

# Exploring Live Cloud Migration On Amazon EC2

Ibrahim Ejdayid A.Mansour  
Dept. of Computing & Informatics  
Bournemouth University  
Bournemouth, UK  
imansour@bournemouth.ac.uk

Hamid Bouchachia  
Dept. of Computing & Informatics  
Bournemouth University  
Bournemouth, UK  
abouchachia@bournemouth.ac.uk

Kendra Cooper  
Independent Scholar  
Kelowna, Canada  
kendra.m.cooper@gmail.com

**Abstract**—Cloud users may decide to live migrate their virtual machines from a public cloud provider to another due to a lower cost or ceasing operations. Currently, it is not possible to install a second virtualization platform on public cloud infrastructure (IaaS) because nested virtualization and hardware-assisted virtualization are disabled by default. As a result, cloud users' VMs are tightly coupled to providers IaaS hindering live migration of VMs to different providers. This paper introduces LivCloud, a solution to live cloud migration. LivCloud is designed based on well-established criteria to live migrate VMs across various cloud IaaS with minimal interruption to the services hosted on these VMs. The paper discusses the basic design of LivCloud which consists of a Virtual Machine manager and IPsec VPN tunnel introduced for the first time within this environment. It is also the first time that the migrated VM architecture (64-bit & 32-bit) is taken into consideration. In this study, we evaluate the implementation of the basic design of LivCloud on Amazon EC2 C4 instance. This instance has a compute optimized instance and has high performance processors. In particular we explore three developed options. These options are being tested for the first time on EC2 to change the value of the EC2 instance's control registers. Changing the values of the registers will significantly help enable nested virtualization on Amazon EC2.

**Keywords**—Virtualization; Virtual Machine; Network Virtualization; Nested Virtualization; Live Cloud Migration; Cloud infrastructure

## I. INTRODUCTION

Virtualization allows cloud users to exploit multi-tenant resources (compute, network and storage) from a secure Cloud IaaS [3]. Public cloud IaaS is often described as a heterogeneous environment due to the fact that each cloud provider has their own hypervisor. Providers such as Amazon EC2 and Rackspace use the hypervisor Xen; while, Fractus and Google Compute Engine rely on KVM. Windows Azure, on the other hand, uses the Microsoft hypervisor, Hyper-V [1][2]. Despite that many providers leverage the same hypervisors for virtualization, for example, Google and HP both use KVM, live cloud migration of VMs between those providers is still challenging [2]. This has made it difficult for cloud users to live migrate VMs to other providers [1]. This is one consequence of vendor lock-in [6]. If cloud users were able to utilize resources and services from various providers, many benefits would be achieved, including, offering service continuity in case of ceasing operation or natural disasters, borrowing resources from different providers in case of over-utilization or limited resources and achieving high flexibility

to change service providers, thereby, alleviating vendor lock-in [6][21]. Furthermore, in 2016, RightScale conducted cloud computing trends survey in which 1060 IT professionals were interviewed about their adoption of cloud infrastructure and related technologies. The survey showed 17% of enterprises had more than 1000 virtual machines (VMs) in public cloud, up from 13% in 2015 [20]. This number of VMs would have been reduced to 250 VMs hosting 4 VMs each if public cloud IaaS had not been deliberately locked (disabled nested virtualization or no hardware-assisted virtualization features enabled).

**LivCloud** is a user-centric live cloud migration approach that does not need the provider's agreement to be applied. It is designed based on live cloud migration criteria to achieve effective live migration of VMs public cloud infrastructure (IaaS) without service interruption [2][21]. It is envisioned into two stages: the basic design and enhancement of the basic design [2]. This paper evaluates the basic design of LivCloud on Amazon EC2 c4.2xlarge instances [2]. By default nested virtualization or hardware-assisted virtualization features (Intel VT-x, Intel VT-d and Extended Page Tables) are *not enabled* on any Amazon instances [4][28]. Nested virtualization is configuring one hypervisor (in the upper layer) within a virtual machine hosted by another hypervisor [25]. Consequently, enhanced QEMU, HQEMU is configured as a second layer hypervisor. HQEMU [29] is an academic project to enhance QEMU performance by using dynamic binary translation (DBT). DBT is similar to binary translation mentioned in Section II-B, but DBT is an open source technology. The implementation process has a number of twisted configurations to overcome Amazon network and KVM configuration challenges. For example, adding a second network interface with Elastic IP [27] is layer 3 networking with detailed steps to correctly enable this interface; whereas, in traditional operating system, adding a second interface is a simple layer 2 networking. Moreover, configuring IPsec VPN tunnel between Amazon VPC and the local network to secure the migration channel. The Virtual Machine (VM) manager is used as GUI interface to connect *Cloud-Host* on the Amazon VPC to *Local-Host* on the local network. IPsec VPN and the virtual manager are the main contributions of implementing LivCloud.

