2 achieve scale-up and -down by switching instance types offline. Since the user has to choose one from several instance types,

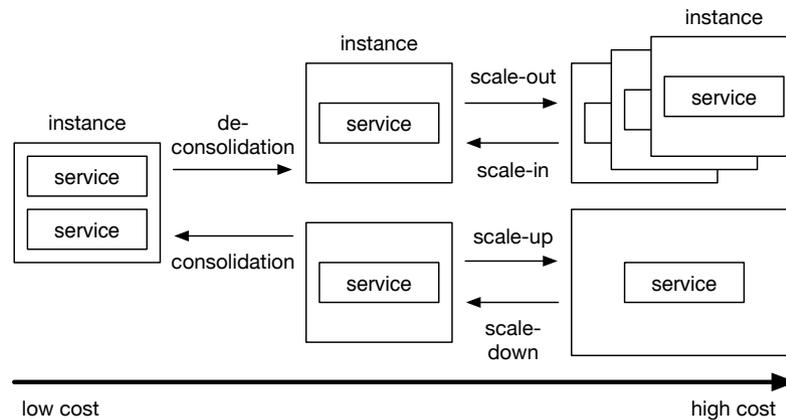


FIGURE 1 Various optimizations of instance deployment.

cost reduction is limited by the cost of the minimum instance type. In addition, when the user switches his current instance to a new one, he has to stop services, move their data to the new instance, and restart these services in the new instance. This duration becomes downtime, for which services cannot be provided.

VMware vCloud Air Virtual Private Cloud OnDemand⁹ supports seamless and flexible scale-up and -down of instances. Users can dynamically increase or decrease the amount of resources assigned to their instances, according to service demands. Unlike most of exiting IaaS clouds, they do not need to create a new instance of appropriate type and move services to it for scale-up and -down. Since users can pay only for assigned resources, not for instances, the cost can be reduced for under-utilized services. This is one implementation of the Resource-as-a-Service cloud¹⁰. However, the number of vCPUs cannot be decreased less than one.

For further optimization, users can consolidate multiple services running in multiple instances into one instance. This is called service consolidation. For example, consider a multi-tier application that consists of multiple services such as a Web server, an application server, and a database. When these services are running across multiple instances and all of them are under-utilized, the user can run these services in one instance to reduce costs. Later, when the instance becomes over-utilized, the user can de-consolidate these services to multiple instances again to perform load balancing. For further cost reduction, even different users could consolidate their services into one instance. However, a security issue arises due to this service consolidation. Since multiple services run in the same instance, isolation among them becomes weaker than when using one instance per service. In addition, consolidation and de-consolidation also cause downtime when the user moves services between instances.

To reduce downtime during service consolidation and de-consolidation and scale-up and -down, process migration¹ can be used. Process migration enables application processes to be moved between hosts without disrupting services. For example, Zap¹¹ provides a thin virtualization layer between processes and the OS and runs a group of processes in a container called a pod. Using pods, Zap enables most of the process state to be maintained on process migration. However, isolation between pods is not strong enough because a pod is protected only by namespaces provided by the OS. PicoCenter¹² runs one service in a Linux container and runs containers inside instances. Instead of process migration, it swaps out idle containers to cloud storage and swaps in them when the services in the containers are used. Since swapped-out containers cannot provide services at all, this mechanism is applicable only to mostly idle services. In addition, isolation between containers is not strong like pods.

2.2 | Nested Virtualization

Nested virtualization⁵ enables VMs and the hypervisor to run inside an outer VM. For Intel x86-based systems, nested virtualization has been developed without hardware support⁶. It virtualizes the CPU instruction set for virtualization by trapping and emulating those instructions. It also virtualizes MMU by using multi-dimensional paging. Nested virtualization is used for various purposes in clouds. For example, xCloud¹³ and Xen-Blanket¹⁴ allow users to use their own hypervisor and implement cross-platform live migration in public clouds. HVX¹⁵ provides a thin virtualization layer using nested virtualization and enables users to run unmodified VMs in almost any existing IaaS clouds. Inception¹⁶ takes one step further and enables users to build nested IaaS clouds on top of public clouds.

Nested virtualization causes larger overhead than traditional single-level virtualization. However, it is reported that the overhead is 6–8% for common workloads⁶. Special-purpose cloud hypervisors as used in CloudVisor¹⁷ and TinyChecker¹⁸ can improve the performance of nested virtualization more. Hardware support for nested virtualization has been also added. For example, Intel VMCS Shadowing¹⁹ can eliminate VM exits due to VMREAD and VMWRITE instructions for accessing VMCS. The ARMv8.3 architecture²⁰ supports nested virtualization and its extension called NEVE has been proposed for coalescing and deferring traps²¹.

2.3 | LibOS

A LibOS² enables each application to link most of the functions of the OS as a library and to perform its own resource management. Developers can customize a LibOS, considering characteristics of each application, and optimize application performance. There are many LibOSes running in VMs on top of the hypervisor. MiniOS in Xen⁷ is a minimal OS and runs only one application in the kernel address space. Libra²², GUK²³, and OS^v⁴ run the Java VM with the LibOS to optimize the execution of Java applications. Libra provides the LibOS with only functions that affect the performance of the Java VM and uses file systems and networks provided externally in Xen. GUK extends MiniOS and improves the memory management of MiniOS. It also adds support for SMP, memory ballooning, and VM suspension and resumption. OS^v can run not only Java applications but also many existing C applications. MirageOS³ specializes the LibOS to OCaml applications and generates a unikernel directly running on the hypervisor.

Although a LibOS can run only one application, there are several studies on multi-process support. Xok/ExOS²⁴ can run the existing Unix applications without modifications using the nano kernel called Exokernel² and the LibOS. ExOS provides mechanisms for multi-process such as process fork and inter-process communication using shared memory. Graphene²⁵ supports multi-process for applications with the Linux-compatible LibOS. It provides an abstraction called a picoprocess²⁶, which runs on top of the host OS. Its applications can perform inter-process communication using RPC via the LibOS. In addition, Graphene achieves process fork and non-live migration by application checkpointing. KylinX²⁷ is based on MiniOS and achieves the fork function by VM fork²⁸. Like normal OS processes, the memory of a VM running the LibOS is shared between parent and child VMs in a copy-on-write manner. KylinX also support shared libraries.

3 | FLEXCAPSULE

This paper proposes FlexCapsule, which enables seamless and secure service consolidation for optimizing instance deployment in IaaS clouds. FlexCapsule runs each service in a lightweight VM called an app VM inside an existing instance. Using VM migration, FlexCapsule can move services between instances with negligible downtime at the optimization time of instance deployment. Thanks to strong isolation between app VMs, FlexCapsule can guarantee security between services consolidated into one instance.

3.1 | System Architecture

Fig. 2 illustrates the system architecture of FlexCapsule. Using nested virtualization⁶, FlexCapsule runs the hypervisor inside each instance, which is usually a VM, and runs app VMs on top of the hypervisor. Since the hypervisor is smaller and simpler than the OS, it is less vulnerable and more secure than the OS. Each app VM runs only one application process and a LibOS, which is linked to the application to provide functions of the OS without any overhead of protection mechanisms. FlexCapsule also runs an OS server in each instance and provides functions that cannot be achieved only by the LibOS inside app VMs.

