EVALUATING LIVE MIGRATION PERFORMANCE OF A KVM-BASED EPICS

Lei Hu, Jiajie Li, Libin Ding, Xu Jun, YaLi Liu, Institute of High Energy Physics, Beijing, China
Dapeng Jin, Jian Zhuang, Institute of High Energy Physics, Beijing; China
State Key laboratory of Particle Detection and Electronics of China, Beijing*

Abstract

In this paper we present some results about live migration performance evaluation of a KVM-Based EPICS on PC.About PC,we care about the performance of storage,network and CPU.EPICS is a control system. we make a demo control system for evaluation, and it is lightweight. For time measurement, we set a monitor PV, and the PV can automatics change its value at regular time intervals.Data Browser can display the values of 'live' PVs and can measure the time. In the end, we get the evaluation value of live migration time using Data Browser.

Keyword: EPICS, CSS, KVM

INTRODUCTION

The virtualization is one of the hottest topics in nowadays distributed computation as they are one of the key elements in cloud computing. The aim is to improve the efficiency of resource utilization, to reduce the energy consumption, to increase the availability of applications etc. One of the most efficient techniques of virtualization is Live Migration. Based on this technique, one application can migrate from one machine to another with a minimal transfer time (about some ms) [1], [2]. In order to reach these facilities, it is offered a software layer called Hypervisor.

When discussing virtualization, there are two main components: the host (the phycsical machine) and the guest (the virtual machine). The software layer that separates them is the hypervisor. There are two major hypervisor types:The Type-1 Hypervisor is a thin software layer that is installed over the hardware components and provides services for the running virtual machines.The second type is mostly used for paravirtualization or software virtualization. The type 2 hypervisor runs virtual machines as processes over a full-featured operating system.

Kernel-based Virtual Machine (KVM)[3] is the default vir-tualization technology starting with version 6 of Red Hat Enterprise Linux. Supported by Red Hat, along with the 2.6.20 Linux Kernel version.The main reason is that KVM is an open source project, hence it is free to integrate it in any other software environment. However, the fact that it is free does not mean that performance is an issue, since recent tests show that it is at least as good as the more experienced technology.

Even if KVM is considered a Type-2 hypervisor, it has an interesting trait: it works only over hardware with support for virtualization, thus delivering performance similar to type-1 hypervisors. As it looks on figure 1. Migration describes the process of moving a guest from one host to another. In a live migration, the guest continues to run on the source host while its memory pages are transferred, in order, to the destination host. The time it takes to complete such a migration depends on network bandwidth and latency. If the network is experiencing heavy use or low bandwidth, the migration will take much longer.

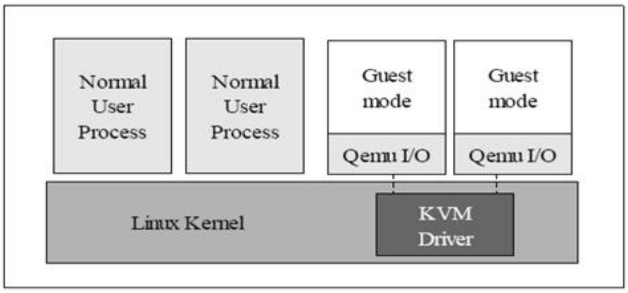

Figure 1: KVM hypervisor.

EPICS(Experimental Physics and Industrial Control System)[4] is a set of Open Source software tools, libraries and applications developed collaboratively and used worldwide to create distributed soft real-time control systems for scientific instruments such as a particle accelerators, telescopes and other large scientific experiments.

In this paper we present some results about live migration performance evaluation of a KVM-Based EPICS on PC. we make a demo control system and it is lightweight. For time measurement, we set a monitor for PV, and the PV can automatics change its value at regular time intervals.In the end, we get the evaluation value of live migration time using Data Browser.

RELATED WORKS

We use DesktopsPC(Personal computer) to finish the whole experimental works. The table 1 lists some import parts of computer. For OS(Operating System), we choose SL(Scientific Linux), The base SL distribution is basically Enterprise Linux(From RedHat Enterprise). Hardwares for experimence are depicted in Table 1.

Table 1: The Tech Specs of DesktopsPC

Processor	Intel(R) Core(TM)2 Duo CPU E7500 @ 2.93GHz
Operating System	Scientific Linux release 6.2 (Carbon), 64bit
Memory	4GB DDR3 SDRAM at 1333MHZ
Hard Drive	500GB 3.5" SATA 3Gb/s with 8MB DataBurst Cache

* E-Mail: hul@ihep.ac.cn, lijj@ihep.ac.cn, xujun@ihep.ac.cn,liuyali@ihep.ac.cn,jindp@ihep.ac.cn, zhuangj@Ihep.ac.cn

Softwares about KVM are depicted in Table 2.KVM (Kernel Virtual Machine) is a Linux kernel module that allows a user space program to utilize the hardware virtualization features of various processors.QEMU[5] can make use of KVM when running a target architecture that is the same as the host architecture. Libvirt[6] is a toolkit to interact with the virtualization capabilities of recent versions of Linux (and other OSes).The virt-manager[7] application is a desktop user interface for managing virtual machines through libvirt .

