Testing Docker Performance for HPC Applications

Alexey Ermakov `ermakov.as@mipt.ru`
Alexey Vasyukov `a.vasyukov@phystech.edu`

April 20, 2017

The main goal for this article is to compare performance penalties when using KVM virtualization and Docker containers for creating isolated environments for HPC applications. The article provides both data obtained using synthetic tests (High Performance Linpack) and real life applications (OpenFOAM). The article highlights the influence on performance of major infrastructure configuration options – CPU type presented to VM, networking connection type used.

Keywords: Docker, KVM, MPI, HPL, OpenFOAM, benchmark

1 Introduction

One of the most important issues related to high performance computing that one may encounter is the availability of certain execution environment. It means that many scientific programs require a specific set of dependencies (such as compilers, runtime libraries etc.), that often may even conflict with dependencies of other software. There is a number of ways to solve the issue, one the most mature technologies that is

used for such a purpose is a virtualization. Despite tha fact that virtualization provides full environment isolation, by-design it has some performance penalty. Another approach to provide isolated environment is operating-system-level virtualization that implies all such environments have common kernel and separate isolated user-space libraries. The main goal for this article is to compare performance penalties when using two mentioned ways of creating isolated environment (KVM and Docker containers, to be precise).

2   Related work

Cloud computing environments for HPC applications are commonly based on KVM for virtualization and isolation and OpenStack for cluster management, auto-provision and user self-service. An example of the system based on these technologies can be found in [1], describing an experience of Technische Universitat Dresden. Similar KVM-based clusters are deployed in different organizations over the world. However, performance penalties for real life applications may be significant when running in virtualized environment [2]. Container-based systems for HPC applications emerge during recent years [3–5] and benchmarks look promising [6, 7]. This article also contributes to public benchmarks of KVM and Docker containers for HPC applications.

3   Virtual machines and containers

As it was mentioned earlier, the main difference (see fig. 1) between virtualization and containerization is that containers share the same kernel and maybe even some host devices, when each virtual machine has its own kernel and virtualized devices (e.g. network card)[1]. More information on used technologies may be found in official documentation for KVM [8, 9] and Docker [10].

---

[1]In this article we do not consider usage of paravirtualization or any «passthrough» technologies to make host devices available to virtual machine
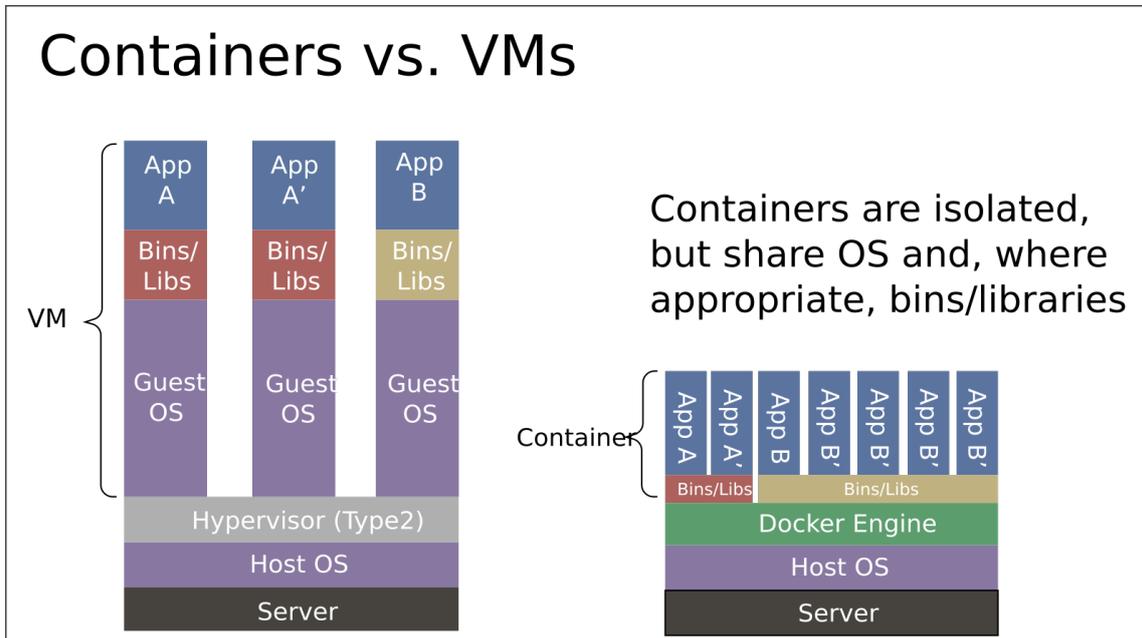
Figure 1: Comparsion of virtualization and containers architecture[2]

4 Benchmark setup and methodics

To perform benchmark the following setup was used: two identical hosts with Intel Core i7-5820K CPU (6 physical cores) and 64 GB RAM, connected with QDR Infiniband and 100 MB/s Ethernet. Hyperthreading was disabled using corresponding BIOS settings, since it drastically decreases performance (see fig. 2). MPICH was used as MPI implementation, because it's a bit faster than OpenMPI and does not require any configuration to execute program on two hosts when they belong to different subnets (this is very important to be able to inter-connect virtual machines and containers.)

