

Assignment 1 – Lightweight vs. Heavyweight Virtualization Techniques

(Due by Midnight, Wednesday, September 20, 10% Grade)

1 Summary

In this assignment, you will play with two virtualization techniques: one is Docker (a light-weight, container-based virtualization technique), and the other is QEMU (a user-level full virtualization technique). You will need to correctly deploy them, be familiar with the operations, measure the performance using benchmarks, and finally compare the results in a comprehensive report.

2 Environment Setup

Please follow this document [1] to start a Google cloud instance. We refer to this instance as the native. We will configure this instance using **2 vCPU, 4 GB memory, and 20 GB disk size** for this assignment. Please follow the document [1] to install the corresponding Linux system.

3 Docker

3.1 Installation & Preparation

Please refer to this instruction [2] for Docker installation. Notice that, please install Docker CE (choose using the repository). In addition, you could get familiar with Docker basic operations by referring to [2].

3.2 Docker Measurement

Please use the following container image, with Sysbench pre-installed: csmhpp/ubuntu-sysbench. You will measure the performance of Docker using Sysbench.

Sysbench is a benchmark tool [4]. We can have some basic ideas about the system performance by running such benchmark tools against the system under test. In this project, we will focus on the cpu and fileio test modes (section 4.1 and 4.5 in [4]) of Sysbench.

Notice that the following instructions present some ideas about how to conduct a meaningful measurement. Please read them carefully, and then produce measurement results (you will lose points, if the results are incorrectly generated and/or can not be well justified).

- You need to figure out all the test cases (under the cpu and fileio test modes).
- For each test case, you need to figure out the “right” parameters. For example, you need to select the parameter for “-cpu-max-prime” to ensure your test case won’t end in a short time period (too small) or will never end (too large). For example, you need to decide the file size (e.g., -file-total-size) for Sysbench fileio test mode. Usually, it’s reasonable to make a test case lasting for at least 30 seconds. Please justify your configurations.
- You may want to repeat/re-run at least 3 times for each test case, and report the average value of your results (the Sysbench will report some user-level performance data, e.g., total time).
- To reduce test variation, you need to test each case under similar test environment. For example, for the fileio test, after one test, the operating system will cache the accessed files. If we don’t manually drop such cache, the following tests will finish much faster, as most I/O data will be served directly from the

memory cache instead of disks. You have to manually drop such cache for example using the command: `echo3 > /proc/sys/vm/drop_caches` in the native (not the container). (Notice that, you can only run this command under the root user. Please figure out how to enable root user under Ubuntu and switch to the root user to drop the cache before each fileio test).

- In addition to the user-level performance data generated by the benchmark, you need to collect system performance data (from the native). For example, you can choose to use `iostat` (a Linux tool for collecting cpu, I/O, network usage [5]). You must provide the cpu and disk I/O performance data. Please report the commands your use and the interpretation of your data.
- Please well organize your experimental setup, configurations, and data (using figures and/or tables); and report them.

4 QEMU

4.1 Installation & Preparation

You can install `qemu` (under Ubuntu) simply by:

```
$sudo apt-get install qemu
```

You will download the Ubuntu iso image to install your QEMU VM:

```
$wget http://mirror.pnl.gov/releases/16.04/ubuntu-16.04.3-server-amd64.iso
```

You need to create an image before installing your QEMU VM. For your convenience, we provide the commands below. Please refer to the QEMU manual for more details [6].

```
$sudo qemu-img create ubuntu.img 10G
```

4.2 Install a QEMU VM

Because installing a VM require a graphical user interface (GUI). Please refer to Section 7 of this document to enable the GUI of your Google instance (i.e., the native). You will do the following operations within GUI (you should use a VNC client). Please install the QEMU VM using the command below (which takes the iso file as a “cdrom” and the `qemu` image as a “hard disk”):

```
$sudo qemu-system-x86_64 -hda ubuntu.img -boot d -cdrom ./ubuntu-16.04.3-server-amd64.iso -m 1536
```

Then, install the Ubuntu Linux system step by step (it may take some time to complete the whole installation).