The rest of the paper is organized as follows: Section II presents a brief summary of related work highlighting existing approaches to achieve live cloud migration. Section III

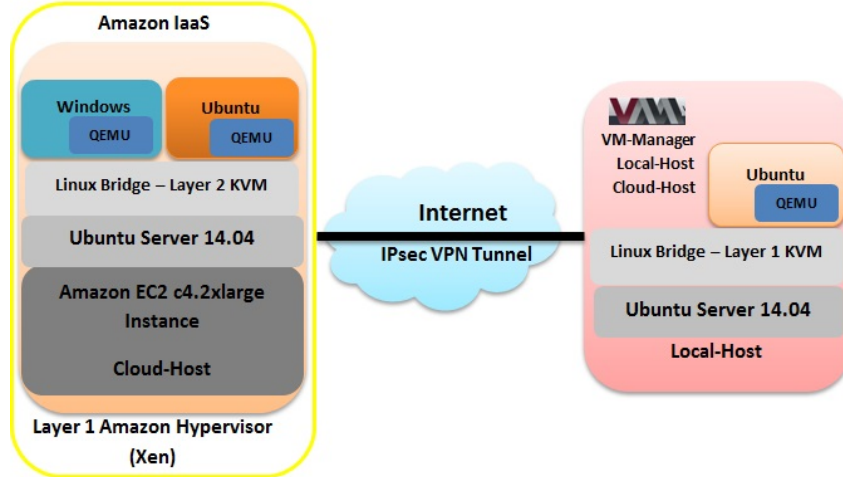


Fig. 1. LivCloud Architecture on Amazon EC2

introduces the LivCloud architecture on Amazon EC2. Section IV explains the implementation of LivCloud on EC2. Section V discusses the implementation results. Section VI outlines the possible solutions to enable nested virtualization on Amazon EC2. The conclusion and future work are presented in Section VII.

## II. RELATED WORK

Nested virtualization has been used to decouple the VM from public IaaS [1][3][7]. Nested virtualization is configuring one hypervisor (in the upper layer) within a virtual machine hosted by another hypervisor [25]. Most of legacy hypervisors, such as KVM, Xen, and VMware can run nested virtualization [3][4][7]. However, public cloud hypervisors do not allow running nested virtualization [1]. Two main techniques have been used to enable nested virtualization on the top of cloud IaaS: paravirtualization and binary translation. The Xen hypervisor can be configured to run paravirtualization concept, while VMware and hypervisor, HVX run binary translation [4][7]. KVM is limited in running paravirtualization. However, OPENFV has been developing KVM for running Network Function Virtualization (NFV), which will help overcoming KVMs limitations [2][9]. In the following sections, two related user-centric approaches are explained, Supercloud [1] that uses paravirtualization and HVX [7] that runs binary translation.

### A. Supercloud

It is developed using resources from a number of major cloud providers, including Amazon EC2, Rackspace, HP Cloud and other private clouds [1]. Supercloud uses nested virtualization (Xen-Blanket [4]) that overcomes cloud heterogeneity. Xen-Blanket leverages the paravirtualization (PV-on-HVM) drivers on Xen. Xen paravirtualization cannot run unmodified operating systems (i.e., Windows) [4]. The approach achieves relatively acceptable performance, about 1.4 seconds migration downtime [4][22]. Disk I/O drivers overhead caused by Xen-Blanket reached 30%, which may affect the physical

machine and the other VMs residing on that machine [1][23]. The approach does not utilize an encryption algorithm. Also, a security mechanism is not used during the process, so it opens the system to security attacks. As a result, the migration channel is insecure and data flow is vulnerable to attacks, such as, ARP/ DHCP/DNS poisoning [26]. The approach relies on Xen as a nested virtualization platform, which has a number of issues. Xensploit tool is developed to execute man-in-the-middle attack during VM migration. It is able to modify the sshd memory segment to circumvent sshd authentication. With such a tool, VM might be accessed and the system confidentiality and integrity may be compromised [13][14].

### B. HVX

It can run unmodified operating systems (i.e., Windows). HVX is similar to VMware because both virtualization platforms use binary translation. However, the lack of a popular open-source binary translation hypervisor has allowed other approaches (such as paravirtualization) to be more popular [7][15]. The binary translation manages to incorporate various virtualization hypervisors, such as Qemu, Xen paravirtualization, KVM and VMware ESX. Therefore, it is able to decouple the VM from underlying hardware [7]. This approach is the only one to run both modified O/S (Linux) and unmodified O/S (Windows). Despite, it is seen as a proprietary product and it cannot be evaluated [15]. There is no quantitative evaluation of the approach's speed, but rather it is mentioned as robust and reliable [7]. As the approach leverages binary translation to achieve a better performance in a nested virtualization environment, many experts do not agree with performance statement as this technique imposes extra overhead on the guest kernel [1][7].

