
DEEP-Hybrid-DataCloud documentation Documentation

Release DEEP-2 (XXX)

DEEP-Hybrid-DataCloud consortium

Feb 07, 2020

Contents

1	User documentation	1
1.1	User documentation	1
1.1.1	Quickstart Guide	1
1.1.2	Overview	4
1.1.3	HowTo's	15
1.1.4	Modules	34
2	Component documentation	37
3	Technical documentation	39
3.1	Technical documentation	39
3.1.1	Mesos	39
3.1.2	Kubernetes	51
3.1.3	OpenStack nova-lxd	67
3.1.4	uDocker	76
3.1.5	Miscellaneous	79
4	Indices and tables	85

If you are a user (current or potential) you should start here.

1.1 User documentation

If you want a quickstart guide, please check the following link.

1.1.1 Quickstart Guide

1. go to [DEEP Marketplace](#)
2. [Browse](#) available modules
3. Find the module you are interested in and get it

Let's explore what we can do with it!

Run a module locally

Requirements

- [docker](#)
- If GPU support is needed:
 - you can install [nvidia-docker](#) along with docker, OR
 - install [udocker](#) instead of docker. [udocker](#) is entirely a user tool, i.e. it can be installed and used without any root privileges, e.g. in a user environment at HPC cluster.

N.B.: Starting from version 19.03 docker supports NVIDIA GPUs, i.e. no need for nvidia-docker (see [Release notes](#) and [moby/moby#38828](#))

Run the container

Run the Docker container directly from Docker Hub:

Via docker command:

```
$ docker run -ti -p 5000:5000 -p 6006:6006 deephdc/deep-oc-module_of_interest
```

Via udocker:

```
$ udocker run -p 5000:5000 -p 6006:6006 deephdc/deep-oc-module_of_interest
```

With GPU support:

```
$ nvidia-docker run -ti -p 5000:5000 -p 6006:6006 deephdc/deep-oc-module_of_
↪interest
```

If docker version is 19.03 or above:

```
$ docker run -ti --gpus all -p 5000:5000 -p 6006:6006 deephdc/deep-oc-module_
↪of_interest
```

Via udocker with GPU support:

```
$ udocker pull deephdc/deep-oc-module_of_interest
$ udocker create --name=module_of_interest deephdc/deep-oc-module_of_interest
$ udocker setup --nvidia module_of_interest
$ udocker run -p 5000:5000 -p 6006:6006 module_of_interest
```

Access the module via API

To access the downloaded module via *the DEEPaaS API*, direct your web browser to <http://0.0.0.0:5000/ui>. If you are training a model, you can go to <http://0.0.0.0:6006> to monitor the training progress (if such monitoring is available for the model).

For more details on particular models, please read *the module's* documentation.

models		
GET	/v2/models/	Return loaded models and its information
GET	/v2/models/imgclas/	Return model's metadata
POST	/v2/models/imgclas/train/	Retrain model with available data
GET	/v2/models/imgclas/train/	Get a list of trainings (running or completed)
GET	/v2/models/imgclas/train/{uuid}	Get status of a training
DELETE	/v2/models/imgclas/train/{uuid}	Cancel a running training
POST	/v2/models/imgclas/predict/	Make a prediction given the input data

Related HowTo's:

- *How to perform inference locally*
- *How to perform training locally*

Train a module on DEEP Pilot Infrastructure**Requirements**

- [DEEP-IAM](#) registration

Sometimes running a module locally is not enough as one may need more powerful computing resources (like GPUs) in order to train a module faster. You may request [DEEP-IAM](#) registration and then use the DEEP Pilot Infrastructure to deploy a module. For that you can use the [DEEP Dashboard](#). There you select a module you want to run and the computing resources you need. Once you have your module deployed, you will be able to train the module and view the training history:

The screenshot shows the DEEP Hybrid-DataCloud dashboard. At the top, there's a navigation bar with links for TRAINING DASHBOARD, DEPLOYMENTS, ADVANCED, and EXTERNAL LINKS, along with a user profile for IGNACIO HEREDIA and the DEEP logo. Below the navigation bar, there's a header with the ID 11ea3ddc-0dd2-7ff3-a966-0242eca7ffec. The main content area is titled 'mods' and contains a table of training history. The table has columns for Training uuid, Date, Status, Message, and Actions. The status column shows 'error', 'done', and 'running' for different training runs. The Actions column contains a 'Delete' button for each row. Below the table, there's a 'Training history' label.