Since FlexCapsule assumes that public IP addresses are assigned only to instances, it assigns private IP addresses to app VMs. This reduces the cost for using public IP addresses in clouds. To provide services of app VMs to the outside, FlexCapsule uses network address port translation (NAPT). Thanks to NAPT, different app VMs can use the same public IP address. The public IP address that each app VM uses is determined at creation time and does not change. Also, it constructs a network with the same segment across multiple instances using a site-to-site virtual private network (VPN). This VPN enables app VMs to continue to use the same public and private IP addresses even after they are migrated to other instances. Each packet is first delivered to the instance with the specified public IP address. Then it is automatically forwarded to an appropriate instance running the target app VM by the VPN.

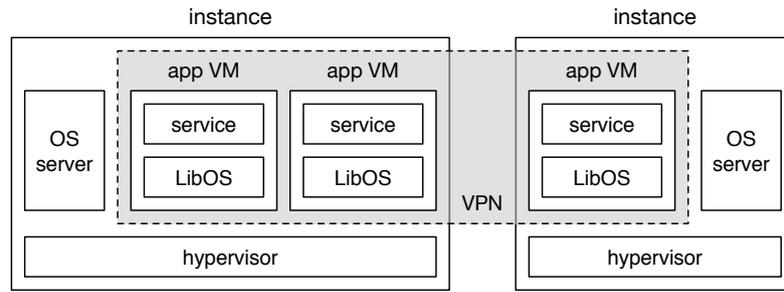


FIGURE 2 The system architecture of FlexCapsule.

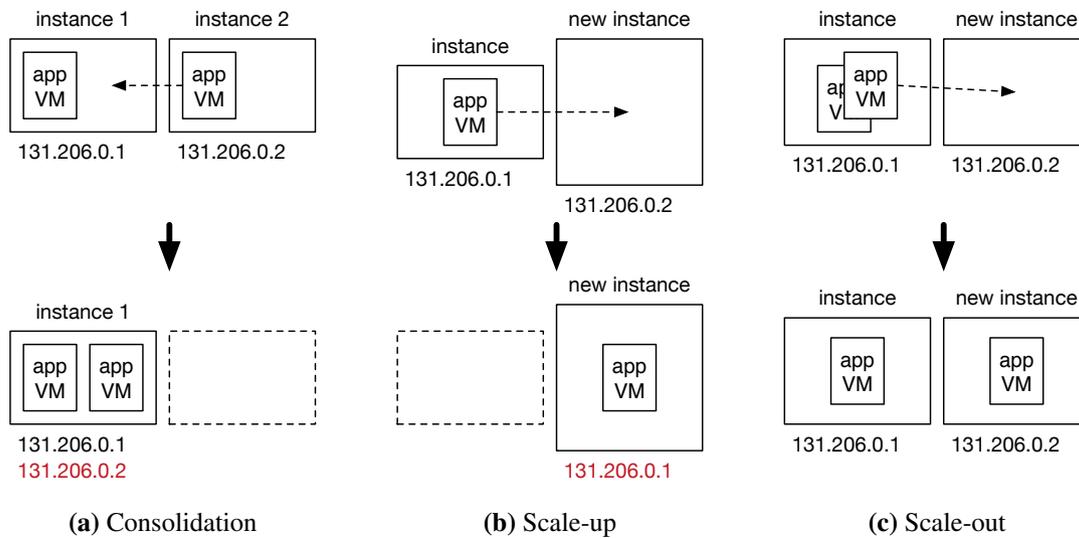


FIGURE 3 The optimization of instance deployment using the migration of app VMs.

3.2 | Optimization Using App VMs

When performing service consolidation, FlexCapsule migrates under-utilized app VMs to one instance, as illustrated in Fig. 3(a). As a result, if the source instances have no app VM, FlexCapsule stops them and re-assigns their public IP addresses to the destination instance, e.g., using Elastic IP addresses in Amazon EC2. Thus migrated app VMs can be reached using the same IP addresses as before service consolidation. In contrast, when performing service de-consolidation, FlexCapsule deploys new instances and migrates over-utilized app VMs to those instances. Before the migration, FlexCapsule re-assigns one of the public IP addresses assigned to the source instance to the destination ones.

To scale a service up and down, FlexCapsule deploys a new instance of appropriate type and migrates app VMs in the original instance to the new one (Fig. 3(b)). Then it stops the original instance and re-assigns that public IP address to the new one. For scaling a service out, on the other hand, FlexCapsule deploys new instances, clones app VMs inside the original instances, and migrates them to the new ones (Fig. 3(c)). At this time, FlexCapsule assigns new private IP addresses to the cloned app VMs but allows them to continue to use the original public IP addresses using NAT. When scaling a service in, FlexCapsule simply stops several instances.

3.3 | FlexCapsule OS

The FlexCapsule OS is a LibOS running in an app VM. It can reduce the memory footprint of an app VM because only necessary functions in the FlexCapsule OS are linked to each application at compile time. Therefore running an app VM per service does

not require extra memory so much, compared with using a general-purpose OS. The small memory footprint of an app VM enables more rapid VM migration by transferring only a smaller amount of memory. Similarly, cloning app VMs becomes faster.

The FlexCapsule OS reduces the overhead of extra virtualization due to app VMs by using para-virtualization. Para-virtualization is a technique that does not emulate real hardware completely to simplify hardware virtualization. Since full virtualization emulates real hardware as it is, the combination with nested virtualization poses larger overhead. Therefore, using para-virtualization can improve the performance of app VMs. In compensation for this performance gain, however, OSes with para-virtualization need to support VM migration by themselves. This is because such OSes are more tightly coupled with the hypervisor and virtual hardware. Specifically, the FlexCapsule OS enables itself to be suspended and resumed.

3.4 | FlexCapsule Server

The FlexCapsule server is an OS server running in each instance. It provides functions related to process management and multi-process, which cannot be supported only by the FlexCapsule OS. For process management, users can obtain the list of running app VMs using the `ps`-like command and terminate specified app VMs using the `kill`-like command. For multi-process support, the FlexCapsule server is used for cooperation between app VMs. For example, when a service in an app VM invokes the `fork` function, the FlexCapsule server clones the entire app VM by VM `fork`²⁸. At this time, the FlexCapsule OS communicates with the FlexCapsule server and then the FlexCapsule server creates a child app VM from the parent app VM. The public IP address that the child app VM uses is the same as that used by the parent app VM. It returns an identifier of the child VM to the parent VM and zero to the child VM, as in the original process `fork`.

The FlexCapsule server also manages NATP rules to forward packets to app VMs. Since an app VM communicates with the outside using NATP, the FlexCapsule server registers a NATP rule when a service invokes the `listen` function. For example, consider a Web server listening to TCP port 80 in an app VM. The FlexCapsule server registers a NATP rule that translates a pair of the public IP address used by the app VM and port 80 into a pair of the private IP address of the app VM and port 80. For load balancing, the FlexCapsule server supports process pooling, which is a technique for preparing multiple processes that wait for the same port using `fork`. The FlexCapsule server configures NATP rules so that packets are delivered to one of the app VMs in a process pool.

4 | IMPLEMENTATION

We have implemented FlexCapsule in Xen 4.2.4⁷. A VM called DomU is an instance provided by a cloud and runs user's virtualized system including VMs and the hypervisor using nested virtualization. In the user's virtualized system based on Xen 4.2.4, DomU is used as an app VM and the FlexCapsule server runs in a privileged VM called Dom0. To share information between Dom0 and VMs, the database called XenStore runs in Dom0.