Table 2: Softwares adout KVM

qemu-kvm.x86_64	2:0.12.1.2-2.355.el6_4.6
libvirt.x86_64	0.10.2-18.el6_4.5
virt-manager.x86_64	0.9.0-7.el6

For virtual machine, we want a way to access to the host, to the internet or to resources available on your local network. At the same time, host system must be able to access the internet or the local network. We want to assign an IP address to virtual machine and make them accessible from your local network. We call this work public bridge, it is dipicted in Figure 2.

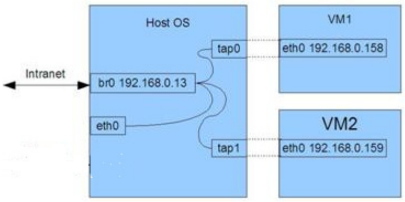

Figure 2: Network public bridge.

As it looks from figure 3 we have used three computers: one is for NFS server (IP address is 10.10.10.1). Shared storage can use one of the following protocols: Fibre Channel-based LUNs, iSCSI, FCoE, NFS. We use NFS for the shared storage.NFS is suitable for demonstration environments but for a production environment a more robust shared storage configuration using Fibre Channel or iSCSI is recommended. The another two computers is physical hosts, virtual machine can be run on one of them. All the computers connected by a Cisco Gigabit Switch with 24 gates. All the computers are connected by Fiber Channel communication and Fast Ethernet Interface 100/1000 Mbit/sec. The network topology is bus and all the packets are transmitted in the Layer 2 of OSI Architecture. All the computers are configured with C class IP, based on TCP protocol.

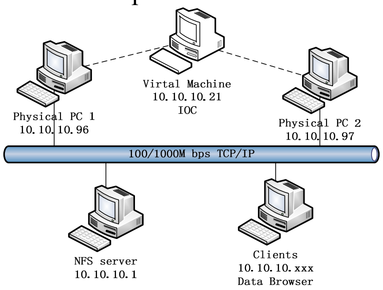

Figure 3: Illustration of the network.

EPICS is a control system architecture and a software toolkit. For EPICS, "client/server" model is Network-based, they speak of Channel Access role. Channel Access clients are programs that require access to Process Variables to carry out their purpose. The "service" that a Channel Access server provides is access to a Process Variable. A Process Variable(PV) is a named piece of data associated with the machine. Network-based "client/server" model where the basic data element is a Process Variable. The Channel Access Protocol defines how Process Variable data is transferred between a server and client.

We set up a EPICS demo which is a Channel Access server Providing a PV to emulate a analog input variable (AI). As a record, the PV processing can be periodic, so we set the period 0.1 second, 1 second and so on. This means the changing of PV's value is regular. The figure 5 shows that PV is changing with 1 second intervals. For EPICS demo, Figure 4 presents the EPICS demo system of two records. PV(rootHost:calcExample) is a CALC record, produces regular values with algorithm. FLNK field of rootHost:calcExample points to the PV rootHost:ai1. Result value of rootHost:calcExample cause VAL field of rootHost:ai1 to change[8]

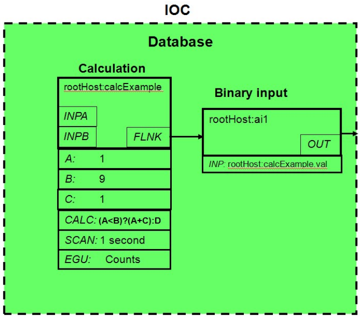

Figure 4: Simple EPICS demo.

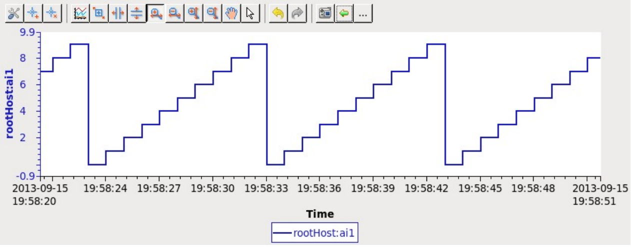

Figure 5: Value of rootHost:ai1.

TEST AND EVALUATING

For network traffic performance of virtual machine, we use Netperf[9]. Netperf can be used to measure the performance of many different types of networks and it provides tests for throughput, and end-to-end latency. It has two separate executable files: one for server side and other for client side. The test network comprises of two computers connected using CISCO switch. All netperf settings are showed in table 3.

ISBN 978-3-95450-139-7

Table 3: Netperf Parameters

Buffer size of local machine	8K
Buffer s ize of remote machine	8K
Running period	30 seconds
Protocol	TCP
Payload size	From 128 to 2048

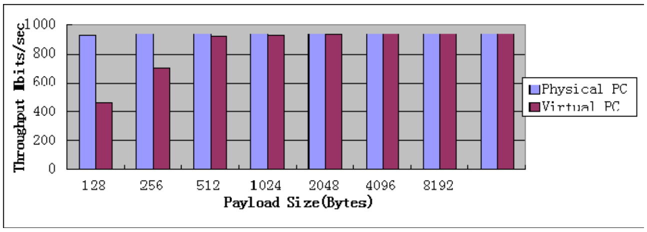

Figure 6: TCP throughput Test.

