

Machine Learning Workshop

슈퍼컴퓨팅 본부
대용량데이터허브실

류건모
2017년 09월 07일



Contents

- Machine Importing Schedule
- Introduction of Singularity
- Tutorial

About GPU machine

■ Server Spec

- ✓ Server Product : Dell R730 (2U)
- ✓ CPU : 2x Xeon 2.6GHz 14Core
- ✓ RAM : 384GB
- ✓ GPU : NVIDIA P100
 - Double-Precision : 4.7TFLOPS
 - Single-Precision : 9.3TFLOPS
- ✓ Provisioning and Available Date : End of Oct.



System information

■ Plan of System Setup

- ✓ OS : CentOS7 but changeable. ex) Ubuntu
- ✓ Software
 - CVMFS
 - Singularity (or Docker)
 - CUDA Toolkit
 - Development Packages
- ✓ Singularity (or Docker) Images
 - tensorflow with GPU support
 - MXNet with GPU support
 - EL6 (or SLC6) image for legacy
- ✓ Mount point
 - /cms : CMS Tier3 home and scratch directories
 - /xrootd : CMS Tier3 XRootD Fuse mount point
 - + additional endpoint if a user requests,

Singularity

■ Introduction

- ✓ “Singularity” is a container-based server virtualization tool like as the “Docker”
- ✓ “Singularity” don’t need to additional daemon for super user privilege
 - Instead of the “Docker”, a regular user can use “Singularity” to run the image and to interactive the shell with container environment
 - This is a good attribute for batch system
 - If you already know admin password, you can not escalate to admin in container
- ✓ “Singularity” can use the “Docker” images **directly**
Web Page : <http://singularity.lbl.gov/>



Tutorial

■ Pull the images from hub

✓ singularity pull **docker://ubuntu:latest**

<-> Docker : docker pull ubuntu:latest

✓ To search image is not yet included (docker search)

– Please, search image on docker hub page

```
[geonmo@ui10 simages]$ singularity pull docker://ubuntu:latest
Initializing Singularity image subsystem
Opening image file: ubuntu-latest.img
Creating 225MiB image
Binding image to loop
Creating file system within image
Image is done: ubuntu-latest.img
Docker image path: index.docker.io/library/ubuntu:latest
Cache folder set to /share/geonmo/.singularity/docker
Importing: base Singularity environment
Importing: /share/geonmo/.singularity/docker/sha256:d5c6f90da05dc7e77d2e5fef63c341ab05ba2a03396ab5ae8f18814a7bbf5265.tar.gz
Importing: /share/geonmo/.singularity/docker/sha256:1300883d87d5ae0582a4e0f3f86ad93985b18676e8d829f6f3b4990403fdf145.tar.gz
Importing: /share/geonmo/.singularity/docker/sha256:c220aa3cfc1bff0c1df733eac869a66476ea35136c69d881da1e297ed8d0af2f.tar.gz
Importing: /share/geonmo/.singularity/docker/sha256:2e9398f099dcc5a3e5521205e1f55ee5d7665e203193cf48b32553eabcca323c.tar.gz
Importing: /share/geonmo/.singularity/docker/sha256:dc27a084064fee108f4d486469939bbfa6a5b16f6d6d3c6b5b2475e938f4c71e.tar.gz
Importing: /share/geonmo/.singularity/metadata/sha256:fe44851d529f465f9aa107b32351c8a0a722fc0619a2a7c22b058084fac068a4.tar.gz
Done. Container is at: ubuntu-latest.img
```

Tutorial

- “shell” command : interactive task in container
 - ✓ singularity shell --home \$HOME:/home/\$USER --pwd /home/\$USER --bind /cvmfs --scratch /var/tmp --scratch /tmp -shell /bin/bash --pid ubuntu-latest.img (or directory)
 - Options
 - --home
 - » it is necessary that if your home dir is not under /home
 - --pwd
 - » Changing current dir after login to container
 - --bind
 - » Bind the directories to container. The directory must be existed in image
 - --ipc
 - » IPC Isolation (Pipe X)
 - --pid
 - » Process Isolation (/bin/bash’s PID=>1)
 - --contain
 - » Disconnect /dev, \$HOME, /tmp dirs between host and container
 - --writable
 - » Read/Write mode to change container contents

