udocker – executing docker containers with GPU applications in non privileged environments

Highlight and status update

Mario David – david@lip.pt

On behalf of DEEP-HybridDatacloud

Main developer Jorge Gomes - jorge@lip.pt

EGI Conference 2019
May 6-8, 2019
Amsterdam

DEEP-Hybrid-DataCloud is funded by the Horizon 2020 Framework Programme of the European union under grant agreement number 777435

Introduction

- Run applications encapsulated in docker containers.
  - without using docker
  - without root privileges
  - without system administrators intervention
  - without additional system software
- Run as:
  - normal user
  - with normal process controls and accounting
  - in interactive or batch systems

Open source, paper in Open Access:
https://github.com/indigo-dc/udocker/
https://doi.org/10.1016/j.cpc.2018.05.021

Improved support for GPU applications

UDOCKER Hybrid DataCloud

- In the framework of DEEP-HybridDatacloud:
  - Automatic use of GPUs: implement code that allows the use of GPUs without the user needing to install and match the NVIDIA drivers in the image or containers.
  - Done in DEEP-Genesis release: udocker 1.1.3

- **Before** (same requirement as if running with docker):
  - Image has same NVIDIA driver and libraries version as host system .
  - User has to build his/hers images, know the host(s) NVIDIA version.
  - Several similar images with the different NVIDIA versions if want to run in multiple systems.
    - **udocker run <my-nvidia-image> <my-application>**
  - Still simpler than running with docker:
    - **docker run --device=/dev/nvidia0:/dev/nvidia0**
      **--device=/dev/nvidiactl:/dev/nvidiactl **
      --device=/dev/nvidia-uvm:/dev/nvidia-uvm <my-nvidia-image> <my-application>

- And nvidia "invented": nvidia-docker:
  - An execution engine for docker that uses/abstracts the underlying host drivers and libraries.
  - Obviously, uses oficial nvidia docker images (CUDA).

Improved support for GPU applications

- In the framework of DEEP-HybridDatacloud:
  - Automatic use of GPUs: implement code that allows the use of GPUs without the user needing to install and match the NVIDIA drivers in the image or containers.
  - Done in DEEP-Genesis release: udocker 1.1.3

- **After**:
  - Images from official nvidia dockerhub: https://hub.docker.com/r/nvidia/cuda
    - as well as others: https://hub.docker.com/r/tensorflow/tensorflow
  - User builds his/her application with the above base images.
  - No more need to know underlying host nvidia driver or library versions.
  - New option in the "setup" CLI of udocker:
    - `udocker setup --execmode=<execution-mode> **--nvidia** <my-container>`
    - `udocker run <my-container> <my-application>`
  - Automagic?
    - Search and copy the host nvidia libraries into correct paths inside the container.
  - Obtain same functionality as nvidia docker execution engine and libraries.

Things of udocker

- Dependencies: only python standard library - no requirements
  - no need to install any additional packages.
    - user may not have them in “all” batch clusters, HPC systems, etc.
    - user level pip install may not be enough, sometimes need to install additional OS packages.
  - no need to have pip, nor pyenv (virtualenv).
- Support - Python 2.6 (Centos6), Python 2.7 (Centos7, Ubuntu16.04).
  - we known Python2 is EOL by the end of the year but.
  - still exists batch clusters, HPC systems, etc. in Centos6 and Centos7.
  - we know that some (or most) support loading Python3.
- Big Python script - 7000 LoC
  - user just downloads/wget/curl the script and runs it.
  - (unit tests single 7700 LoC)
- Continuous Delivery: Jenkins pipeline now include publication to Pypi

udocker: Work in Progress

- Git fork: https://github.com/mariojmdavid/udocker/tree/devel3
  - udocker script and unit tests - breakdown into modules/classes.
  - higher maintainability.
  - several modules already compatible with Python3.
  - Support Ubuntu 18.04.
  - but try not to break Python 2.6 and 2.7
  - Distribution:
    - tarball: download and untar and ready to use
    - Pypi: pip install or git clone and python setup install

udocker: New features

- Better support and automatization for MPI.
  - full automatization may not be possible, as was implemented for GPUs.
  - even if host drivers and libraries (infiniband and mpi flavor) are available inside the container, applications will need to be compiled with them.
- There was a request from charliecloud to implement it as a new execution engine in udocker:
  - evaluation and decision to be done.

Thank You

Mario David – david@lip.pt