The PING command is used to test the connection and latency between two network connections. The PING command sends packets of information to a specified IP Address and then measures the time it takes to get a response from the specified computer or device.

Here, we make a test in a local area network. When you migrate a virtual machine from one physical Host to another, the servers such as OS, EPICS demo may be stop for a while. In figure 7, we use PING command to test the response time. In local area network, a client sends PING command to virtual machine in 0.1 second interval. If virtual machine can not response a PING command from client, then this PING is lost, and the client can not receive the feedback from virtual machine. We count the lost PING command and emulate the response time.

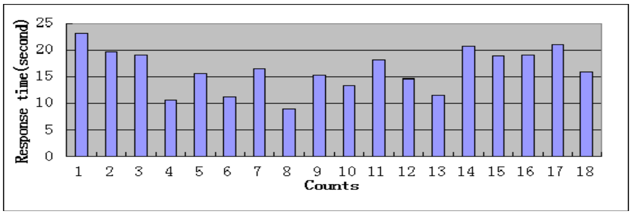

Figure 7: Response Time for Virtual machine.

The last, we test the EPICS demo. The PV rootHost:ai1 counts from 0 to 9 every 10 seconds or 1 seconds, it's depends on your settings. As figure 5, we use databrowser from CSS(Control System Studio) tools to illustrate the lines. First, stop the virtual machine for a while and restart the virtual machine, then figure 8 show the results. Second, when live migration completes from physical machine(10.10.10.96) to physical machine(10.10.10.97) or from physical machine(10.10.10.97) to physical machine(10.10.10.96), the IOC demo keeps running and was not interrupted. In figure 9 indicates that running state when live migrating.

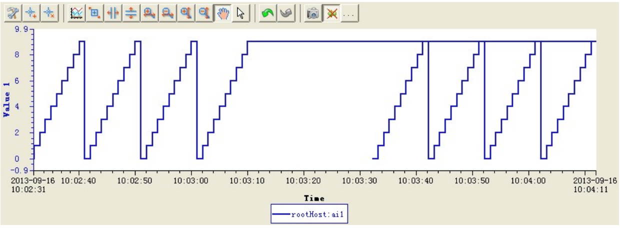

Figure 8: EPICS Demo when IOC stops for a while.

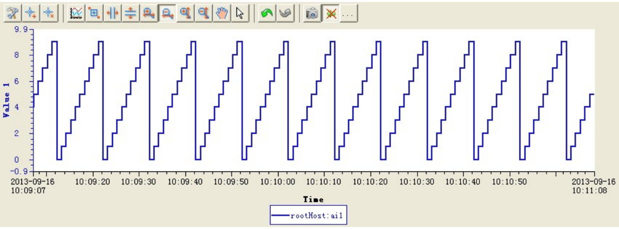

Figure 9: Live migration with virtual machine.

CONCLUSIONS AND FUTURE WORK

For TCP throughout test when transporting big TCP/IP packages, the performance of virtual machine is same as physical machine. For CPU load, virtual machine have used more source than physical machine. As figure 7, when a live migration keeps the EPICS demo running on a physical machine host to another physical machine host, the response time of OS is about 20 seconds. If we improve share storage and local network, the response time may be shorter.

When we make sure that stopping the epics demo, then we can illustrate the line that interrupts for a while. For live migration, When we migrate from one physical machine to another, then the phenomenon that IOC demo was interrupted was observed.

If the original guest modifies pages faster than KVM can transfer them to the destination host, offline migration must be used, as live migration would never complete. This test demo is lightweight. If the CPU load is heavy, we need more test for live migration with KVM. There are many options for improvement.

REFERENCES

[1] Christopher Clark, Keir Fraser, Steven Hand, Jacob Gorm Hanseny, Eric July, Christian Limpach, Ian Pratt, Andrew War_eld, "Live Migration of Virtual Machines",2010

[2] http://en.wikipedia.org/wiki/Live_migration

[3] Kivity, Y.Kamay, D.Laor,and U.Lublin, "kvm the Linux virtual machine monitor," Proceedings of the Linux,2007.

[4] EPICS Home Page [Online]. Available: http://www.aps.anl.gov/epics/

[5] http://wiki.qemu.org/Main_Page

[6] http://libvirt.org/docs.html

[7] http://virt-manager.org/

[8] Martin R. Kraimer, Janet B. Anderson, Andrew N. Johnson. EPICS Application Developer's Guide[EB/OL]. http://www.aps.anl.gov/epics/base/R3-14/11-docs/AppDevGuide.pdf,2010-01-25。

[9] R. Jones , " Netperf 2.4.3," http://www.netperf.org/netperf/ .

ISBN 978-3-95450-139-7