4.1 | Two Types of FlexCapsule OSes

We have implemented two types of FlexCapsule OSes based on OS^v 0.21⁴ and MiniOS in Xen. The reason why we selected these two is that these LibOSes have different characteristics. The system architectures of the two LibOSes are illustrated in Fig. 4. OS^v is an OS optimized for virtualized systems. It is fully virtualized and is called an HVM guest. When using Intel's hardware virtualization extensions, CPUs and memory are virtualized using VT-x and Extended Page Tables (EPT), respectively. To reduce the overhead of device emulation, OS^v provides para-virtual (PV) device drivers to use PV devices. A PV driver running in the OS is called a frontend driver in Xen and communicates with a backend driver running in Dom0. For this communication, these two drivers establish event channels and shared memory at initialization time. OS^v can run many existing C applications with no or slight modification. In addition, It can run custom applications more efficiently.

On the other hand, MiniOS is an OS used for separating components in Dom0 into independent VMs. It is para-virtualized to minimize the virtualization overhead and is called a PV guest. CPUs and memory are virtualized without using hardware extensions. MiniOS directly deals with physical memory called machine memory assigned to a VM. Its page tables are configured to translate virtual addresses into machine memory frame numbers (MFNs). In addition, MiniOS manages pseudo-physical memory as virtualized physical memory in a VM. To translate pseudo-physical memory frame numbers (PFNs) into MFNs, it maintains the P2M table. The reverse translation is performed using the global M2P table. Like OS^v, MiniOS also uses PV

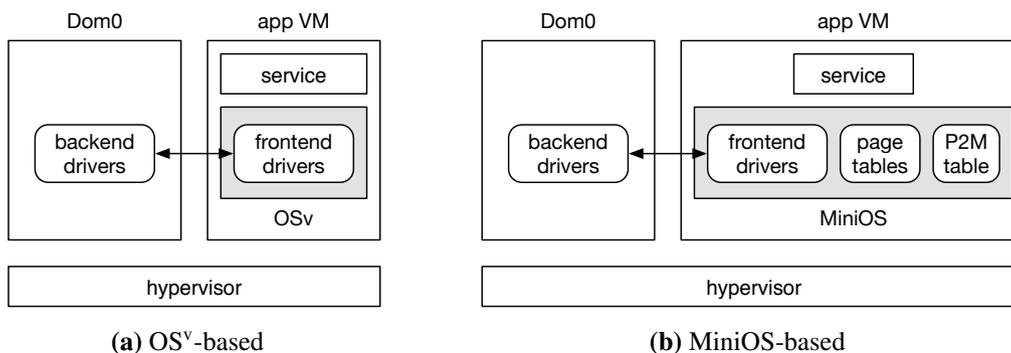


FIGURE 4 Two different FlexCapsule OSes.

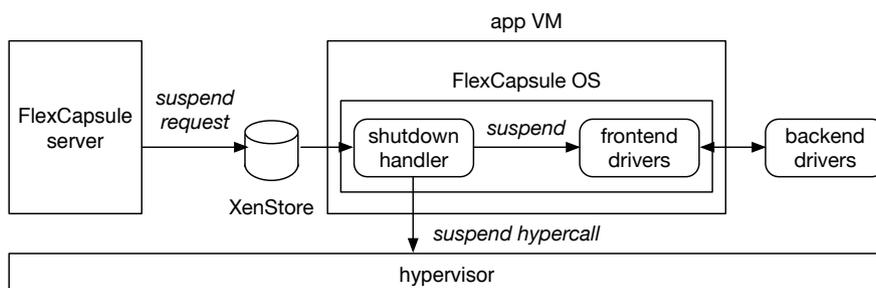


FIGURE 5 Suspending PV devices.

drivers. It provides not only the C runtime but also the OCaml runtime and the GHC runtime for Haskell. However, it cannot run the existing applications easily.

4.2 | Migration of App VMs

To migrate an app VM, the FlexCapsule server first transfers the memory contents of an app VM running in the source instance to the destination instance. Since the memory of the app VM continues to be modified during the memory transfer, the FlexCapsule server re-transfers modified memory contents. If the source and destination instances are in the same host, this memory transfer can be optimized²⁹. It repeats the re-transfers until the amount of memory to be re-transferred is small enough and finally suspends the app VM. At this time, it saves the states of CPUs and virtual devices and transfers them. In the destination instance, it restores these states and resumes the app VM.

Since the FlexCapsule OS uses para-virtualization techniques, it needs migration support at the OS level to migrate an app VM. When suspending an app VM at the final stage of VM migration, the FlexCapsule server writes a request for power management to the control/shutdown node in XenStore, as illustrated in Fig. 5. To monitor the node, the FlexCapsule OS starts a dedicated thread at boot time and registers the shutdown handler as a callback function. When the node value is changed, XenStore sends an event to the corresponding app VM using an event channel and the FlexCapsule OS invokes the registered shutdown handler.

The shutdown handler first suspends PV devices such as block and network devices. Event channels are established between the frontend and backend drivers, but they are not reused after VM migration. Therefore, the shutdown handler destroys all the event channels. Next, it invokes the suspend hypercall and then the hypervisor stops the app VM. In the destination instance, the hypercall returns when the app VM is restarted. Then the shutdown handler resumes PV devices and the frontend drivers re-establish new event channels with the backend drivers in Dom0.

In addition, the MiniOS-based FlexCapsule OS maintains bindings to the hypervisor after VM migration. Since MiniOS is a para-virtualized OS, it is more tightly coupled with the hypervisor than OS^v. For example, MiniOS maintains a grant table for managing memory sharing between VMs. Since shared memory between an app VM and the other VMs is destroyed after VM migration, the FlexCapsule OS removes entries in the grant table on VM suspension and recreates them on VM resumption.

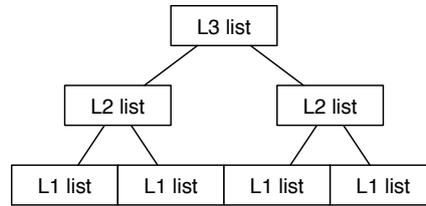


FIGURE 6 The P2M table with a tree structure.

Also, MiniOS manages the P2M table with a three-level tree structure, as shown in Fig. 6. The L1 lists construct a contiguous array and is restored by the FlexCapsule server when a VM is resumed. However, the L2 and L3 lists become invalid in the destination instance because they maintain links using MFNs in the source instance. Therefore the FlexCapsule OS temporarily replaces MFNs in the L2 and L3 lists to virtual addresses on VM suspension. Upon VM resumption, it replaces the virtual addresses to newly assigned MFNs using information in the L1 lists.

When the MiniOS-based FlexCapsule OS issues a hypercall for suspending the app VM, it passes the `start_info` structure to the hypervisor. This data structure contains information on restarting the app VM such as a wallclock and the `shared_info` structure including the state of vCPUs. The OS^v-based FlexCapsule OS does not need to pass this data structure.

In the current implementation, the state of the FlexCapsule server does not need to be transferred together with an app VM. For example, the NAPT rules for the app VM continue to be applied in the source instance, as explained in Section 4.4.

4.3 | Fork of App VMs

When a service invokes the fork function, the FlexCapsule OS sends a fork request to the FlexCapsule server via XenStore. The FlexCapsule server suspends a parent app VM and creates a child app VM. It configures a newly allocated IP address for the child app VM and registers new NAPT rules based on those for the parent app VM. Then it copies the states of the parent app VM to the child app VM. In addition, the FlexCapsule server makes the parent and child app VMs share the disk of the parent app VM in a copy-on-write manner. Finally, it starts to run the child app VM, whereas it resumes the parent app VM by canceling the VM suspension.