## III. LIVCLOUD ARCHITECTURE ON AMAZON EC2

Figure 1 illustrates LivCloud architecture on Amazon EC2. LivCloud is designed based on live cloud migration criteria

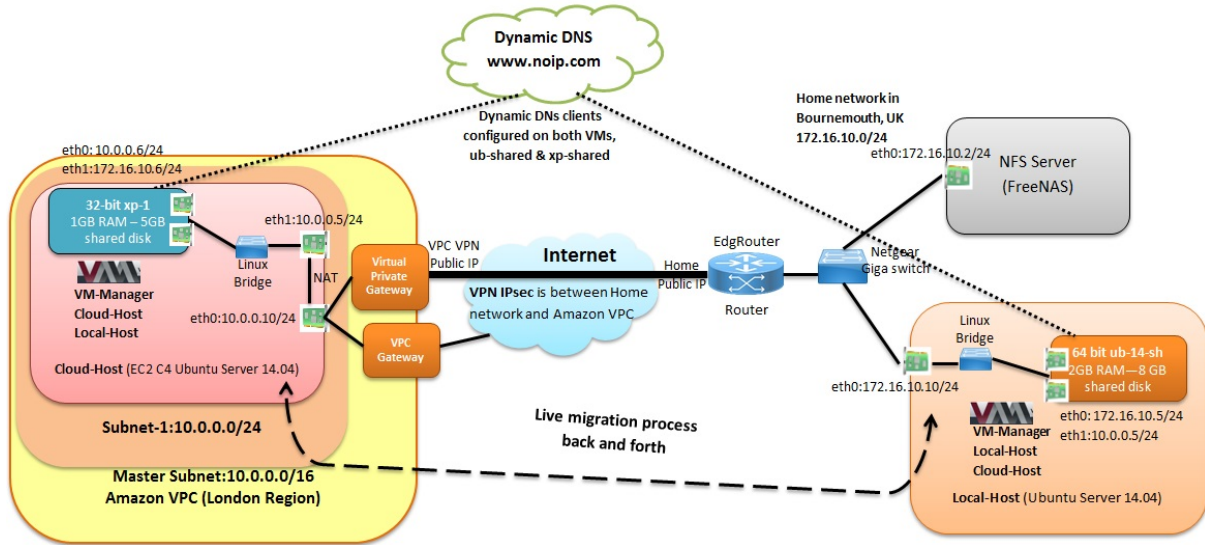


Fig. 2. LivCloud's implementation on Amazon EC2

published in [21], which are classified in three general categories: performance, flexibility and security. There are three performance criteria, denoted as:

- P1 live migration must be imperceptible to the migrated VM and its users;
- P2 predicting the required resources to decide whether or not to proceed with live migration;
- P3 monitoring resource utilization to avoid overutilization and to predict any possible failure.

There are two flexibility criteria:

- F1 decoupling the migrated VM from underlying system by supporting wide range of hardware drivers, such as CPU drivers;
- F2 supporting various OS on the migrated VM, for instance, Windows.

With respect to security, there are security criteria:

- S1 maintaining data privacy during live migration using encryption;
- S2 imposing authentication during migration .

To support effective live cloud migration, the design needs a foundation that supports nested virtualization in order to decouple VMs from the cloud IaaS and connect hypervisors on the IaaS in order to facilitate live migration back and forth. In addition to this, the design needs to optimize live migration performance, prevent any potential failure, and protect the process against hijacking and penetration [2]. The basic requirements help fulfill F1, F2, and P1. In the basic design stage, *Dynamic DNS* is used to maintain the migrated VM's connections and configurations (P1). Dynamic DNS is used to keep a domain name pointing to the same physical or virtual server connected to the Internet regardless of any IP addresses changes [30]. Also, IPsec VPN tunnel is used to fulfill S1 and S2. The secure connection between LivCloud and IaaS is an essential part of live cloud migration.

## IV. LIVCLOUD IMPLEMENTATION ON EC2

### A. Experiment motivation

The main motivation behind conducting HQEMU live migration between Local-Host and Cloud-Host across the Internet is to illustrate that LivCloud basic design can be implemented on uncontrolled environment, Amazon's datacenter without any enhancements from the next stage of LivCloud.