Benchmarks described in this article use Intel Linpack benchmark [11], High-performance linpack [12] and interFoam solver as «real world» application from OpenFOAM [13]. All experiments were run 10 times to reduce statistical errors, so each plot shows mean value for measurement

[2]Image courtesy of Linux Magazine http://www.linux-magazine.com/Issues/2015/171/Docker#article_f2

and error bars for confidence interval of 0.95.
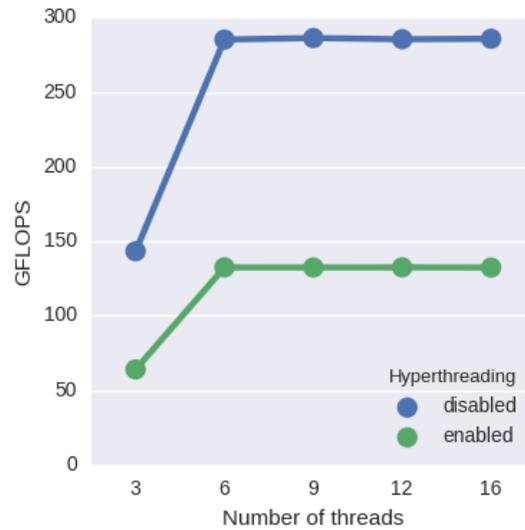
Figure 2: Hyperthreading performance impact according to Intel Linpack benchmark results
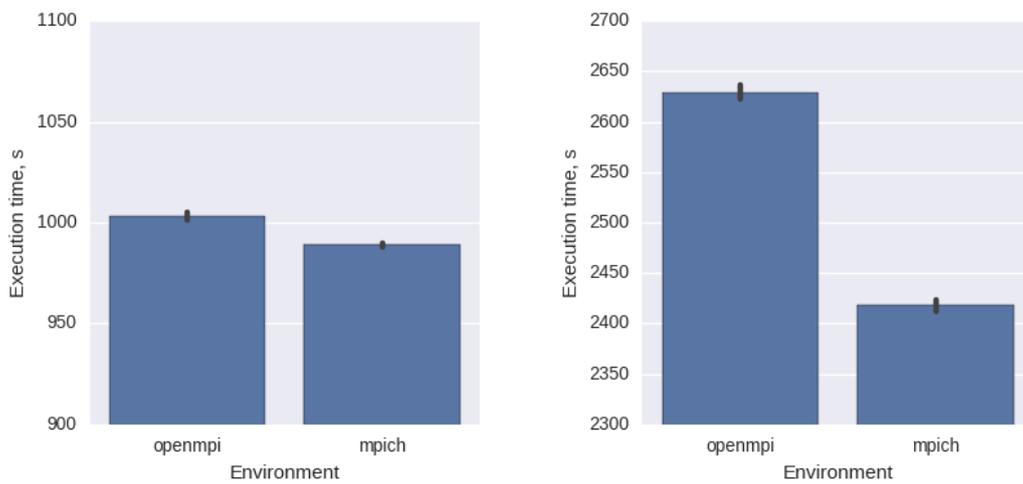
(a) QDR infiniband connection

(b) 100 Mb/s Ethernet connection

Figure 3: Comparsion of OpenMPI and MPICH performance using HPL

5 Benchmark

5.1 1 host benchmarks

First of all let's see how usage of containers and virtualization impacts performance. The following tests were performed on a single to host to avoid networking influence on performance results.
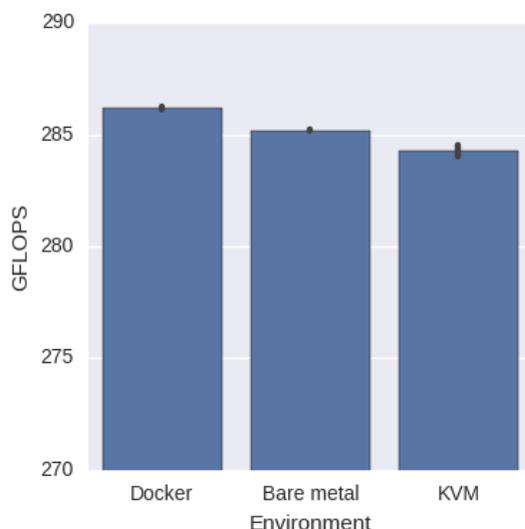
Figure 4: Performance comparsion of KVM and Docker on a single host using Intel Linpack