Our VM fork is similar to SnowFlock²⁸ and KylinX²⁷, but there are two differences. First, FlexCapsule supports cloning of VMs for not only para-virtualization but also full virtualization. In full virtualization, VM states to be copied are different from those in para-virtualization. Second, FlexCapsule eagerly copies the entire memory of a VM at fork time although SnowFlock and KylinX copy it on demand. This is because the cost for constructing EPT for copy-on-write was similar to that for memory copies in our experiment.

4.3.1 | Duplicating VM States

The FlexCapsule server issues a newly created hypercall to duplicate VM states. First, the hypercall copies the memory contents of the parent app VM to the child app VM. For each page of the parent app VM, it allocates a new page for the child app VM and registers the mapping from its MFN to the corresponding PFN. For PV guests, the FlexCapsule server needs special cares for the page tables and the P2M table, as illustrated in Fig. 7. When the hypercall copies pages used for page table entries (PTEs), it replaces MFNs stored in them. It first canonicalizes an MFN for the parent app VM into the corresponding PFN using the M2P table and then uncanonicalizes the PFN into an MFN used for the child app VM using the P2M table. After all the PTEs are modified, the hypercall sets the page type to the L1 to L4 page table.

For the P2M table, the hypercall replaces MFNs similarly. Since the P2M table has a tree structure as in Fig. 6, the hypercall first finds the L1 lists by traversing the tree from the root. Then it replaces MFNs in the L1 lists using the M2P table. The M2P table is constructed by the hypervisor when the memory of the child app VM is allocated. The L2 and L3 lists also contain MFNs for the parent app VM, but they are reconstructed by the MiniOS-based FlexCapsule OS on resuming an app VM, as described in Section 4.2.

Next, the hypercall copies the states of the parent app VM to the child app VM. For HVM guests, it copies CPU states such as PAE, the TSS in the virtual 8086 mode, the identity-map page directory, and the location of ACPI control blocks. Also, it copies information on the ring buffers used for memory events, console, XenStore, and I/O requests to the child app VM. Then

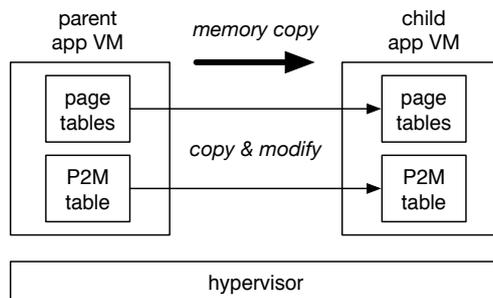


FIGURE 7 Memory duplication for a PV guest.

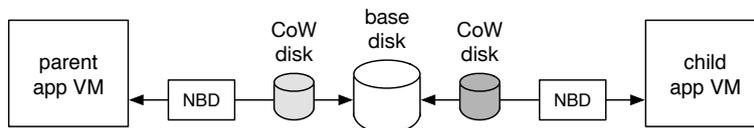


FIGURE 8 Dynamic copy-on-write sharing of a disk.

it saves the HVM context in the parent app VM and loads it to the child app VM. For PV guests, the hypercall copies the vCPU context of the parent app VM to the child app VM. At that time, it replaces the MFNs used for the `start_info`, GDT frames, and the page table base pointer for the child app VM. In addition, it copies the `shared_info` structure. For both HVM and PV guests, the hypercall copies the time stamp counter.

Finally, the FlexCapsule server duplicates device states only for HVM guests. It saves the device states in `qemu-dm` for the parent app VM by sending a QMP command. Xen-specific QEMU called `qemu-dm` is a device emulator and runs in `Dom0`. Then the FlexCapsule server starts `qemu-dm` for the child app VM with the saved states.

4.3.2 | Sharing a Disk

To share the disk of the parent app VM with the child app VM in a copy-on-write manner, the FlexCapsule server creates two copy-on-write (CoW) disks for these app VMs, respectively, as illustrated in Fig. 8. A CoW disk is a disk that stores only updates to the base disk, which is a disk used by the parent app VM before fork. In the `qcow2` format, writes are performed to a CoW disk, whereas reads are performed from the disk if disk blocks exist; otherwise, they are done from the base disk.

To enable dynamically changing a disk of a running app VM, the FlexCapsule server indirectly attaches a disk to an app VM via a network block device (NBD)³⁰. For the child app VM, the FlexCapsule server connects one CoW disk to an unused NBD device and then attaches the device to the child app VM. For the parent app VM, the FlexCapsule server first disconnects the base disk from the NBD device already attached. Then it re-connects the other CoW disk to the original NBD device. As such, the parent app VM can use the CoW disk seamlessly. If the NBD is not used for indirect disk attachment, this is difficult to achieve because a local disk directly attached to a VM cannot be detached as long as the VM is not stopped.

4.4 | Networking

Fig. 9 illustrates networking in FlexCapsule. When a service invokes the `listen` function to wait for new network connections, the FlexCapsule OS obtains a listening port number from the specified socket and sends it to the FlexCapsule server via XenStore. Then the FlexCapsule server adds a new NAPT rule to Linux iptables. The added rule forwards packets sent to the public IP address and port number used by the app VM to the app VM. We used the `libiptc` library³¹ for manipulating netfilter, which is a packet filtering framework in Linux. When a service invokes the `close` function for a socket, the FlexCapsule OS sends a listening port obtained from the socket to the FlexCapsule server via XenStore. Then the FlexCapsule server deletes the NAPT rule corresponding to the port.

To achieve load balancing with a process pool for app VMs, the FlexCapsule server uses the `nth` mode of the `statistic` module for iptables. The `nth` mode is used for simple stateful load balancing in a round-robin fashion. When a service invokes the `fork`

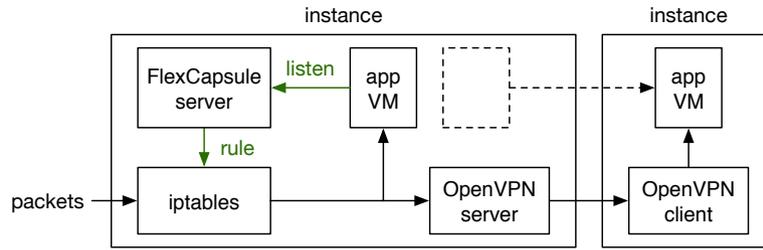


FIGURE 9 Networking in FlexCapsule.

function or the FlexCapsule server clones an app VM for scale-out, the FlexCapsule server examines whether the app VM is listening to network ports. If there are such ports, the FlexCapsule server translates the corresponding NAT rules into rules for the nth mode and adds new rules for a child app VM. Using the nth mode, packets are delivered to one of the app VMs included in the same process pool. Note that all the packets in one connection are delivered to the same app VM.

The FlexCapsule server constructs one VPN that connects all app VMs inside user's instances using Ethernet bridging of OpenVPN 2.3.2³². One instance runs the OpenVPN server, whereas the others run the OpenVPN clients. Thanks to the VPN, even after app VMs are migrated to other instances, the NAT rules in the source instance are still applied. If the source instance has no app VM and is therefore stopped, the FlexCapsule server transfers the NAT rules to the destination instance.