### B. experiment setup

To achieve the basic design stage, a private network (172.16.10.0/24) based in Bournemouth (UK), which has two physical servers (Local-Host and NFS server) is connected to a Ubuntu server 14.04 (private address, 10.0.0.10/24) on Amazon's datacentre in London, UK. The experiment aims to evaluate LivCloud's basic design within the mentioned environment. Thus, the lab setup as shown in Figure 2 consists of one HP Z440 workstation is connected to the Internet through EdgRouter X and Netgear L2 switch providing 1Gbps. The workstation has 32GB of RAM, 1TB disk and 4-core 2.8GHz Intel(R) Xeon(R) E5-1603 v3 CPU. 64-bit Ubuntu Server 14.04 LTS, HQEMU (Layer 1 hypervisor) and HQEMU routed network are installed and configured on the machine, *Local-Host*. The other machine on the private network is configured as an NFS server (FreeNAS 9.3) for the lab. The Amazon c4.2xlarge instance 64-bit Ubuntu server 14.04, *Cloud-Host* is connected through a network card providing 1Gbps. The server has 15GB of RAM, 100GB disk and 8 vCPU 2.9GHz Intel(R) Xeon(R) Xeon E5-2666 v3. HQEMU (Layer 2 hypervisor) and HQEMU routed network are installed and configured on the instance. Any VM on either Local-Host or Cloud-Host can be configured with a local disk or a disk hosted on the lab NFS server. Using HQEMU, VMs, 2 VMs, 32-bit Windows XP, xp-1 and 64-bit Ubuntu server 14.04, ub-14-sh used as the migrated VMs between both hosts. Their disks are hosted

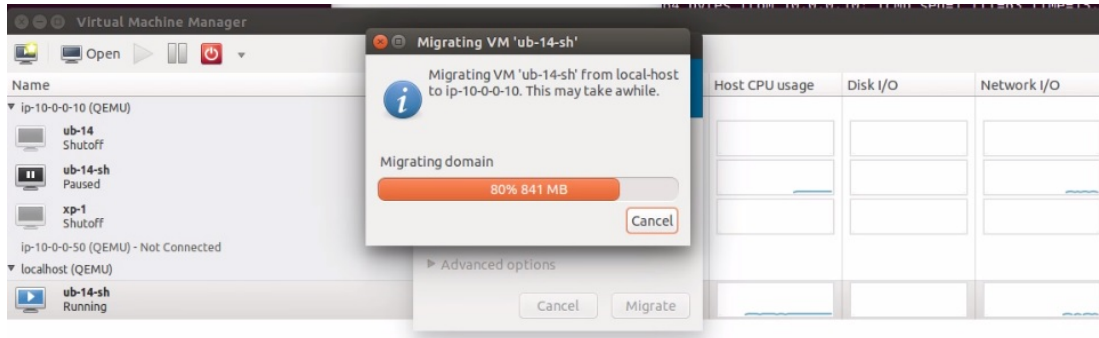


Fig. 3. Virtual manager's connection to both hosts

```

root@local-host:~# ping
PING
64 bytes from Amazon Public: icmp_seq=1 ttl=49 time=14.5 ms
64 bytes from IP: icmp_seq=2 ttl=49 time=12.8 ms
64 bytes from IP: icmp_seq=3 ttl=49 time=12.9 ms
^C
--- ping statistics ---
4 packets transmitted, 3 received, 25% packet loss, time 3004ms
rtt min/avg/max/mdev = 12.896/13.455/14.503/0.741 ms
root@local-host:~# ping cloud-host
PING cloud-host.local.lan (10.0.0.10): 56(84) bytes of data.
64 bytes from cloud-host.local.lan (10.0.0.10): icmp_seq=1 ttl=63 time=17.0 ms
64 bytes from cloud-host.local.lan (10.0.0.10): icmp_seq=2 ttl=63 time=16.1 ms
64 bytes from cloud-host.local.lan (10.0.0.10): icmp_seq=3 ttl=63 time=14.5 ms
64 bytes from cloud-host.local.lan (10.0.0.10): icmp_seq=4 ttl=63 time=14.9 ms
64 bytes from cloud-host.local.lan (10.0.0.10): icmp_seq=5 ttl=63 time=15.8 ms
64 bytes from cloud-host.local.lan (10.0.0.10): icmp_seq=6 ttl=63 time=15.9 ms