Results shown on a fig. 4 should be treated as follows: there is is no significant difference in performance when running CPU-intensive highly-optimized application in KVM, Docker or on bare metal. It should be noted that in these tests Intel Linpack demonstrates 90% of theoretical CPU performance, thus performance comparsion may be considered reliable. We can see that Docker shows a bit better performance even than bare metal, but it should be considered as statistical error. Another cause for this may be operating system scheduler that for some reason gives a bit more priority for containerized processes.

In previos test QEMU was run with host-model CPU set, that's why

it showed pretty good performance. In case of different CPU type there is a huge performance spread: depending on exact CPU model used to run a virtual machine KVM may be up to 5 times slower than a bare metal.
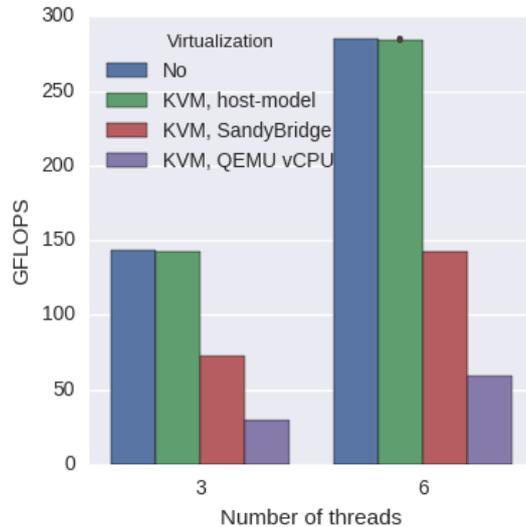
Figure 5: KVM performance spread

5.2   2 hosts benchmarks

Another component, besides CPU, that has a major performance influence is networking. The next series of tests demonstrate what performance impact one may have when using unappropriate networking stack.

First thing to note is quite obvious but anyway should be mentioned: networking type matters. According to fig. 6 HPL performs about 2.5 slower on 100 Mb/s Ethernet than on QDR Infiniband.

Another thing that should be taken into account is the way virtual machine or container is connected to network. For HPL tests there is almost no difference (see fig. 7) in performance of dockerized network application between bridged connection and host networking stack, KVM
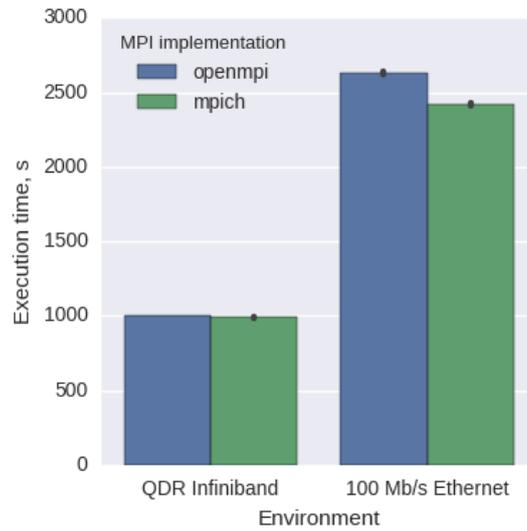
Figure 6: HPL performance comparsion using different inter-connection
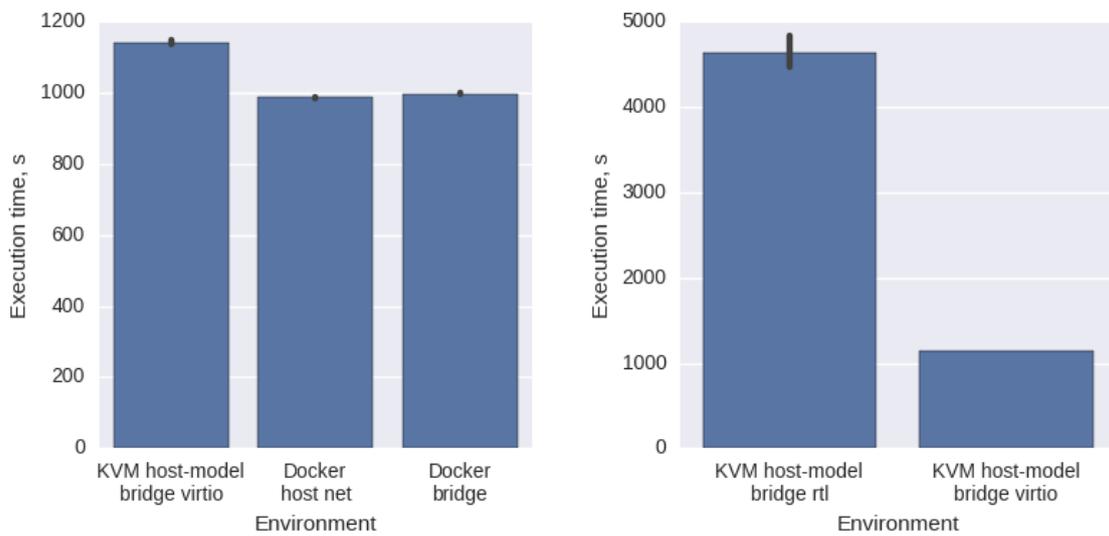
Figure 7: Docker and KVM networking performance comparsion