5 | EXPERIMENTS

We conducted experiments to confirm the effectiveness of FlexCapsule. In our experiments, we used four app VMs running the OS^v-based FlexCapsule OS (OS^v VMs). For each app VM, we ran one of the four services: `lighttpd` 1.4.35³³, `memcached` 1.4.21³⁴, `Redis` 3.0.1³⁵, and `Dhrystone` 2.1³⁶. `lighttpd` is a single-threaded web server, `memcached` is a memory object caching system, `Redis` is an in-memory database, and `Dhrystone` is a computing benchmark. For `memcached`, we used a version customized for OS^v. For `Redis`, we used pipelining for sending new requests without waiting for responses. We measured the server throughput using `httperf` 0.9.0³⁷ for `lighttpd`, `memslap`³⁸ for `memcached`, and `redis-benchmark`³⁹ for `Redis`. For an app VM running the MiniOS-based FlexCapsule OS (MiniOS VM), we ran only `Dhrystone` because it was difficult to port the other existing services.

We used two PCs with an Intel Xeon E3-1290 v2 processor and 8 GB of memory. We ran Xen 4.2.4 for the hypervisor and Linux 3.13.0 in Dom0. We created several DomUs as instances in a cloud and assigned four vCPUs and 2 GB of memory for each. Inside these instances, we ran several app VMs, each of which was assigned one vCPU and 4 to 1024 MB of memory.

5.1 | Service Consolidation

To show the effectiveness of service consolidation using app VMs, we measured changes in service performance after consolidation and de-consolidation. In this experiment, we used three OS^v VMs running `lighttpd`, `memcached`, and `Redis`, respectively. When consolidating these app VMs, we ran them in one instance with one vCPU. When de-consolidating them, we used three instances, each of which has one vCPU.

Fig. 10 shows the performance after service consolidation and de-consolidation. In all the services, the performance after de-consolidation was 1.9 to 2.7 times higher than that after consolidation. This means that service de-consolidation can improve the service performance even when services run in app VMs. In other words, service consolidation is useful if each app VM does not use so large amount of resources. When only one of the three app VMs was busy in one instance, it could gain the performance similar to that after de-consolidation.

5.2 | Scale-out

First, we investigated whether app VM-level scale-out was effective. We increased the number of app VMs running in one instance and measured the total performance of all the app VMs. For OS^v VMs, we ran `lighttpd` and measured the throughput when we sent requests for 1 KB files. Fig. 11(a) shows the total throughput of app VMs. When we increased the number of app

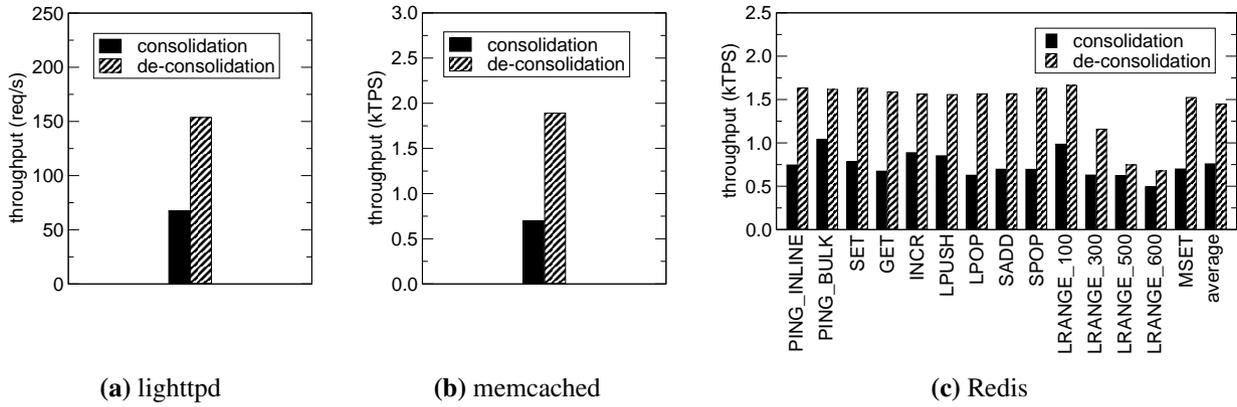


FIGURE 10 The performance after service consolidation and de-consolidation.

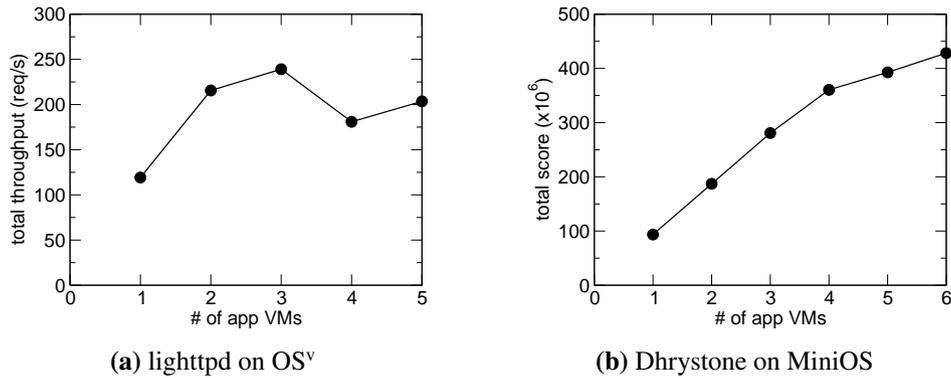


FIGURE 11 The performance improvement by app VM-level scale-out.

VMs, the total throughput also increased until three app VMs. It was doubled by using two app VMs, but the increase was small in three app VMs. When we consolidated four app VMs, the throughput rather decreased. This is because three app VMs and one Dom0 inside the instance used up four vCPUs. The throughput slightly increased in five app VMs by using resources more efficiently.

For MiniOS VMs, we ran Dhrystone and measured its score. Like lighttpd, Dhrystone uses a single thread. Fig. 11(b) shows the total score, which linearly increased to four app VMs because executing Dhrystone did not need Dom0. The total score continued to increase for more app VMs although the increase became smaller. This means that one app VM running Dhrystone could not use up one vCPU.

Second, we investigated whether instance-level scale-out was effective. We increased the number of instances and measured the throughput of lighttpd. When we ran one OS^v VM per instance, the total throughput in two instances became twice. Even when we fixed the number of app VMs to four in total, using two instances improved the total throughput. This is because each instance could run two app VMs and one Dom0 using four vCPUs.

5.3 | Scale-up

We investigated whether instance-level scale-up improved the performance of an app VM. We ran one OS^v VM running lighttpd in one instance and changed the number of vCPUs assigned to the instance. Fig. 12(a) shows the throughput of lighttpd when we sent requests for 1 MB files. As the number of vCPUs was increasing, the throughput was improved. For two vCPUs, the app VM and Dom0 could use one vCPUs for each. The performance was still improved in four vCPUs although the app VM was assigned only one vCPU. This is because Dom0 used more than one vCPUs.

