

CSC 2012 – Virtualization Exercises

Part I

Note: These exercises are available online under <http://cvmrepo.web.cern.ch/cvmrepo/csc12/cernvm-intro.pdf>

Today’s Synopsis

- Instantiate (start) a virtual machine
- Bootstrap and contextualize a fresh virtual machine
- Run LHC software inside the virtual machine
- Start and stop a cluster of virtual machines

Exercise 1 – Getting CernVM Out of the Box

CernVM is a Linux virtual machine that is designed to provide a complete and portable environment for developing and running LHC data analyses (Figure 1). CernVM runs on end user computers (laptop, desktop) with a graphical user interface as well as in a pure batch job environment on grid sites and on cloud resources. One of its biggest advantages is that it provides a *uniform* environment no matter if an analysis is being developed or executed and independent of the operating system of the host (Windows, Mac, or Linux).

CernVM comes with the CernVM File System, which runs inside the virtual machine and provides on-demand access to all the releases of the LHC experiment software. All files and directories under `/cvmfs` will only be downloaded if they are accessed. Once accessed, they are kept in a local hard disk cache. That allows for small virtual machine bootstrap images that grow only by the files and directories that are actually needed.

In this exercise, we will instantiate a “Desktop CernVM”, which you could use, for instance, to develop physics algorithms on your laptop. Here, we use VirtualBox as a hypervisor.

1. We will start your virtual machine on your machine in the computing center. This computer is what mimics your laptop in real life. We use

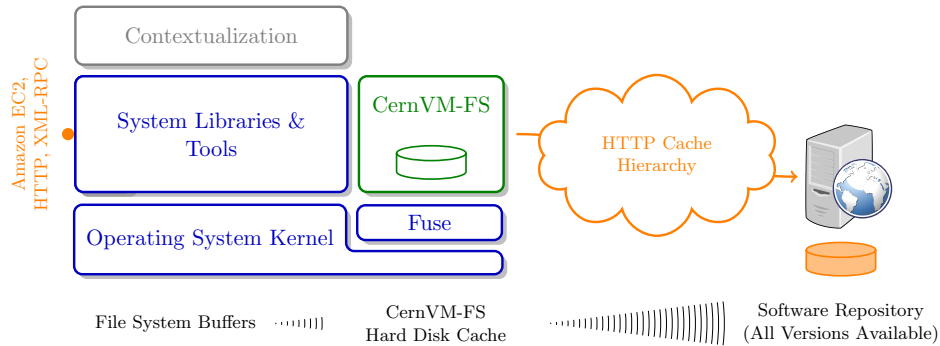


Figure 1: The key building blocks of a CernVM image: 1) Minimal Operating System, 2) CernVM-FS – read only network file system using HTTP protocol, 3) contextualization interfaces.

a remote machine because the processor of your local computer does not support virtualization extensions.

We will open a shell on your remote computer in the computing center. We will then clone the initial virtual machine image into the exercise directory, customize the virtual machine resources, and boot it with VirtualBox. Use the following commands to do so:

```
ssh -X <CSC-GROUP>@<REMOTE MACHINE>
cd cernvm/01cernvm
# cernvm-clone.sh makes a fresh, writable hard disk image
# using a read-only image from an image store as a template
cernvm-clone.sh first-cernvm.vdi
echo NAME=cernvm-desktop > virtual-hardware.conf
echo MEMORY=1024 >> virtual-hardware.conf # note the double ">"
```

Make a note of the image size. **Hint:** the Unix command `ls -lh` shows file sizes.

2. Use the provided script `cernvm-start.sh` in order to start VirtualBox with the CernVM image `first-cernvm.vdi` and the virtual hardware specification in `virtual-hardware.conf`:

```
cernvm-start.sh --disk=first-cernvm.vdi \
--hardware=virtual-hardware.conf
```

Don't worry about the VirtualBox info messages that might appear during the boot of your virtual machines.

3. We will now initially configure the CernVM. This is only necessary on the very first boot. The configuration is done through the CernVM Online web interface. Here, we customize the virtual machine and then we *pair* the specification with the virtual machine.

- Open a browser window and go to `http://cernvm-online.cern.ch`
- Login using your CSC group id as a user name and password (e. g. `csc01/csc01`). You will see the CernVM dashboard (Figure 2).