Figure 8: virtio and rtl networking performance comparsion

virtio performs about 15% slower while KVM rtl is almost 5 times slower (see fig. 8) than host netwoking stack.

Earlier we've noticed that CPU type used to run virtual machine has significant influence on overall performance. Let's see if it still applies to distributed MPI application. According to pic. 9 CPU used to run virtual machine does not affect resulting performance. Reason for this behaviour is the fact that HPL is not as CPU-intensive as Intel Linpack: network performance is more important for HPL rather than CPU performance.
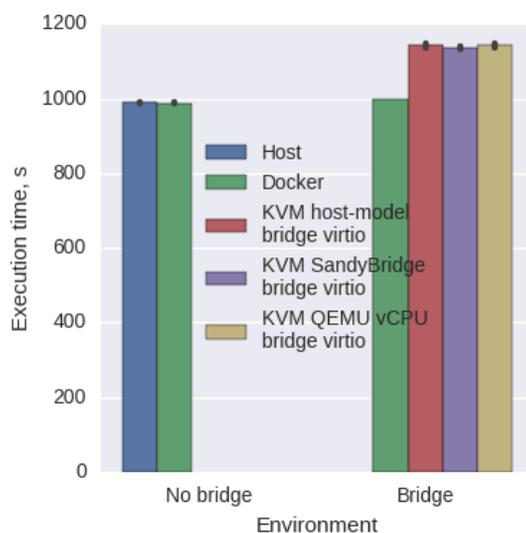
Figure 9: CPU type influence on performance of HPL

Finally, let's see how «non-synthetic» distributed MPI application performs depending on used environment. We used interFoam solver from OpenFOAM in few scenarios to realize how virtualization type and kind on networking connection influences overall performance. Results are shown on figs. 10 – 12. As we can see interFoam performs almost 10 times slower on 100 Mb/s Ethernet rather than on QDR Infiniband, bridge performance impact is about 20%.

Most important results are shown on fig. 12: a real «non-synthetic» application is more than 2 times slower in KVM with virtio networking than the same application in Docker with host or bridged networking.
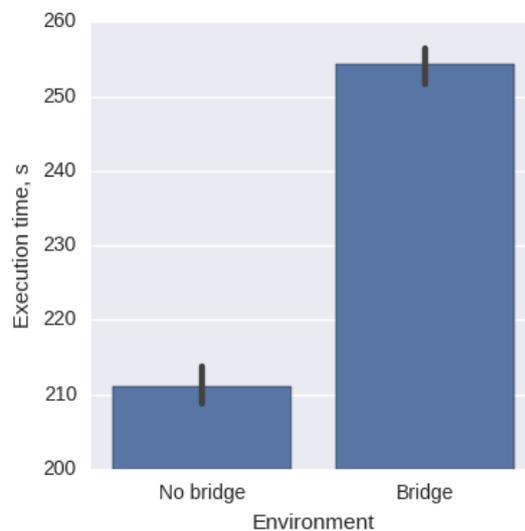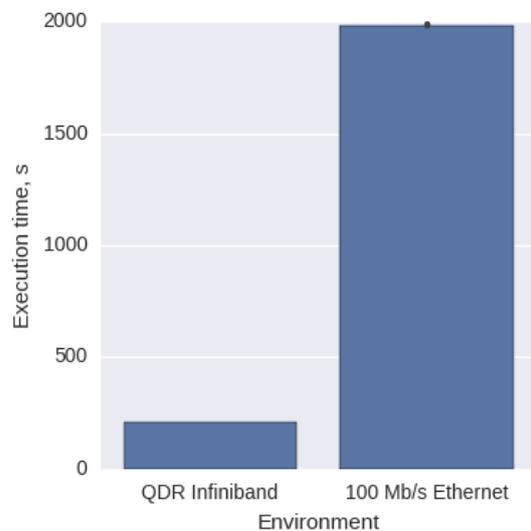
Figure 10: Networking type influence