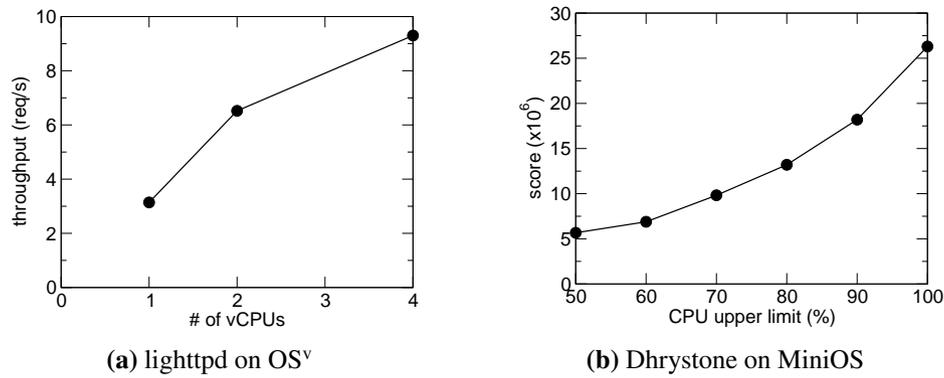


FIGURE 12 The performance improvement by scale-up.

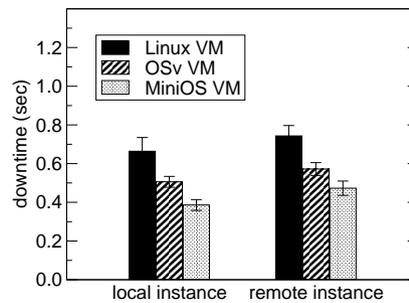


FIGURE 13 The downtime during VM migration between instances.

For a MiniOS VM, we ran Dhrystone and changed the CPU upper limit of an instance. Fig. 12(b) shows that the score was improved as the limit was relaxed. However, the performance was not proportional to the assigned CPU time. This is probably due to the vCPU scheduler in the hypervisor.

5.4 | Downtime

We measured the downtime during the migration of an app VM between instances. In this experiment, the downtime is the time until an app VM is restarted in the destination instance after the suspend request is sent to the app VM in the source instance. We migrated an OS^v VM and a MiniOS VM. For comparison, we ran a regular VM running Linux with PV drivers (Linux VM) in an instance and migrated it. We measured the downtime for the VMs with various memory sizes when we migrated a VM to another instance at the same host (local instance) and at the different host (remote instance).

Fig. 13 shows the average and standard deviation of the downtime. The downtime did not depend on the memory size of a VM and was almost the same. Overall, the downtime of the app VMs was sufficiently short. The downtime of the OS^v VM was shorter than that of the Linux VM. One reason is that OS^v supports only the smaller number of virtual devices to be suspended. Another is that suspending the state of PV devices was faster in OS^v. Compared with the OS^v VM, the downtime of the MiniOS VM was shorter. This means that complete para-virtualization makes the suspension of a VM faster. When VMs were migrated to a remote instance, the downtime became only 0.1 seconds longer.

5.5 | Migration Time

We measured the migration time between instances when we changed the memory sizes of app VMs. The migration time is the time needed for the execution of the migration command. Like the experiment in Section 5.4, we migrated the two app VMs and the Linux VM. Fig. 14(a) shows the results when we migrated a VM to another instance at the same host. The migration time was proportional to the memory size of a VM. The app VMs had an advantage over the Linux VM because they could run with the smaller amount of memory. The MiniOS VM needed only 4 MB at minimum, whereas the Linux VM needed 128 MB

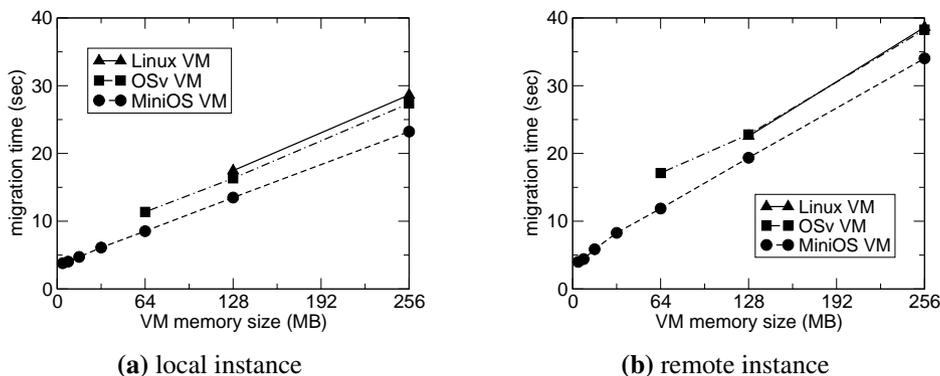


FIGURE 14 The migration time between instances.

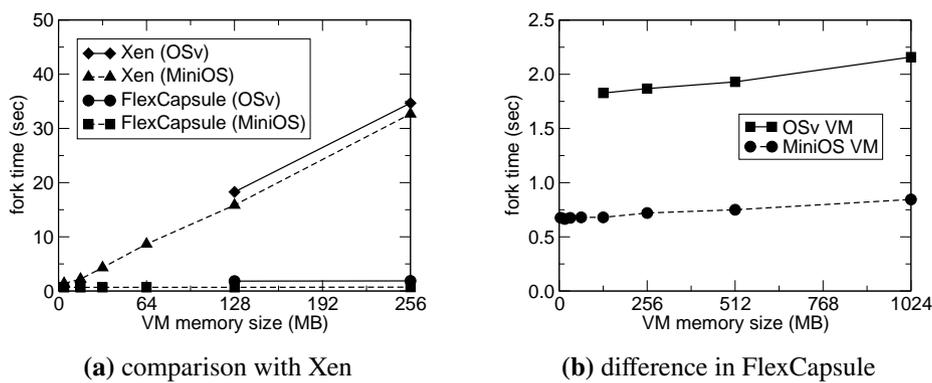


FIGURE 15 The fork time of app VMs.

at least. In this case, the app VM could be migrated 4.6 times faster than the Linux VM. On the other hand, the OS^v VM needed 64 MB. In this minimal memory size, the migration time was still 1.5 times faster than the Linux VM. However, it was 3 times slower than that of the MiniOS VM.

Fig. 14(b) show the migration time when we migrated VMs to another instance at a different host. Due to network overhead, the increase in migration time was more rapid. As a result, the migration of the MiniOS VM was 5.6 times faster than that of the Linux VM when these VMs were assigned the minimum amount of memory.

5.6 | Fork Time

We measured the time needed for the execution of the fork function in an app VM. The fork function in the FlexCapsule OS communicates with the FlexCapsule server and performs VM fork. For comparison, we measured the fork time when we naively achieved VM fork using Xen's standard save and restore commands. Fig. 15(a) shows the comparison of the fork time in an OS^v VM and a MiniOS VM. Our VM fork was much faster than that using Xen's commands. With VM fork using Xen's commands, the fork time increased significantly as the memory size of an app VM was increasing. For app VMs with 256 MB of memory, our VM fork was 26 and 45 times faster for OS^v and MiniOS, respectively. Compared with the fork function in Linux, our VM fork needed much more time. In a Linux VM, it took only 5.4 ms because process fork is much faster than VM fork.

Fig. 15(b) compares the fork time in our VM fork in detail. The fork time did almost not depend on the memory size of an app VM. This means that even eager memory copies from a parent VM to a child app VM was very efficient. The fork of the OS^v VM was much slower than that of the MiniOS VM. One reason is that it took a long time to restore the state of qemu-dm in OS^v. Another is the difference of the overhead of nested virtualization: 0.8 seconds for OS^v and 0.2 seconds for MiniOS. This is due to the difference between HVM and PV guests and that of the implementation complexity.