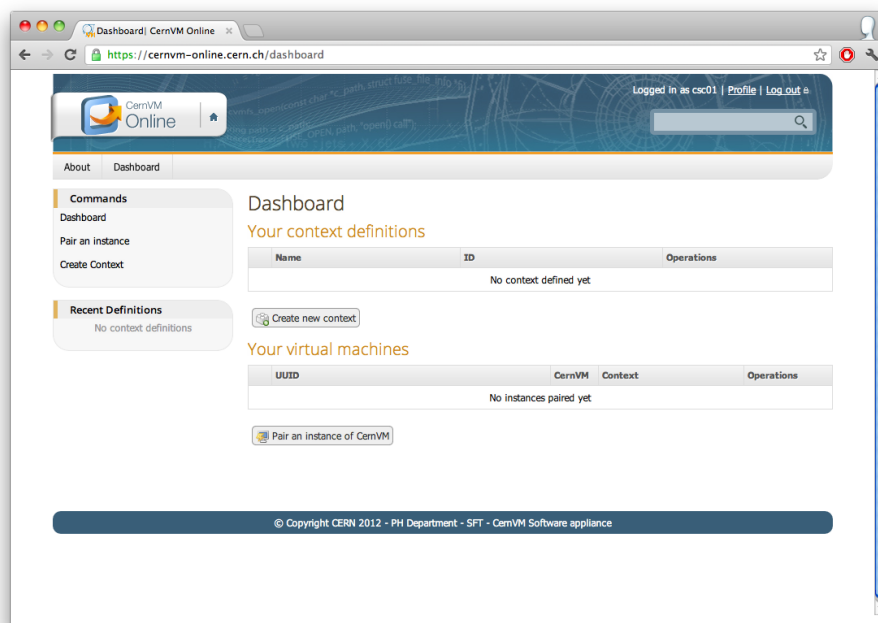
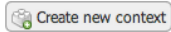

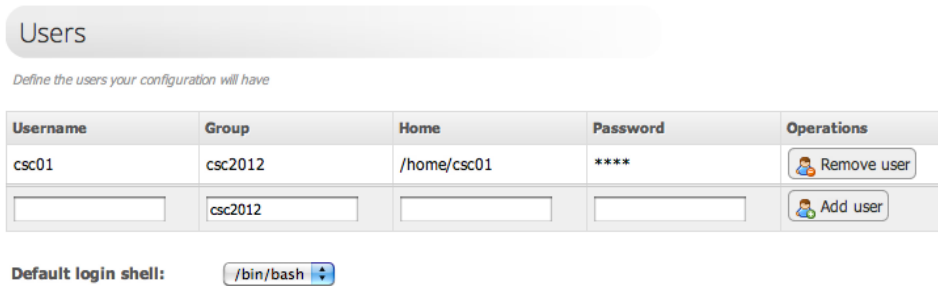
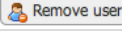
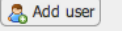


Figure 2: CernVM Online dashboard.

- You need to create a new context, so click on the  button. You will see the *Context template* interface. From here you can specify various customizations to your Virtual Machine. In this exercise we are going to customize a *Desktop* machine that we can login to with pre-defined credentials.
- In the *General* section specify a name for your new context. Name your context **CSC-Desktop**.
- In the *Repository* section select **CSC2012** from the drop-down menu.
- In the *Users* section, add a main user for you virtual machine:

- Use **cernvm** as a user name
- Type **csc2012** in the *Group* field.
- Use **csc2012** as a password
- Click the  button.

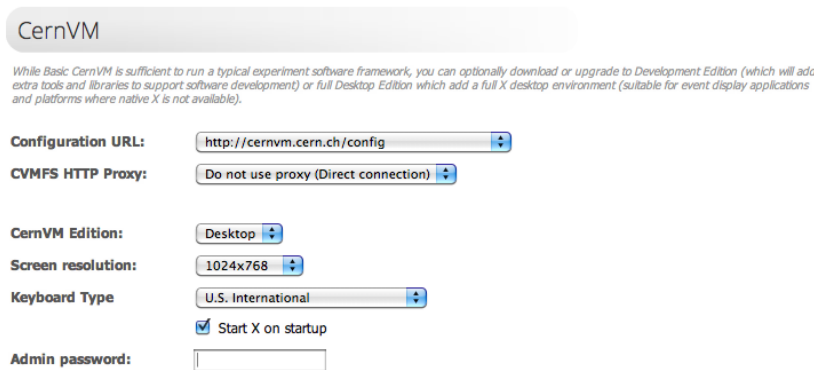


| Username | Group | Home | Password | Operations |
|----------------------|---------|----------------------|----------------------|---|
| csc01 | csc2012 | /home/csc01 | **** |  |
| <input type="text"/> | csc2012 | <input type="text"/> | <input type="text"/> |  |

Default login shell:

Figure 3: You User configuration should look like this.

- Finally, we need a desktop environment. To do so, click on the *CernVM* section.
 - Select **Desktop** from the *CernVM Edition* drop-down menu.
 - Click on the **Start X on startup** checkbox. This will tell CernVM to boot into graphics mode.



Configuration URL:

CVMFS HTTP Proxy:

CernVM Edition:

Screen resolution:

Keyboard Type:

Start X on startup

Admin password:

Figure 4: You CernVM section should look like this.

- Click on the button *Create context* in the bottom of the page.
- Now that you have the context definition in place you will need to tell your virtual machine to use it. You will *pair* the virtual machine with the context. From the side-bar *Commands* (Figure 5) click the **Pair an instance** link.

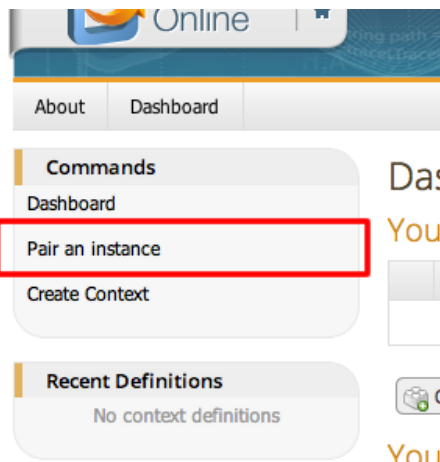



Figure 5: The commands group is located on the left side.