```

Fig. 4. Latency comparison between Internet connection and IPsec VPN

on the NFS server. The Windows VM has 1GB of RAM, 2vCPUs and 5GB of disk. Whereas, the Ubuntu VM has 2GB of RAM, 2vCPU and 8GB of disk. The private network and the Amazon VPC network are securely connected via IPsec VPN tunnel. Local-Host and Cloud-host are connected through the tunnel via the VM manager that is installed on Local-Host as shown in Figure 3. VPC is Amazon Virtual Network that helps building user-defined private network subnets inside the cloud in order to facilitate controlling IP address changes [24]. Furthermore, Dynamic DNS is used to maintain the migrated VMs' connections and configurations (P1). Dynamic DNS is used to keep a domain name pointing to the same physical or virtual server connected to the Internet regardless of any IP addresses changes [30]. no.ip is a dynamic DNS provider that has been chosen to register ub-14-sh and xp-1 under the DNS name records, ub-14-sh.ddns.net and xp-1.ddns.net respectively. Dynamic DNS clients (noip-2.1.9-1) are installed and configured on both migrated VMs [10]. Also, to prove that it can achieve flexibility and security despite that it is not possible to conduct a successful migration at this stage. Moreover, the migrated VMs' hardware specifications in respect to RAM and disks sizes are larger than the migrated VMs in previous approaches [1][3][7].

## V. EXPERIMENT RESULTS DISCUSSION

In this scenario, both hosts have HQEMU bridged or routed network installed and configured because KVM modules cannot be loaded on Amazon EC2 instances. IPsec VPN tunnel

is configured between the Local-Host's private network and Amazon VPC. Local-Host and EC2 Cloud-host are connected through the tunnel via the virtual manager that is installed on Local-Host as shown in Figure 3. The migration process of ub-14-sh is also shown in this figure. Amazon VPC provides two public IPs to VPN tunnel for load-balancing. Figure 4 shows a comparison between the latency (RTT) of a direct ping from Local-Host to Cloud-Host's public IP and the latency of a ping through the IPsec VPN from Local-Host to Cloud-Host's private IP. The private network and the Amazon VPC network are securely connected via IPsec VPN tunnel.

### A. Live migration with shared disk

Despite the successful completion of the migration of 2 VMs with shared disks (xp-1 & ub-14-sh) from Local-Host to Cloud-Host, it is necessary to restart both VMs to fix the halt state on xp-1 and the kernel panic on ub-14-sh. The average total migration time of ub-14-sh is just above 3 minutes, whereas, it is about 2 minutes in xp-1 migration. Furthermore, the performance of both VMS is rather slow despite compiling HQEMU [29] instead of QEMU. Figure 5 shows the kernel panic of ub-14-sh.

### B. Live migration without shared disk

Live migration of VMs disks has been considered in many studies [3][16]. However, it is considered to be unreliable and needs synchronization between CPU processing speed and network bandwidth [18]. Moreover, many cloud users prefer keeping VMs disks in-house for more control and privacy [1][8][22]. As mentioned earlier, LivCloud uses HQEMU that is an enhancement of QEMU. QEMU has a live block migration feature that allows migrating the disks state [18].

```

[ 964.295276] [ffffffffffb16040] ? net_init_array+0x50/0x55
[ 964.295276] [ffffffffffb160120] ? early_idt_handler_array+0x120/0x120
[ 964.295276] [ffffffffffb1605e0] x86_64_start_reservations+0x2a/0x2c
[ 964.295276] [ffffffffffb16072d] x86_64_start_kernel+0x13d/0x14c
[ 964.295276] Code: e2 fc 74 b1 48 89 c1 48 89 40 48 8b 50 08 48 39 ca 74 44 f6 02 01 75 b3 48 8b 7
a 10 48 89 c1 48 83 c9 01 48 89 78 08 48 69 42 10 <48> 89 0f 48 8b 06 48 89 0a 48 63 e1 fc 48 89 10
8f 84 bc 80 80
[ 964.295276] RIP: [ffffffffffb13e24e0] rb_erase+0x110/0x320
[ 964.295276] RSP: [ffffb0803fc03e08]
[ 964.295276] CR2: 0000000000000000
[ 964.295276] ---[ end trace e1f189e1e38c37a ]---
[ 964.295276] Kernel panic - not syncing: Fatal exception in interrupt

```

Fig. 5. Migrated Ubuntu VM's kernel panic



```

64 bytes from 1hr26s04-in-f14.1e100.net (216.58.198.238): icmp_seq=525 ttl=52 time=389 ms
64 bytes from 1hr26s04-in-f14.1e100.net (216.58.198.238): icmp_seq=526 ttl=52 time=411 ms
ping: sendmsg: No buffer space available
[ 1012.752216] NMI watchdog: BUG: soft lockup - CPU#1 stuck for 23s! [ping:1230]
[ 1040.484107] NMI watchdog: BUG: soft lockup - CPU#1 stuck for 23s! [ping:1230]

```

Fig. 6. Migrated Ubuntu VM's halt state

However, during the evaluation of LivCloud, this feature showed instability and the process crashed many times. However, before crashing both VMs continue working for almost 2 minutes and Dynamic DNS's records are correctly update with the new public IP. The total migration time is approximately 15 minutes of both VMs due to the disks sizes. As a result, live migration of the VM's disk is cloud users' decision to either use this feature or leave the disk on the shared storage in LivCloud. Figure 6 shows the crushing of the migrated Ubuntu VM. In the next section, a number of solutions to enable nested virtualization on Amazon EC2 are discussed. Some of them have been tested during the implementation process and some of them are still being evaluated.