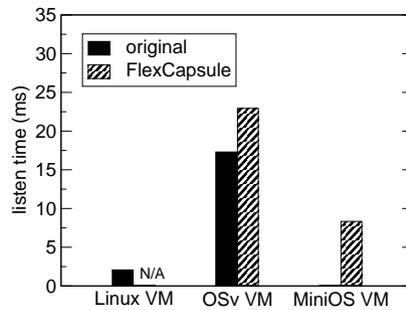


FIGURE 16 The listen time.

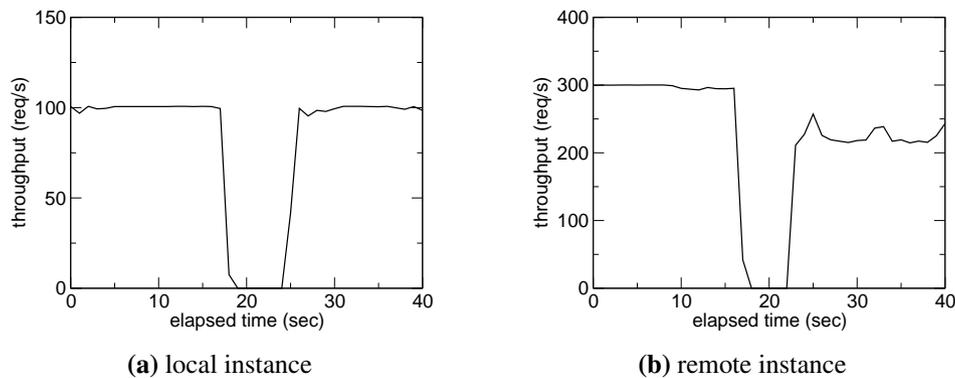


FIGURE 17 The performance after VM migration.

5.7 | Listen Time

We measured the time needed for the execution of the listen function in an app VM. The listen function in the FlexCapsule OS communicates with the FlexCapsule server and registers a new NAPT rule. Fig. 16 shows the listen time in an OS^v VM, a MiniOS VM, and a Linux VM. The overhead of the registration was 5.6 and 8.3 ms in OS^v and MiniOS, respectively. This is not large for OS^v but is too large for MiniOS because the execution time of the listen function was very small in MiniOS. However, this does not degrade application performance because the listen function is not invoked so frequently.

5.8 | Overhead of the VPN

To examine the overhead of the VPN across instances, we migrated an OS^v VM to another instance and measured the throughput change of lighttpd in the app VM. After we migrated the app VM, packets were forwarded to the destination instance by the VPN server. Fig. 17 shows the changes in throughput. When we migrated the app VM to the instance at the same host, the throughput was degraded only by 1%. In contrast, the performance degradation was 24% when we migrated the app VM to a remote instance. This is the overhead of packet forwarding between hosts.

5.9 | Service Performance

We ran various services in app VMs and compared its performance with that in a Linux VM. We ran Dhrystone in both an OS^v VM and a MiniOS VM, whereas we ran the other services only in OS^v VMs. For memcached, we also used an unoptimized version, which uses the general-purpose socket API, for example. To examine the maximum performance, we assigned four vCPUs to each app VM. Fig. 18 shows the performance. For Dhrystone, the performance of the MiniOS VM was the best due to less overhead and was 31% higher than that of the Linux VM. That of the OS^v VM was slightly worse than that of the Linux VM. For lighttpd, the throughput of the app VM was 9% higher than that of the Linux VM. The performance degradation due to using NAPT was 3%. For memcached, the performance of the app VM was 4.3 times higher when we used a version customized

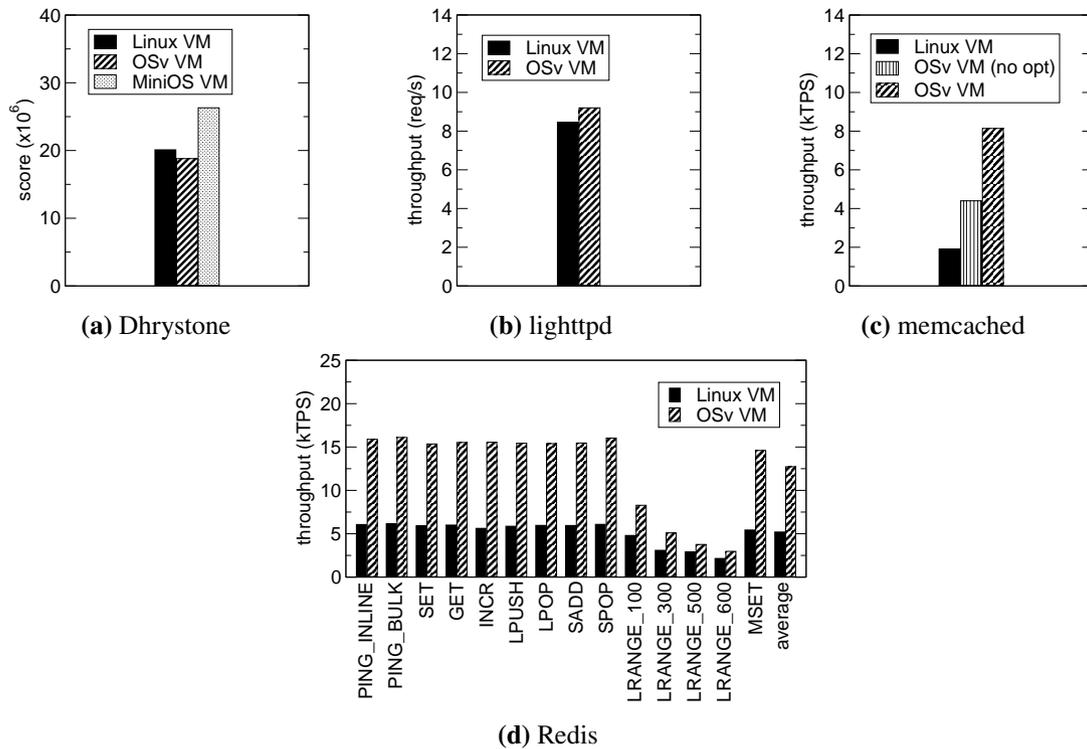


FIGURE 18 The performance of various services in app VMs.

for OS^v. Even using the original version, the performance was 2.3 times higher. For Redis, the performance of the app VM was 2.5 times higher on average.

6 | CONCLUSION

This paper proposed FlexCapsule, which runs each service in a lightweight VM called an app VM using a LibOS. FlexCapsule can optimize instance deployment at service granularity. The migration of app VMs enables seamless service consolidation and de-consolidation and scale-up and -down. Strong isolation among app VMs enables secure service consolidation. We have implemented FlexCapsule in Xen using OS^v and MiniOS. The FlexCapsule OS cooperates with the FlexCapsule server to support VM migration, networking, process fork, and process pools. Experimental results show that FlexCapsule is effective for the optimization of instance deployment.

One of our future work is to enable various services to run in app VMs. For example, we need to advance multi-process support such as inter-process communication. Another direction is performance improvement of app VMs. Since OS^v running in app VMs is fully virtualized except for PV drivers, virtualization overhead is relatively large in nested virtualization. To reduce that overhead, we need to develop para-virtualized OS^v. The network performance of app VMs should be improved because the emulation of NICs used by user's virtualized system is heavyweight. This overhead can be mitigated by using network optimization in Xen-Blanket¹⁴.

