

A. Artifact Appendix

A.1 Abstract

This artifact is composed of Makefiles and scripts that can be used to measure the latency experienced in KVM-based and Xen-based VMs (latency measurements are performed by using `cyclictest`). In particular, these scripts have been used to perform the experiments presented in the paper and can be used to reproduce similar experiments on different hardware/software configurations.

A.2 Artifact check-list (meta-information)

- **Program:** `cyclictest`, the Linux kernel, and `busybox` are downloaded during compilation; all the used programs are open-source software
- **Compilation:** the scripts download and compile some software (`cyclictest`, the Linux kernel with Preempt-RT patch, `busybox`, `sudo`) to run the experiments
- **Binary:** no binaries are contained in the artifact (everything is compiled from source code)
- **Run-time environment:** any modern Linux distribution (with the dependencies mentioned below) should be OK
- **Hardware:** Intel-based systems have been used in the paper, but in theory, the artifact can be used on every hardware supported by Linux and Xen
- **Run-time state:** the measured latencies depend on the runtime state, and the scripts provide some way to generate a “kernel stress” that trigger high latencies
- **Execution:** the scripts take care of properly pinning tasks and setting priorities when requested. The execution time depends on the host speed (a Linux kernel is recompiled at least 2 times), and several reboots (to select the appropriate host kernel, and to boot with the Xen hypervisor when needed) are required
- **Metrics:** latencies
- **Output:** some text files containing the eCDFs of the measured latencies, and some figures in eps format
- **Experiments:** typing “`make`” repeats the experiments presented in Figures 1 and 2 of the paper. The generated figure obviously depends on the host hardware, but the results should be consistent with the paper (KVM shows smaller latencies - and much smaller worst-case latencies in particular, on Xen using Preempt-RT on the guest increases the latency instead of decreasing it, etc...)
- **How much disk space required (approximately)?:** about 2.5GB (without compiling a host kernel)
- **How much time is needed to complete experiments (approximately)?:** several hours (the first 2 runs compile 2 Linux kernels)
- **Publicly available?:** yes
- **Code licenses (if publicly available)?:** GPL
- **Archived (provide DOI)?:**
<http://retis.santannapisa.it/luca/VMLatencies>

A.3 Description

A.3.1 How delivered

The scripts and Makefiles can be downloaded from the artifact website (the web page also provides some documentation).

A.3.2 Hardware dependencies

Although the artifact has been tested on various Intel and AMD CPUs (including an Intel Core Duo, various versions of Intel Core i7 and i5, several Intel Xeon, an AMD Ryzen Threadripper, and some other), every CPU supported by KVM and Xen should be usable. The achieved results will be comparable with the ones presented in the paper (and consistent with them).

A.3.3 Software dependencies

The Makefiles use the same computer (the host) to compile the needed software and to run the experiments. Hence, both the tools needed for compilation and the two hypervisors under testing have to be installed on the host.

This requires to install a Linux distribution on the host (every recent Linux distribution should work; Debian stable, Debian testing, Ubuntu 18.04, Linux Mint and OpenSUSE have been tested).

The tools needed for compilation are GNU `make`, `git`, `patch`, `gcc`, `binutils`, `flex`, `bison`, `bc`, and can be installed by using the package manager provided by the host distribution. Moreover, to compile the Linux kernel the “`libelf`” and “`libssl`” libraries (including the header files) must be installed. On a Debian, Mint, or Ubuntu system they can be installed with “`apt-get install libelf-dev libssl-dev`”.

The hypervisors needed to run the experiments are QEMU/KVM and Xen and are assumed to be installed on the host distribution. Both of them can be installed by using the package manager provided by the host distribution. However, some distributions provide a version of Xen that is too old for these tests (Xen 4.11 has been used in the paper). In this case, Xen can be installed from source as indicated in the Xen wiki (note that running “`make install`” is not recommended; use “`make debball`” or “`make rpmball`” instead, and then use the package manager).