- Find the context definition you made in the previous step and click the  button.
 - You will see a 6-digit pin. With this pin you can pair any CernVM instance.
 - In your virtual machine, press Right-Ctrl+F10 to change the console. Enter the 6-digit pin and press enter to pair the virtual machine.
4. If the configuration went successful, after some time CernVM will automatically switch into graphical mode. Login to your virtual machine with the cernvm user, open a terminal and use the following commands to start the CMS event display

```
. /cvmfs/cms.cern.ch/cmsset_default.sh
cmsrel CMSSW_3_8_0
cd CMSSW_3_8_0
cmsenv
cmsShow
```

Once opened, you can open sample events by clicking File → Open Remote Data File → Fireworks at Cern. Select the ROOT file from the Cosmics sub directory. Close the event display and open it again. Why is it faster the second time?

5. Check again the size of your virtual machine image. As you can see, the image got larger as CernVM-FS has downloaded the necessary files to start the CMS event display. You can check which files have been loaded by running (inside the virtual machine)

```
sudo cvmfs-talk -i cms cache list
```

How many files (approximately) did CernVM-FS load in order to start the CMS event viewer? **Hint:** The `wc -l` utility can count lines.

6. Why is it a very bad idea to recursively list all files under `/cvmfs/cms.cern.ch` (e. g. by tools like `find` or `du`)?

Exercise 2 – Running a Zoo

Now that we have happily running a Desktop CernVM, we also want to use virtual machines to run computations on multiple computers. In order to do so, we will instantiate a couple of “batch node” CernVMs. Conceptually, three artifacts are required to instantiate a batch CernVM on a cloud infrastructure: the virtual machine image, the specification of the virtual hardware, and the so-called *contextualization* specification.

By contextualization, a virtual machine image can be adapted to its runtime environment. You remember the configuration procedure you did in CernVM Online for the previous exercise? This configuration can be given to the virtual machine in the form of a text file containing key-value pairs. The virtual machine can then contextualize itself according to the given information. For instance, the virtual machine can act in different roles (“master node”, “worker node”) or it can be ordered to start specific services such as the Condor scheduler. The advantage of contextualization is that the image itself does not need to change. A single virtual machine can cover many use cases.

Disclaimer: In this exercise, we will use a couple of hand-written scripts to start virtual machines, list running virtual machines, and so on. While the exact way of managing clusters of virtual machines varies from provider to provider (CernVM Online, Cern LxCloud, Amazon EC2, ...), the concepts are the same everywhere. Every cloud provider gives you means to instantiate virtual machines, contextualize them, list them, and destroy them.

1. In folder `cernvm/02zoo` you find three files (`file1`, `file2`, `file3`). Can you figure out, which of them is the virtual machine image, which is the contextualization specification, and which is the virtual hardware specification? **Hint:** You can use the `ls -l` tool again to check for file sizes. You can read text files with the `less` tool.
2. How much RAM will your virtual machine use according to the provided artifacts?
3. Create disk images for your zoo starting from the template `fileX`:

```
cernvm-clone.sh --in=fileX cernvm01.vdi
cernvm-clone.sh --in=fileX cernvm02.vdi
```

4. Use the script `cernvm-start.sh` to start two batch CernVMs. Name them `cernvm-batch[01-03]`, e. g. `cernvm-batch01`. **Hint:** You need the `--user-data`, `--name`, `--disk`, and `--nogui` parameters of the `cernvm-start.sh` script. If anything goes wrong, you can use the `cernvm-destroy.sh` script to start from scratch.
5. Use the script `cernvm-list.sh` to list your CernVMs and to check if they all have booted.
6. Use the script `cernvm-destroy.sh` to shut all the virtual machines down.

Bonus Exercise. You might wonder how you can tell the batch CernVMs what to do. In fact, it is common for batch virtual machines that you cannot login to them (sometimes, such virtual machines are also called “headless”). In this exercise, the batch CernVM is pre-contextualized in a way that it can do something useful on request. Find out, how the batch CernVM connects to the outside and how to submit work to the batch CernVMs. Use your favorite search engine if you like.

Up to this point, we didn’t do anything useful with our CernVM zoo. Stay tuned.

Further Reading

- Kivity et al., *kvm: the Linux Virtual Machine Monitor*, Proc. of the 2007 Linux Symposium
- VirtualBox: <https://www.virtualbox.org>
- CernVM: <http://cernvm.cern.ch>
- Blomer et al., *Distributing LHC application software and conditions databases using the CernVM file system*, Journal of Physics: Conference Series **331**, 2011
- CernVM Online: <http://cernvm-online.cern.ch>
- Thain, Tannenbaum, Livny, *Distributed Computing in Practice: The Condor Experience*, Concurrency – Practice and Experience **17**, 2005