## VI. POSSIBLE SOLUTIONS TO ENABLE NESTED VIRTUALIZATION

To load KVM modules on public IaaS, the hardware-assisted virtualization features must be enabled. To check if the IaaS has these features, the KVM acceleration must be enabled, VMX or SVM flags's number should be larger than 0 and the nested virtualization of `kvm_intel` must be set to 'Y'. Figure 7 proves that these features are not enabled on Amazon EC2 instances. There is a number of possible solutions to

```

root@ip-10-0-0-100:~
root@ip-10-0-0-100:~# kvm-ok
INFO: Your CPU does not support KVM extensions
KVM acceleration can NOT be used
root@ip-10-0-0-100:~# egrep -c '(vmx|svm)' /proc/cpuinfo
0
root@ip-10-0-0-100:~# cat /sys/module/kvm_intel/parameters/nested
N

```

Fig. 7. Hardware-assisted virtualization features disabled on EC2

enable these values and consequently, enable the hardware-assisted virtualization on public cloud IaaS. The following solutions have been explored as part of the implementation process.

### A. Recompiling Amazon C4 Linux instance's kernel

This solution aims to recompile the instance kernel with specific features enabled such as, KVM acceleration support as shown in Figure 8 using the latest version of Linux kernel [11] and `menuconfig` command [12]. The `menuconfig` command is a menu-based user interface that rebuilds Linux kernel with selected options. Because Amazon instances' BIOS cannot be reached, `menuconfig` is an alternative tool to enable many hardware features on the instances. The rebuilding process of the kernel takes about two hours to finish and includes also upgrading the grub file. However, the result of this process changes only the nested virtualization of `kvm_intel` from 'N'

```

config - Linux/x86 4.0.15 Kernel Configuration
* Virtualization
Virtualization
Arrow keys navigate the menu. <Enter> selects submenus --- (or empty
submenu ---). Highlighted letters are hotkeys. Pressing <Y> includes,
<D> excludes, <O> modularizes features. Press <Esc><Esc> to exit, <?> for
Help, </> for Search. Legend: [*] built-in [ ] excluded <D> module <>
[*] Virtualization
<> kernel-based Virtual Machine (KVM) support
<+> KVM for Intel processors support
<+> KVM for AMD processors support
[*] audit KVM MMIO
[*] KVM legacy PCI device assignment support (DEPRECATED)
<+> Host kernel accelerator for virtio net
<D> vhost_vsock TCM fabric driver
<D> vhost virtio-vsock driver
[*] Cross-endian support for vhost
<Select> < Exit > < Help > < Save > < Load >

```

Fig. 8. Recompiling the EC2 instance's kernel

to 'Y'. However, the other two features, the KVM acceleration and the VMX flags number, the process cannot change their values. This solution does not help improve the live migration process.

### B. Compiling Bareflank on Amazon EC2

Bareflank [17] is an open source, lightweight hypervisor, lead by Assured Information Security, Inc. which provides the minimum requirements to install other complete/complicated hypervisors, such as Xen, KVM and VirtualBox. To enhance Bareflank development, it is written in C++, and includes support for exceptions and the C++ Standard Template Library (STL) via `libc++`. It can be run on various operating systems, including Windows and Linux. If the compilation of the hypervisor is successful, it converts the operating system into a VM [17]. It is installed and configured on Amazon c4.2xlarge instance because it can force enabling hardware-assisted virtualization, but the configuration process to convert the operating system to a VM has been repeatedly interrupted and stopped by the Amazon hypervisor, Xen.

### C. Run a C script on Amazon EC2

It is a script written in C language that had been used to enable Intel VT-x on Mac Pro and other operating systems in 2008 [5]. This code is programmed to enable hardware-assisted virtualization on the Intel based machines that have these features in the CPU architecture, but there is no BIOS support for them provided that the BIOS does not lock these features. The result of running the script shows that the BIOS

```

root@ip-10-0-0-10:/home/ubuntu# ./msrmagic
MSR Magic 1.0 (C) 2009 George Styles www.georgestyles.co.uk

cpuid 0 is locked ON :)
cpuid 1 is locked ON :)
cpuid 2 is locked ON :)
cpuid 3 is locked ON :)

```

Fig. 9. The output of running the script on C4 instance

locks these features as shown in Figure 9.