Training uuid	Date	Status	Message	Actions
4473224736a24c869fce3da38c03f288	2020-01-28 12:19:31.943220	error	This is a module error.	Delete
55dc244cee04204a1dd8e79e59190a	2020-01-28 12:20:46.987462	done		Delete
691d7d5b52744c4bac8b5eff8de5b3	2020-01-28 12:08:00.437706	error	This is a module error.	Delete
b6686ab3de104563b0500e1a0caa4e3f	2020-01-28 12:22:52.640658	done		Delete
df2d13e3da13420496fbc3bd022c7f89	2020-01-28 12:22:58.749524	done		Delete
26a906a603a047859b5ad5aa79c35ca7	2020-01-28 12:23:07.996410	done		Delete
ac4dac0c06854913a964a97126f29163	2020-01-28 12:23:36.010563	done		Delete
0e012dea06f948f9b094b980d022996	2020-01-28 12:23:36.808146	running		Delete

Related HowTo's:

- [How to train a model remotely](#)
- [How to deploy with CLI](#)

Develop and share your own module

The best way to develop a module is to start from *the DEEP Data Science template*. It will create a project structure and files necessary for an easy *integration with the DEEPaaS API*. The *DEEPaaS API* enables a user-friendly interaction with the underlying Deep Learning modules and can be used both for training models and doing inference with the services. The integration with the API is based on the definition of endpoints to the model and the creation of standard API methods (eg. train, predict, etc).

Related HowTo's:

- [How to use the DEEP Data Science template for model development](#)
- [How to develop your own machine learning model](#)
- [How to integrate your model with the DEEPaaS API](#)
- [How to add your model to the DEEP Marketplace](#)

A more in depth documentation, with detailed description on the architecture and components is provided in the following sections.

1.1.2 Overview

Architecture overview

There are several different components in the DEEP-HybridDataCloud project that are relevant for the users. Later on you will see how each *different type of user* can take advantage of the different components.

The marketplace

The DEEP Marketplace is a catalogue of modules that every user can have access to. Modules can be:

- **Trained:** Those are modules that an user can train on their own data to create a new service. Like training an [image classifier](#) with a plants dataset to create a [plant classifier](#) service. Look for the `trainable` tag in the marketplace to find those modules.
- **Deployed for inference:** Those are modules that have been pre-trained for a specific task (like the [plant classifier](#) mentioned earlier). Look for the `inference` and `pre-trained` tags in the marketplace to find those modules.

Some modules can both be trained and deployed for inference. For example the [image classifier](#) can be trained to create other image classifiers but can also be deployed for inference as it comes pre-trained with a general image classifier.

For more information have a look at the [marketplace](#).

The API

The [DEEPaaS API](#) is a key component for making the modules accessible to everybody (including non-experts), as it provides a consistent and easy to use way to access the model's functionality. It is available for both inference and training.

Advanced users that want to create new modules can make them *compatible with the API* to make them available to the whole community. This can be easily done, since it only requires minor changes in user's code and adding additional entry points.

For more information take a look on the full [API guide](#).

The data storage resources

Storage is essential for user that want to create new services by training modules on their custom data. For the moment we support hosting data in [DEEP-Nextcloud](#). We are currently working on adding additional storage support, as well as more advanced features such as data caching (see [OneData](#)), in cooperation with the [eXtreme-DataCloud](#) project.

