

Sandbox containers

Container basics

For how to use Singularity/Apptainer containers please see our course at: http://dcsrcs-courses.ad.unil.ch/r_python_singularity/r_python_singularity.html

Sandboxes

A container image (the .sif file) is read only and its contents cannot be changed which makes them perfect for distributing safe in the knowledge that they should run exactly as they were created.

Sometimes, especially when developing things, it's very useful to be able to interactively modify a container and this is what sandboxes are for.

Please be aware that anything done by hand is not reproducible so all steps should be transferred to the container definition file.

Creating and modifying a sandbox

Note that the steps here should be run on the cluster login node (curnagl.dcsr.unil.ch) as it is currently the only machine with the configuration in place to allow containers to be built.

To start you need a basic definition file - this can be an empty OS or something more complicated that already has some configuration.

In the following example we will use a definition that installs the latest version of R. We will then try and install extra packages before creating the immutable SIF image.

Here's our file which we save as `newR.def`

```
Bootstrap: docker
From: ubuntu:20.04
```

```

%post
  apt update
  apt install -y locales gnupg-agent
  sed -i '/^#.* en_*.UTF-8 /s/^#//' /etc/locale.gen
  sed -i '/^#.* fr_*.UTF-8 /s/^#//' /etc/locale.gen
  locale-gen

  # install two helper packages we need
  apt install -y --no-install-recommends software-properties-common dirmngr

  # add the signing key (by Michael Rutter) for these repos
  wget -q0- https://cloud.r-project.org/bin/linux/ubuntu/marutter_pubkey.asc | tee -a
/etc/apt/trusted.gpg.d/cran_ubuntu_key.asc

  apt-key adv --keyserver hkp://keyserver.ubuntu.com:80 --recv-keys 51716619E084DAB9

  # add the R 4.0 repo from CRAN -- adjust 'focal' to 'groovy' or 'bionic' as needed
  add-apt-repository "deb https://cloud.r-project.org/bin/linux/ubuntu $(lsb_release -cs)-
cran40/"

  apt install -y --no-install-recommends r-base

```

Create the sandbox

Change to your scratch space /scratch/username and:

```

$ module load singularityce

$ singularity build --fakeroot --sandbox newR newR.def

WARNING: The underlying filesystem on which resides "/scratch/username/myR" won't allow to set
ownership, as a consequence the sandbox could not preserve image's files/directories
ownerships
INFO: Starting build...
Getting image source signatures
Copying blob d7bfe07ed847 [-----] 0.0b / 0.0b
Copying config 2772dfba34 done
..

```

```
..  
..  
Processing triggers for libc-bin (2.31-0ubuntu9.9) ...  
Processing triggers for systemd (245.4-4ubuntu3.17) ...  
Processing triggers for mime-support (3.64ubuntu1) ...  
INFO:   Creating sandbox directory...  
INFO:   Build complete: myR
```

This will create a directory called newR which is the writable container image. Have a look inside and see what there is!

Run and edit the image

Before running the container we need to set up the filesystems that will be visible inside - here we want /users and /scratch to be visible

```
$ export SINGULARITY_BINDPATH="/users,/scratch"  
  
$ mkdir newR/users  
$ mkdir newR/scratch
```

Now we launch the image with an interactive shell

```
$ singularity shell --writable --fakeroot newR/  
  
Singularity>
```

On the command line we can then work interactively with the image.

As we are going to be installing R packages we know that we need some extra tools:

```
Singularity> apt-get install make gcc g++ gfortran
```

Now we can launch R and install some packages

```
Singularity> R
```

```
R version 4.2.1 (2022-06-23) -- "Funny-Looking Kid"
Copyright (C) 2022 The R Foundation for Statistical Computing
Platform: x86_64-pc-linux-gnu (64-bit)
..

> install.packages('tibble')
Installing package into '/usr/local/lib/R/site-library'
(as 'lib' is unspecified)
also installing the dependencies 'glue', 'cli', 'utf8', 'ellipsis', 'fansitools', 'lifecycle',
'magrittr', 'pillar', 'rlang', 'vctrs'

trying URL 'https://cloud.r-project.org/src/contrib/glue_1.6.2.tar.gz'
Content type 'application/x-gzip' length 106510 bytes (104 KB)
=====
downloaded 104 KB

..

..

** testing if installed package can be loaded from final location
** testing if installed package keeps a record of temporary installation path
* DONE (tibble)
```

Keep iterating until things are correct but don't forget to write down all the steps and transfer them to the definition file to allow for future reproducible builds.

Sandbox to SIF

```
$ singularity build --fakeroot R-4.2.1-production.sif newR/
```

You will now have a SIF file that can be used in the normal way

```
$ singularity run R-4.2.1-production.sif R

R version 4.2.1 (2022-06-23) -- "Funny-Looking Kid"
Copyright (C) 2022 The R Foundation for Statistical Computing
Platform: x86_64-pc-linux-gnu (64-bit)
..
```

>

Remember that files on /scratch will be automatically deleted if there isn't enough free space so save your definitions in a git repository and move the SIF images to your project space in /work

Révision #14

Créé 7 juillet 2022 15:14:01 par Ewan Roche

Mis à jour 10 décembre 2024 08:06:40 par Emmanuel Jeanvoine