There is a potential solution that is still under evaluation. It consists of reprogramming an existing Assembly code written to enable Intel VT-x and EPT on a Windows physical machine [19]. It switches on the values of the machine's CPU

control registers to enable nested virtualization features on Amazon. The reprogramming is based on enabling the code to modify the vCPU's control registers values on Amazon Ubuntu instance.

## VII. CONCLUSION

Given the current state of public cloud IaaS in terms of hardware-assisted virtualization features, VMs live migration is still challenging to cloud users. LivCloud approach is introduced to help successfully live migrate cloud users' VMs without services disruption across different public cloud providers. The basic design stage of this approach is implemented and evaluated on Amazon C4 instance. Although, the connectivity is securely maintained between Local-Host and Cloud-Host through Virtual Machine manager and IPsec tunnel, the migration process is not successfully completed due to the lack of nested virtualization feature on Amazon IaaS. We explore 3 developed options to enable nested virtualization on Amazon EC2. None of them have yielded the desired results. In our future work, we will investigate other alternatives, such as using open source libraries or hard coding approaches to enable virtualization features.

## REFERENCES

- [1] Qin Jia, Zhiming Shen, Weijia Song, Robbert van Renesse, and Hakim Weatherspoon. Supercloud: Opportunities and Challenges. *SIGOPS Oper. Syst. Rev.* 49, 1, January 2015, 137-141.
- [2] Ibrahim Mansour, Kendra Cooper and Hamid Bouchachia. "Effective Live Cloud Migration". In proceedings of The IEEE 4th International Conference on Future Internet of Things and Cloud (FiCloud 2016). DOI:10.1109/FiCloud.2016.54
- [3] Kaveh Razavi, Ana Ion, Genc Tato, Kyuho Jeong, Renato Figueiredo, Guillaume Pierre and Thilo Kielmann. Kangaroo: A Tenant-Centric Software-Defined Cloud Infrastructure. *IEEE International Conference on Cloud Computing (IC2E)*, March, 2015. DOI: 10.1109/IC2E.2015.19
- [4] Dan Williams, Hani Jamjoom, and Hakim Weatherspoon. The Xen-Blanket: virtualize once, run everywhere. *The 7th ACM European Conference on Computer Systems (EuroSys '12)*. ACM, New York, NY, USA, 2012, 113-126. DOI=10.1145/2168836.2168849
- [5] KVM. 2008. "Enable VT-X on Mac Pro (Early 2008)". Available at: [http://www.linux-kvm.org/page/Enable\\_VT-X\\_on\\_Mac\\_Pro\\_\(Early\\_2008\)](http://www.linux-kvm.org/page/Enable_VT-X_on_Mac_Pro_(Early_2008)).(Accessed: 23-02-2017).
- [6] Adel Nadjaran Toosi, Rodrigo N. Calheiros, and Rajkumar Buyya. Interconnected Cloud Computing Environments: Challenges, Taxonomy, and Survey. *ACM Comput. Surv.* 47, 1, Article 7, May 2014, 47 pages. DOI=10.1145/2593512
- [7] Alex Fishman, Mike Rapoport, Evgeny Budilovsky and Izik Eidus. HVX: Virtualizing the Cloud. *5th USENIX Workshop on Hot Topics in Cloud Computing*, 2013. USENIX Association, Berkeley, CA, USA.
- [8] Dan Williams, Hani Jamjoom, Zhefu Jiang and Hakim Weatherspoon. VirtualWires for Live Migration Virtual Networks across Clouds. Report by Cornell University and IBM T.J Watson Research Center, New York, April 2013.
- [9] OPENFV. Available at: <https://www.opnfv.org/>. (Accessed: 06-04-2016)
- [10] no.ip. 2017. Available at: <https://www.noip.com/>.(Accessed: 06-04-2017).
- [11] The Linux Kernel Archives. 2017. Available at: <https://www.kernel.org/>.(Accessed: 06-04-2017).(Accessed: 23-02-2017).
- [12] Creating custom kernels with Debian's kernel-package system. 2003. "Configuring the kernel". Available at: <http://newbiedoc.sourceforge.net/tutorials/kernel-pkg/config-kernel-pkg.html.en>.
- [13] Rajesaheb R. Kadam1 and Manoj Bangare. 2014. "A Survey on Security Issues and Solutions in Live Virtual Machine Migration". *The International Journal of Advance Foundation and Research in Computer (IAJFRC)*. ISSN: 2348 4853.