The dashboards

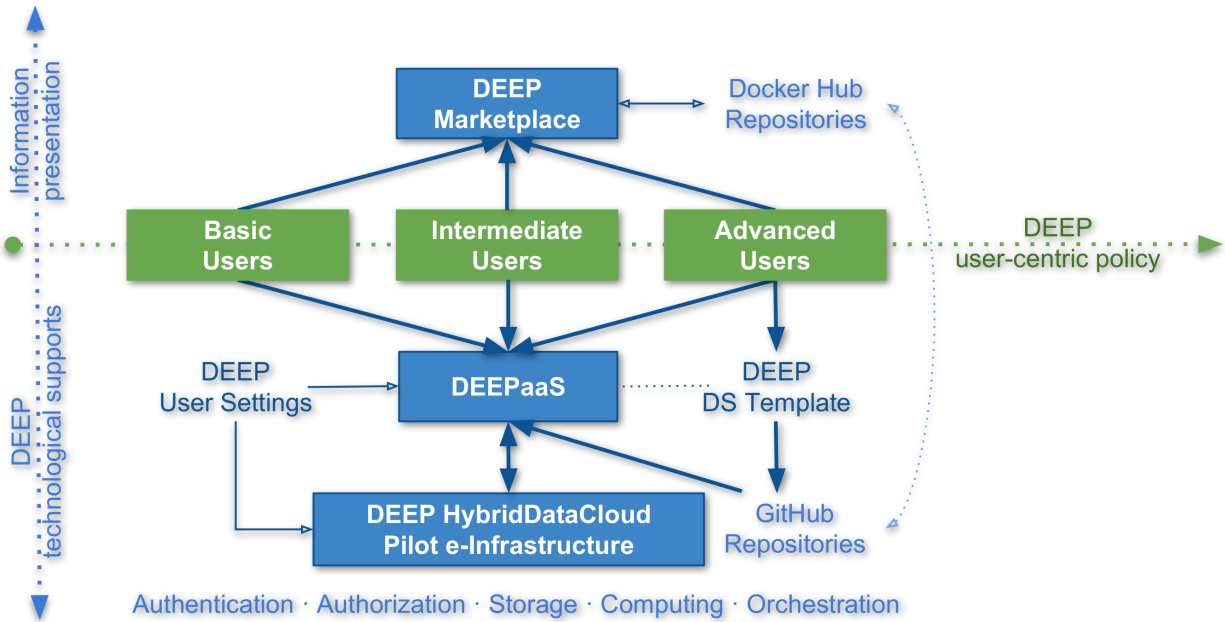
The DEEP dashboard allow users to access computing resources to deploy, perform inference, and train their modules. To be able to access the Dashboard you need [IAM credentials](#). There are two versions of the Dashboard:

- **Training dashboard** This dashboard allows you to interact with the modules hosted at the [DEEP Open Catalog](#), as well as deploying external Docker images hosted in Dockerhub. It simplifies the deployment and hides some of the technical parts that most users do not need to worry about. *Most of DEEP users would use this dashboard.*
- **General purpose dashboard** This dashboard allows you to interact with the underling [TOSCA templates](#) (which configure the job requirements) instead of modules and deploy more complex topologies (e.g. a kubernetes cluster). Modules can either use a [general template](#) or create a dedicated one based on the [existing templates](#).

For more information take a look on the full [Dashboard guide](#).

Our different user roles

The DEEP-HybridDataCloud project is focused on three different types of users, depending on what you want to achieve you should enter into one or more of the following categories:



The basic user

This user wants to use modules that are already pre-trained and *test them with their data*, and therefore don't need to have any machine learning knowledge. For example, they can take an already trained module for *plant classification* that has been containerized, and use it to classify their own plant images.

What DEEP can offer to you:

- a *catalogue* full of ready-to-use modules to perform inference with your data
- an *API* to easily interact with the services
- solutions to run the inference in the Cloud or in your local resources
- the ability to develop complex topologies by composing different modules

Related HowTo's:

- *How to perform inference locally*

The intermediate user

The intermediate user wants to *retrain an available module* to perform the same task but fine tuning it to their own data. They still might not need high level knowledge on modelling of machine learning problems, but typically do need basic programming skills to prepare their own data into the appropriate format. Nevertheless, they can re-use the knowledge being captured in a trained network and adjust the network to their problem at hand by re-training the network on their own dataset. An example could be a user who takes the generic *image classifier* model and retrains it to perform *seed classification*.

What DEEP can offer to you:

- the ability to train out-of-the-box a module of the *catalogue* on your personal dataset
- an *API* to easily interact with the model
- data *storage resources* to access your dataset (DEEP-Nextcloud, OneData, ...)