Since some of the experiments compare the latencies experienced when using an RT host kernel (Linux with the Preempt-RT patches) with the ones experienced when using an NRT kernel (vanilla Linux), a Preempt-RT kernel (4.19.8-rt6 has been used in the paper) and a vanilla Linux kernel (as provided by the host distribution) must be installed on the host. The artifact provides a way to build a basic Preempt-RT kernel for the host (see next section).

Some scripts and Makefiles use the “`sudo` command, that must also be installed in the host (if it is configured not to require a password, the execution of the experiments is simplified).

Finally, the “`gnuplot`” program is used to generate the figures and must be installed in the host (the package manager provided by the host distribution can be used).

A.4 Installation

The artifact does not require any particular installation: just download the tarball and uncompress it (with “`tar xvf isorc-artifact.tgz`”) in a place with enough free disk space (from 1.5GB to 3GB, depending on the experiments you plan to run). The artifact will automatically download and compile most of the needed software.

The only software components that the artifact does not automatically install are the hypervisors (see the previous section) and the host's kernel. If you want to reproduce the experiments from the paper and a Preempt-RT kernel is not already installed on the host, it is possible to install a basic Preempt-RT kernel by following these instructions:

- Unpack the artifact: `tar xvf isorc-artifact.tgz`
- `cd Artifact/LatencyTests/VMTest`
- `make host` (this will download the Linux kernel sources, patch them with the Preempt-RT patch, and compile a “default” kernel with RT preemption enabled)
- `cd BuildHost; sudo make modules_install; sudo make install`

Notice that the configuration used to compile the Preempt-RT kernel is a simple one that should boot on standard PCs, if you use a non-Intel host or if you have particular devices (laptops often fall in this category) you might need to change the configuration in “BuildHost/.config” and recompile the kernel.

After this, the artifact is usable.

A.5 Experiment workflow

The latency experiments are driven by some Makefiles: the Makefiles in `Artifact/LatencyTests` allow to run generic latency measurements on the host or in KVM-based or Xen-based VMs, while `Artifact/Makefile` can be used for reproducing the run-time conditions of some experiments presented in the paper.

For example, by typing “`make exp12`” in the “`Artifact`” directory, it is possible to produce a figure similar to Figures 1 and 2 of the paper, by running the same experiments. Basically, this command uses “`Artifact/LatencyTests/VMTest`” to measure the latencies in Xen and KVM machines with various host and guest kernels (RT and NRT). Before starting the experiment, the Makefile checks if the host configuration (kernel and hypervisor) is correct, and if the configuration is not correct the Makefile prints a message (asking to reboot the host) and stops. To run all the experiments, 3 or 4 reboots are needed (to get the appropriate host configurations: no Xen NRT, no Xen RT, Xen NRT and Xen RT). The experiments are performed in a pre-defined order (KVM NRT host / NRT guest, KVM NRT host / RT guest, KVM RT host / NRT guest, KVM RT host / RT guest, Xen NRT host / NRT guest, etc...). Notice that the Makefile first checks if the needed hypervisor (KVM or Xen) can be used, and then checks if the needed kernel (NRT or RT) has been booted. Hence, if “Please reboot with a Preempt-RT Kernel” is (for example) printed, then the system has been booted with the correct hypervisor and the wrong kernel¹.

By typing by typing “`make exp3`” or “`make exp4`” in the “`Artifact`” directory it is possible to produce figures similar to Figure3 and Figure 4 of the paper.