- [14] Naveed Ahmad, Ayesha Kanwal and Muhammad Awais Shibli. 2013. "Survey on Secure Live Virtual Machine (VM) Migration in Cloud". In *Proceedings of the 2nd IEEE National Conference on Information Assurance (NCIA)*. DOI: 10.1109/NCIA.2013.6725332
- [15] Sandor Acs, Miklos Kozlovsky and Peter Kacsuk. 2014. "A Novel Cloud Bursting Technique". In *Proceedings of the 9th IEEE International Symposium on Applied Computational Intelligence and Informatics (SACI)*. DOI: 10.1109/SACI.2014.68400050.
- [16] Konstantinos Tsakalozos, Vasilis Verroios, Mema Roussopoulos and Alex Delis. Time-Constrained Live VM Migration in Share-Nothing IaaS-Clouds. *The 7th IEEE International Conference on Cloud Computing (CLOUD)* (July, 2014). DOI: 10.1109/CLOUD.2014.18.
- [17] Assured Information Security, Inc. 2017. "BitVisor". Available at: <https://github.com/Bareflank/hypervisor>.(Access:12-04-2017)
- [18] OpenStack. Issues with Live Migration. Available at: [http://docs.openstack.org/openstack-ops/content/compute\\_nodes.html](http://docs.openstack.org/openstack-ops/content/compute_nodes.html). (Accessed:07-04-2017)
- [19] Code Project For those who code. 2015. "Virtualization for System Programmers". Available at: <https://www.codeproject.com/Articles/215458/Virtualization-for-System-Programmers>. (Accessed: 23-02-2017)
- [20] RIGHTSCALE. 2016. "Cloud Computing Trends: 2016 State of the Cloud Survey". Available at: <http://www.rightscale.com/blog/cloud-industry-insights/cloud-computing-trends-2016-state-cloud-survey>.(Accessed: 06-04-2017).
- [21] Ibrahim Mansour, Reza Sahandi, Kendra Cooper and Adrian Warman. 2016. Interoperability in the Heterogeneous Cloud Environment: A Survey of Recent User-centric Approaches. In *proceedings of the ACM International Conference on Internet of things and Cloud Computing (IC2016)*. DOI: <http://dx.doi.org/10.1145/2896387.2896447>
- [22] Dan Williams, Hani Jamjoom, and Hakim Weatherspoon. 2014. "Software defining system devices with the 'Banana' double-split driver model". In *Proceedings of the 6th USENIX conference on Hot Topics in Cloud Computing (HotCloud'14)*.
- [23] Mark Shtern, Bradley Simmons, Michael Smit and Marin Litoiu. 2012. "An architecture for Overlaying Private Clouds on Public Providers". In *Proceedings of the 8th International Conference and Workshop on Network and Service Management (CNSM) and System Virtualization Management (SVM)*. E-ISBN: 978-1-4673-3134-0
- [24] Amazon Web Services. "Amazon Virtual Private Cloud Connectivity Options". Available at: [http://media.amazonwebservices.com/AWS\\_Amazon\\_VPC\\_Connectivity\\_Options.pdf](http://media.amazonwebservices.com/AWS_Amazon_VPC_Connectivity_Options.pdf). (Accessed: 06-04-2017).
- [25] Zhenhao Pan, Qing He, Wei Jiang, Yu Chen and Yaozu Dong. 2011. "NetCloud: Towards Practical Nested Virtualization". In *Proceedings of the IEEE international Conference on Cloud and Service Computing (CSC)*. DOI:10.1109/CSC.2011.6138541
- [26] Mahdi Aiash, Glenford Mapp and Orhan Gemikonakli. 2014. "Secure Live Virtual Machines Migration: Issues and Solutions". In *Proceedings of the 28th International Conference on Advanced Information Networking and Applications Workshops (WAINA)*. DOI: 10.1109/WAINA.2014.35.
- [27] Amazon Web Services. 2017. "Elastic IP Addresses". Available at: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html>.(Access: 12-04-2017)
- [28] Changbin Liu and Yun Mao. 2013. "Inception: Towards a Nested Cloud Architecture". In *proceedings of the 5th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 13)*.
- [29] Ding-Yong Hong, Jan-Jan Wu, Pen-Chung Yew, Wei-Chung Hsu, Chun-Chen Hsu, Pangfeng Liu, Chien-Min Wang and Yeh-Ching Chung. 2014. "Efficient and Retargetable Dynamic Binary Translation on Multicores.". In *Proceedings of the IEEE Transactions on Parallel and Distributed Systems*. DOI Bookmark: <http://doi.ieeeecomputersociety.org/10.1109/TPDS.2013.56>.
- [30] Steve Russell. 2015. "DynamicDNS". Available: <https://help.ubuntu.com/community/DynamicDNS>. (Access: 10-01-2017)