- the ability to deploy the developed service on Cloud resources
- the ability to share the service with other users in the user's catalogue

Related HowTo's:

- [How to train a model locally](#)
- [How to train a model remotely](#)
- [How to add your module to the DEEP Marketplace](#)

The advanced user

The advanced users are the ones that will *develop their own machine learning models* and therefore need to be competent in machine learning. This would be the case for example if we provided an image classification model but the users wanted to perform object localization, which is a fundamentally different task. Therefore they will design their own neural network architecture, potentially re-using parts of the code from other models.

What DEEP can offer to you:

- a ready-to-use environment with the main DL frameworks running in a dockerized solution running on different types of hardware (CPUs, GPUs, etc)
- data *storage resources* to access your dataset ([DEEP-Nextcloud](#), [OneData](#), ...)
- the ability to deploy the developed module on Cloud resources
- the ability to share the module with other users in the open *catalogue*
- the possibility to *integrate your module with the API* to enable easier user interaction

Related HowTo's:

- [How to use the DEEP Cookiecutter template for model development](#)
- [How to develop your own machine learning model](#)
- [How to integrate your model with the DEEPaaS API](#)
- [How to add your model to the DEEP Marketplace](#)

DEEP Data Science template

To simplify the development and in an easy way integrate your model with the [DEEPaaS API](#), a project template, `cookiecutter-data-science`⁰, is provided in our GitHub.

In order to create your project based on the template, one has to [install](#) and then run `cookiecutter` tool as follows:

```
$ cookiecutter https://github.com/indigo-dc/cookiecutter-data-science
```

You are first provided with [Info] line about the parameter and in the next line you configure this parameter. You will be asked to configure:

- Remote URL to host your new repositories (git), e.g. <https://github.com/deephdc>, `git_base_url`
- Project name, `project_name`
- Name of your new repository, to be added after “`git_base_url`” (see above), `repo_name` (aka `<your_project>` in the following)

⁰ The more general `cookiecutter-data-science` template was adapted for the purpose of DEEP.

- Author name(s) (and/or your organization/company/team). If many, separate by comma, `author_name`
- E-Mail(s) of main author(s) (or contact person). If many, separate by comma, `author_email`
- Short description of the project, `description`
- Application version (expects X.Y.Z (Major.Minor.Patch)), `app_version`
- Choose open source license, default is MIT. For more info: <https://opensource.org/licenses>, `open_source_license`
- User account at hub.docker.com, e.g. 'deephdc' in <https://hub.docker.com/u/deephdc>, `dockerhub_user`
- Docker image your Dockerfile starts from (FROM <docker_baseimage>) (don't provide the tag here), e.g. tensorflow/tensorflow, `docker_baseimage`
- CPU tag for the baseimage, e.g. 1.14.0-py3. Has to match python3!, `baseimage_cpu_tag`
- GPU tag for the baseimage, e.g. 1.14.0-gpu-py3. Has to match python3!, `baseimage_gpu_tag`

Note: These parameters are defined in `cookiecutter.json` in the [cookiecutter-data-science](#) source.

When these questions are answered, following two repositories will be created locally and immediately linked to your `git_base_url`:

```
~/DEEP-OC-your_project
~/your_project
```

each repository has two branches: 'master' and 'test'.

<your_project> repo

Main repository to integrate model with the following structure:

— data	Placeholder for the data
— raw	The original, immutable data dump.
— docs	Documentation on the project; see sphinx-doc.org for
↪details	
— models	Trained and serialized models, model predictions, or model
↪summaries	
— notebooks	Jupyter notebooks. Naming convention is a number (for
↪ordering),	
	the creator's initials (if many user development),
	and a short `_` delimited description,
	e.g. `1.0-jqp-initial_data_exploration.ipynb`.
— references	Data dictionaries, manuals, and all other explanatory
↪materials.	
— reports	Generated analysis as HTML, PDF, LaTeX, etc.
— your_project	Main source code of the project
— __init__.py	Makes <code>your_project</code> a Python module

(continues on next page)