The `Artifact/LatencyTests/VMTest` Makefile is more powerful (allowing to run different kinds of experiments), but more complex to use. The simplest way to invoke it is by typing “`make cdf-xen-<id>.txt`”, for measuring the latencies in a Xen VM, or “`make cdf-kvm-<id>.txt`”, for measuring the latencies

in a KVM VM (here, “<id>” is just an identifier for the experiment, at the user choice). These commands automatically compile a Preempt-RT kernel for the guest, build a simple guest image (a simple initramfs containing busybox, cyclicttest, and some scripts for running the experiment), run the experiment, and compute the eCDF of the latencies experienced on the guest. To run “`make cdf-xen-<id>.txt`”, the host must have been booted with the Xen hypervisor, and to run “`make cdf-kvm-<id>.txt`” the host must have been booted without the Xen hypervisor; if these prerequisites are not respected, the test fails. The Makefiles in the `Artifact` directory take care of verifying these prerequisites and notifying the user (asking to reboot with or without Xen, or with the correct host kernel) instead of just failing.

Running “`make xen-test`” or “`make kvm-test`” the experiment will be repeated 10 times.

The experiments are performed using an RT guest kernel, to run them with an NRT kernel, you can run “`make bzImage-nrt; cp bzImage-nrt bzImage`” before starting the experiments. The host kernel, instead, is selectable at boot time.

Using “`make cdf-xenstress-<id>.txt`” instead of “`cdf-xen-<id>.txt`” (or “`make cdf-kvmstress-<id>.txt`” instead of “`cdf-kvm-<id>.txt`”) it is possible to run the experiments with some background kernel stress load on the host (or on Dom0), used to trigger higher latencies.

Finally, `Artifact/LatencyTests/Makefile` allows to run some latency measurement on the host kernel (or on the Dom0 kernel, if the host has been booted with the Xen hypervisor), with or without background stress. This Makefile is used by the “`exp3`” and “`exp4`” targets of `Artifact/Makefile`.

A.6 Evaluation and expected result

After successfully running “`make exp12`” (or similar) in the “`Artifact`” directory, several files appear in the “`Figure12`” (or similar) directory. Notice that it might be necessary to run “`make exp12`” and reboot several times; after the last time, the “`make`” command is successful, and an eps file is generated. This file is consistent with the corresponding figure in the paper. The generated text files (for example, “`Figure12/cdf-***.txt`”) contain the results of the experiments executed for the various host/guest configurations.

As said, the obtained results are consistent with the ones presented in the paper, but the numbers cannot be equal. For example, looking at the results in “`Figure12`” it is possible to notice that:

- The KVM latencies are lower than the Xen ones
- Looking at the KVM results, changing the host and guest kernel from NRT to RT does not affect the shape of the CDF (all the KVM curves look overlapping)
 - However, the host and guest kernel affect the worst-case latencies (noticeable by looking at the “`Figure12/cdf-kvm-***.txt`” files): using an RT host and guest kernel noticeably reduce the worst-case latency
- On Xen, the host and guest kernel might affect the shape of the CDFs; moreover, using an RT kernel in the guest has the very strange effect of increasing the worst-case latencies (this is a strange result investigated in the paper)

Finally, notice that some details of the figures in the paper (for example, the X scale and the frequency of displayed points) have been adjusted to make the figures more readable. This manual

¹ if, instead, “Please reboot without Xen” is printed then we know that the system has been booted with Xen — while it should have been booted without Xen — but we do not know if the used kernel is the correct one — the kernel is checked only after the hypervisor check is successful

tuning requires to see the values generated by the experiments and is not performed on the figures generated by the artifact's Makefile (for example, the maximum X value on the figures has been set to an arbitrary value — $1000us$ for Figure12). The figures can be tuned by editing the various "plotcdf.gp" files contained in the "Figure*" directories.

A.7 Experiment customization

The experiments can be easily customized by modifying the test scripts used in the guest (in the "Artifact/LatencyTests/VMTest/GuestScripts" directory) or the Makefiles. Modifying the guest scripts allow changing the cyclicttest parameters (for example, the cyclicttest period, or duration) while modifying the host Makefiles allow to change the VM setup (for example, creating a VM with multiple CPUs, or changing the VM scheduling).