4.3 QEMU VM Measurement

Restart your QEMU VM after installation. Install Sysbench workload in your QEMU VM, and re-do all measurements (those in Section 3.2) under the QEMU VM by following all the guidelines.

5 The Native

We refer to your Google cloud instance as the native. You will need to re-do all the measurements (those in Section 1.3 and 2.6) under the native by following all the guidelines.

6 Submission & Grading

Submit your assignment report on the blackboard as one pdf file. We will grade your assignment based this report. You should not only follow the instructions to do your assignment, but also fully understand what you are doing. The report should include (but not limited to):

- Detailed configurations of your experimental setup (your Google Cloud instance). – 5 points
- Present main steps to enable a Docker container. In addition, please describe the operations you use to manage Docker containers (and some other operations which you think are also important). – 10 points
- Present main steps to enable a QEMU VM. In addition, please present the detailed QEMU commands, and VM configurations. – 10 points
- Proofs (screen snapshots of your Docker and QEMU running environments). – 5 points
- Present how you conduct your measurements in three different scenarios (the native, Docker, and QEMU). – 20 points
- Present how you use performance tools to collect performance data. For CPU utilization, you should at least divide them into two parts including user-level and kernel-level. For I/O, you should present I/O throughput, latency, and disk utilization – 20 points
- Please understand the performance data, analyze the data, and then present them in an understandable way in your report. – 20 points
- Writing – 10 points
- Bonus – 10 points. If you want to have more in-depth understanding about what is happening in your Linux system, you probably don't want to miss Perf (a powerful Linux profiler tool). You can get the bonus points by demonstrating two profiling usages of perf, and explaining collected data (5 points for each).

7 Appendix

Please follow the instructions in this article [7] to enable GUI for the native (your Google instance). Notice that, however, there are two places you need to pay attention, and following the instructions below may solve possible problems.

7.1 .vnc/xstartup

Replace the whole content of .vnc/xstartup with the following content (DO NOT follow the article above).

```
#!/bin/sh
def
export XKL_XMODMAP_DISABLE=1
unset SESSION_MANAGER
unset DBUS_SESSION_BUS_ADDRESS

metacity &
gnome-settings-daemon &
gnome-panel &
nautilus &
gnome-terminal &
```

7.2 Open the firewall

Navigate to the “Firewall rules” item (in Google Cloud Console) shown as in Figure 1. Create a new firewall rule as in Figure 2. Notice that though this is an easy way to open ports for vncserver, it is not a safe solution (as we open too many ports). You can choose to use a safer way to do so.




Figure 1: Open firewall rules.

References

- [1] How to Start an Instance in GCP. http://www.cs.binghamton.edu/~huilu/slides/How_to_use_google_cloud.pdf.
- [2] Docker Installation. <https://docs.docker.com/engine/installation/linux/docker-ce/ubuntu/#install-docker-ce>.
- [3] Running your first container. <https://github.com/docker/labs/blob/master/beginner/chapters/alpine.md>.
- [4] Sysbench. <http://imysql.com/wp-content/uploads/2014/10/sysbench-manual.pdf>.
- [5] IOSTAT. <https://linux.die.net/man/1/iostat>.
- [6] QEMU <https://qemu.weilnetz.de/doc/qemu-doc.html> <https://qemu.weilnetz.de/doc/qemu-doc.html>.
- [7] Graphical user interface (GUI) for Google Compute Engine instance. <https://medium.com/google-cloud/graphical-user-interface-gui-for-google-compute-engine-instance-78fccda09e5c>.

← Create a firewall rule

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. [Learn more](#)

Name ?

Description (Optional)

Network ?

Priority ?

Priority can be 0 - 65535 [Check priority of other firewall rules](#)

Direction of traffic ?

- Ingress
 Egress

Action on match ?

- Allow
 Deny

Targets ?

Source filter ?

Source IP ranges ?

Second source filter ?

Protocols and ports ?

- Allow all
 Specified protocols and ports

Figure 2: Firewall rules.