Figure 11: Network connection type influence
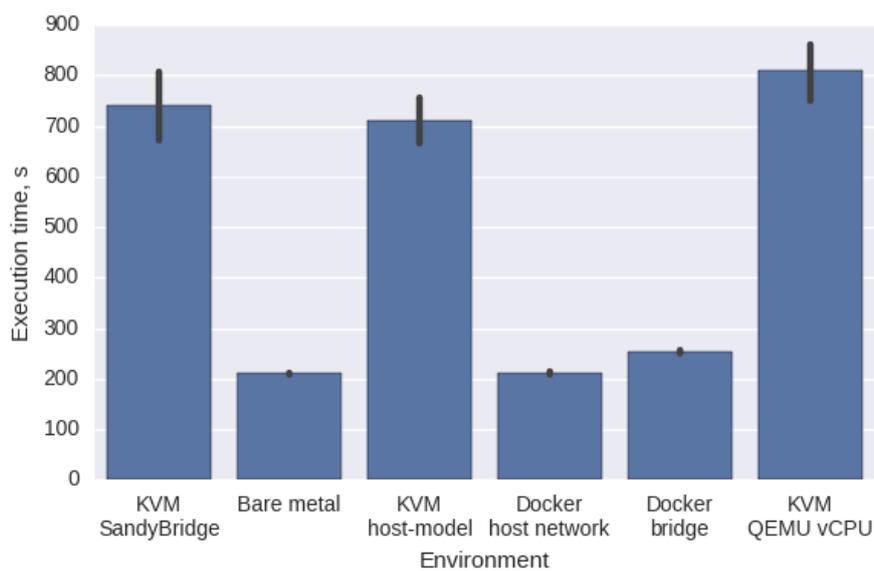
Figure 12: Comparsion of interFoam performance depending on used environment

9

6   Conclusion

Performance of virtualized or containerized applications depends on many factors, such as CPU type (for virtualization) and networking type. In some cases performance may degrade up to 10 times, thus environment to run application in must be carefully selected and verified. A really important thing is that «synthetic» benchmark does not provide you with a full exhaustive information necessary to decide if environment is suitable for an application or not. That's why it's strongly recommended to run benchmarking on exact application you're going to run when considering virtualization or containerization as an option for HPC.

References

[1] Ulf Markwardt. Running virtual machines in a slurm batch system. Slurm User Group, September 2015.

[2] Wes Felter, Alexandre Ferreira, Ram Rajamony, and Juan Rubio. An updated performance comparison of virtual machines and linux containers. Technical Report Technical Report RC25482 (AUS1407-001), IBM Research Division, Austin Research Laboratory, July 2014.

[3] Douglas M. Jacobsen and Richard Shane Canon. Contain this, unleashing docker for hpc. In *Cray User Group Conference Proceedings*, April 2015.

[4] Douglas Jacobsen, James Botts, and Shane Canon. Never port your code again – docker functionality with shifter using slurm. Slurm User Group, September 2015.

[5] H.-E. Yu and W. Huang. Building a Virtual HPC Cluster with Auto Scaling by the Docker. *ArXiv e-prints*, September 2015, 1509.08231.

[6] Christian Kniep. Containerization of high performance compute

workloads using docker. Technical report, QNIB Solutions, November 2014.

[7] B. Varghese, L. Thamsuhang Subba, L. Thai, and A. Barker. Container-Based Cloud Virtual Machine Benchmarking. *ArXiv e-prints*, January 2016, 1601.03872.

[8] Kernel virtual machine project. `http://www.linux-kvm.org/page/Main_Page`. Accessed: 2016-05-12.

[9] Qemu process emulator. `http://www.qemu-project.org/`. Accessed: 2016-03-17.

[10] Docker documentation. `https://docs.docker.com/`. Accessed: 2016-02-25.

[11] Intel math kernel library benchmarks. `http://software.intel.com/en-us/articles/intel-math-kernel-library-linpack-download`. Accessed: 2016-01-11.

[12] High-performance linpack benchmark for distributed-memory computers. `http://www.netlib.org/benchmark/hpl/`. Accessed: 2016-01-11.

[13] Openfoam, open source cfd software. `http://www.openfoam.com`. Accessed: 2016-03-02.