Tutorial

■ “shell” command

- ✓ singularity shell –help : Display full help message for shell command
- ✓ “—writeable” option is meaningless if you do not run the contain as root

```
[geonmo@ui10 simages]$ singularity shell --home $HOME:/home/$USER --pwd /home/$USER --bind /cvmfs --scratch /var/tmp --scratch /tmp --shell /bin/bash --ip
c --pid -w ubuntu-latest.img
WARNING: Skipping user bind, non existant bind point (directory) in container: '/cvmfs'
Singularity: Invoking an interactive shell within container...

geonmo@ui10:~$ cat /etc/lsb-release
DISTRIB_ID=Ubuntu
DISTRIB_RELEASE=16.04
DISTRIB_CODENAME=xenial
DISTRIB_DESCRIPTION="Ubuntu 16.04.3 LTS"
geonmo@ui10:~$ apt search gimp
Sorting... Done
Full Text Search... Done
libgtk2-perl/xenial-updates 2:1.2498-1 amd64
  Perl interface to the 2.x series of the Gimp Toolkit library

geonmo@ui10:~$ id
uid=556800422(geonmo) gid=556800422(geonmo) groups=556800422(geonmo),556800000(admins),556800004(foreman_users),556800020(cms_admins),556800023(jira_users
),556800025(stash_users),556800026(confluence_users),556800038(hcp_users),556800039(hcp_admins),556800061(cms_managers),556800062(cms_users),556800065(hcp
_managers),556800068(infra_managers)
geonmo@ui10:~$ ps
PID TTY          TIME CMD
  1 pts/23      00:00:00 bash
 10 pts/23      00:00:00 ps
geonmo@ui10:~$
```

```
geonmo@ui10:~$ apt install libgtk2-perl
E: Could not open lock file /var/lib/dpkg/lock - open (13: Permission denied)
E: Unable to lock the administration directory (/var/lib/dpkg/), are you root?
```


Tutorial

■ Create images

- ✓ `singularity create -s 3000 testimage.img`
=> create about 3000MByte testimage.img which is blank.

■ Import container

- ✓ `singularity import testimage.img docker://ubuntu:latest`
=> Fill the contents from docker image

■ Expand images

- ✓ `singularity expand -s 6000 testimage.img`
=> Expand the image's size to 6000MB

Tutorial

■ Bootstrap (admin only)

- ✓ singularity bootstrap <dir path> <recipe file>
- ✓ egs) singularity bootstrap ubuntu-latest-test Singularity.txt
- ✓ "Docker" is better than "Singularity" during bootstrapping

■ An Example of recipe file (Singularity.txt)

```
Bootstrap: docker
From: ubuntu:latest
%runscript  ## only use "singularity run~~"
exec echo "The runscript is the containers default runtime command!"
%files      ## copy files from host to container
%environment
VARIABLE=MEATBALLVALUE
export VARIABLE
%labels
AUTHOR your@email.address
%post      ## Build command to your own container. Same as bash script.
apt-get update && apt-get -y install python3 git wget
mkdir /cvmfs
echo "The post section is where you can install, and configure your container."
```

Final Tutorial

■ Simple tensorflow

- ✓ Login to ui10,
- ✓ run "s_tensorflow"
- ✓ bash
- ✓ mkdir test_ml
- ✓ cd test_ml
- ✓ git clone <https://github.com/aymericdamien/TensorFlow-Examples.git>
- ✓ cd TensorFlow-Examples
- ✓ python examples/3_NeuralNetworks/neural_network.py

Conclusion

- GPU machine will be arrived on October
- “Singularity” is a very nice tool to setup the machine learning software
- “Singularity” can use the “Docker” images directly
- Don’t use “singularity” to make your container at this time. Use “Docker”.

감사합니다.