References

1. Milošević DS, Douglass F, Paindaveine Y, Wheeler R, Zhou S. Process Migration. *ACM Comput Surv* 2000; 32(3): 241–299.
2. Engler DR, Kaashoek MF, J. O' Toole J. Exokernel: An Operating System Architecture for Application-level Resource Management. In: *Proceedings of the 15th ACM Symposium on Operating Systems Principles.*; 1995: 251–266.

3. Madhavapeddy A, Mortier R, Rotsos C, et al. Unikernels: Library Operating Systems for the Cloud. In: Proceedings of the 18th International Conference on Architectural Support for Programming Languages and Operating Systems. ; 2013: 461–472.
4. Kivity A, Laor D, Costa G, et al. OSv – Optimizing the Operating System for Virtual Machines. In: Proceedings of the 2014 USENIX Annual Technical Conference. ; 2014: 61–72.
5. Goldberg R. Architecture of Virtual Machines. In: Proceedings of Workshop on Virtual Computer Systems. ; 1973: 74–112.
6. Ben-Yehuda M, Day MD, Dubitzky Z, et al. The Turtles Project: Design and Implementation of Nested Virtualization. In: Proceedings of the 9th Symposium on Operating Systems Design and Implementation. ; 2010: 423–436.
7. Barham P, Dragovic B, Fraser K, et al. Xen and the Art of Virtualization. In: Proceedings of the 19th ACM Symposium on Operating Systems Principles. ; 2003: 164–177.
8. Kourai K, Sannomiya K. Seamless and Secure Application Consolidation for Optimizing Instance Deployment in Clouds. In: Proceedings of the 8th IEEE International Conference on Cloud Computing Technology and Science. ; 2016: 318–325.
9. VMware, Inc. . VMware vCloud Air – Public Cloud Computing from VMware. <http://vcloud.vmware.com>; Accessed September 7, 2018.
10. Ben-Yehuda OA, Ben-Yehuda M, Schuster A, Tsafirir D. The Resource-as-a-service (RaaS) Cloud. In: Proceedings of the 4th USENIX Workshop on Hot Topics in Cloud Computing. ; 2012.
11. Osman S, Subhraveti D, Su G, Nieh J. The Design and Implementation of Zap: A System for Migrating Computing Environments. In: Proceedings of the 5th Symposium on Operating Systems Design and Implementation. ; 2002: 361–367.
12. Zhang L, Litton J, Cangialosi F, Benson T, Levin D, Mislove A. Picocenter: Supporting Long-lived, Mostly-idle Applications in Cloud Environments. In: Proceedings of the 11th European Conference on Computer Systems. ; 2016.
13. Williams D, Elnikety E, Eldehry M, Jamjoom H, Huang H, Weatherspoon H. Unshackle the Cloud!. In: Proceedings of the 3rd USENIX Workshop on Hot Topics in Cloud Computing. ; 2011.
14. Williams D, Jamjoom H, Weatherspoon H. The Xen-Blanket: Virtualize Once, Run Everywhere. In: Proceedings of the 7th ACM European Conference on Computer Systems. ; 2012: 113–126.
15. Fishman A, Rapoport M, Budilovsky E, Eidus I. HVX: Virtualizing the Cloud. In: Proceedings of the 5th USENIX Workshop on Hot Topics in Cloud Computing. ; 2013.
16. Liu C, Mao Y. Inception: Towards a Nested Cloud Architecture. In: Proceedings of the 5th USENIX Workshop on Hot Topics in Cloud Computing. ; 2013.
17. Zhang F, Chen J, Chen H, Zang B. CloudVisor: Retrofitting Protection of Virtual Machines in Multi-tenant Cloud with Nested Virtualization. In: Proceedings of the 23rd ACM Symposium on Operating Systems Principles. ; 2011: 203–216.
18. Tan C, Xia Y, Chen H, Zang B. TinyChecker: Transparent Protection of VMs against Hypervisor Failures with Nested Virtualization. In: Proceedings of IEEE/IFIP International Workshop on Dependability of Clouds, Data Centers and Virtual Machine Technology. ; 2012.
19. Intel Corp. . 4th Generation Intel Core vPro Processors with Intel VMCS Shadowing. 2013.
20. ARM Ltd. . *ARM Architecture Reference Manual – ARMv8, for ARMv8-A Architecture Profile*. 2017.
21. Lim JT, Dall C, Li S, Nieh J, Zyngier M. NEVE: Nested Virtualization Extensions for ARM. In: Proceedings of the 26th ACM Symposium on Operating Systems Principles. ; 2017: 201–217.
22. Ammons G, Silva DD, Krieger O, et al. Libra: A Library Operating System for a JVM in a Virtualized Execution Environment. In: Proceedings of the 3rd International Conference on Virtual Execution Environments. ; 2007: 44–54.
23. Jordan M, Roeck H. Guest VM Microkernel, GUK. <https://kenai.com/projects/guestvm>; Accessed March 8, 2017.

24. Kaashoek MF, Engler DR, Ganger GR, et al. Application Performance and Flexibility on Exokernel Systems. In: Proceedings of the 16th ACM Symposium on Operating Systems Principles. ; 1997: 52–65.
25. Tsai CC, Arora KS, Bandi N, et al. Cooperation and Security Isolation of Library OSES for Multi-process Applications. In: Proceedings of the 9th European Conference on Computer Systems. ; 2014.
26. Douceur J, Elson J, Howell J, Lorch J. Leveraging Legacy Code to Deploy Desktop Applications on the Web. In: Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation. ; 2008: 339–354.
27. Zhang Y, Crowcroft J, Li D, et al. KylinX: A Dynamic Library Operating System for Simplified and Efficient Cloud Virtualization. In: Proceedings of the 2018 USENIX Annual Technical Conference. ; 2018: 173–185.
28. Lagar-Cavilla HA, Whitney JA, Scannell AM, et al. SnowFlock: Rapid Virtual Machine Cloning for Cloud Computing. In: Proceedings of the 4th ACM European Conference on Computer Systems. ; 2009: 1–12.
29. Kourai K, Ooba H. VMBeam: Zero-copy Migration of Virtual Machines for Virtual IaaS Clouds. In: Proceedings of the 35th IEEE Symposium on Reliable Distributed Systems. ; 2016: 121–126.
30. Verhelst W. Network Block Device. <https://nbd.sourceforge.io/>; Accessed September 7, 2018.
31. Balliache L. Querying libiptc HOWTO. <http://www.tldp.org/HOWTO/Querying-libiptc-HOWTO/>; Accessed September 7, 2018.
32. OpenVPN Technologies, Inc. . OpenVPN – Open Source VPN. <https://openvpn.net/>; Accessed September 7, 2018.
33. Lighty Team . Lighttpd - fly light. <https://www.lighttpd.net/>; Accessed September 7, 2018.
34. Fitzpatrick B. memcached – A Distributed Memory Object Caching System. <http://memcached.org/>; Accessed September 7, 2018.
35. Redis Labs, Inc. . Redis. <https://redis.io/>; Accessed September 7, 2018.
36. Weicker R. Dhrystone: A Synthetic Systems Programming Benchmark. *Communications of the ACM* 1984; 27(10): 1013–1030.
37. Mosberger D, Jin T. httpperf – A Tool for Measuring Web Server Performance. tech. rep., Hewlett-Packard Company; 1998.
38. Aker B. memaslap – Load Testing and Benchmarking a Server. <http://docs.libmemcached.org/bin/memaslap.html>; Accessed September 7, 2018.
39. Redis Labs, Inc. . How fast is Redis?. <https://redis.io/topics/benchmarks>; Accessed September 7, 2018.