(continued from previous page)

dataset	Scripts to download and manipulate raw data
└─ make_dataset.py	
features	Scripts to prepare raw data into features for modeling
└─ build_features.py	
models	Scripts to train models and then use trained models to
→make predictions	
└─ deep_api.py	Main script for the integration with DEEP API
tests	Scripts to perform code testing
visualization	Scripts to create exploratory and results oriented
→visualizations	
└─ visualize.py	
.dockerignore	Describes what files and directories to exclude for
→building a Docker image	
.gitignore	Specifies intentionally untracked files that Git should
→ignore	
Jenkinsfile	Describes basic Jenkins CI/CD pipeline
LICENSE	License file
README.md	The top-level README for developers using this project.
requirements.txt	The requirements file for reproducing the analysis
→environment,	
	e.g. generated with `pip freeze > requirements.txt`
setup.cfg	makes project pip installable (pip install -e .)
setup.py	makes project pip installable (pip install -e .)
test-requirements.txt	The requirements file for the test environment
tox.ini	tox file with settings for running tox; see tox.testrun.org

Certain files, e.g. README.md, Jenkinsfile, setup.cfg, tox.ini, etc are pre-populated based on the answers you provided during cookiecutter call (see above).

<DEEP-OC-your_project>

Repository for the integration of the *DEEPaaS API* and your_project in one Docker image.

Dockerfile	Describes main steps on integrationg DEEPaaS API and your_project application in one Docker image
Jenkinsfile	Describes basic Jenkins CI/CD pipeline
LICENSE	License file

(continues on next page)

(continued from previous page)

```
└─ README.md          README for developers and users.
└─ docker-compose.yml  Allows running the application with various configurations,
↳via docker-compose
└─ metadata.json       Defines information propagated to the DEEP Open Catalog,
↳https://marketplace.deep-hybrid-datacloud.eu
```

All files get filled with the info provided during cookiecutter execution (see above).

Step-by-step guide

1. (if not yet done) install cookiecutter, as e.g. `pip install cookiecutter`
2. run `cookiecutter https://github.com/indigo-dc/cookiecutter-data-science`
3. answer all the questions, pay attention about docker tags!
4. two directories will be created: `<your_project>` and `<DEEP-OC-your_project>` (each with two git branches: master and test)
5. go to `github.com/user_account` and create corresponding repositories `<your_project>` and `<DEEP-OC-your_project>`
6. go to your terminal, `<your_project>`, `git push origin --all`
7. go to your terminal, `<DEEP-OC-your_project>`, `git push origin --all`
8. your github repositories are now updated with initial commits
9. you can build `<deep-oc-your_project>` Docker image locally: go to `<DEEP-OC-your_project>` directory, do `docker build -t dockerhubuser/deep-oc-your_project .`
10. you can now run deepaas as `docker run -p 5000:5000 dockerhubuser/deep-oc-your_project`

DEEPaaS API

The [DEEPaaS API](#) enables a user friendly interaction with the underlying Deep Learning modules and can be used both for training models and doing inference with services.

For a detailed up-to-date documentation please refer to the [official DEEPaaS documentation](#).

Integrate your model with the API

To make your Deep Learning model compatible with the DEEPaaS API you have to:

1. Define the API methods for your model

Create a Python file (named for example `deep_api.py`) inside your package. In this file you can define any of the [API methods](#). You don't need to define all the methods, just the ones you need. Every other method will return a `NotImplementError` when queried from the API. For example:

- **Enable prediction:** implement `get_predict_args` and `predict`.
- **Enable training:** implement `get_train_args` and `train`.
- **Enable model weights preloading:** implement `warm`.
- **Enable model info:** implement `get_metadata`.

Here is [an example of the implementation](#) of the methods. You can also browse our [github repository](#) for more examples.

2. Define the entrypoints to your model

You must define the entrypoints pointing to this file in the `setup.cfg` as following:

```
[entry_points]
deepaas.v2.model =
    pkg_name = pkg_name.deep_api
```

Here is an [example](#) of the entrypoint definition in the `setup.cfg` file.

Tip: When developing a model with the *DEEP Data Science template*, the Python file with the API methods will automatically be created at `pkg_name/models/deep_api.py`, as well as the entrypoints pointing to it. This path can of course be modified.

Running the API

To start the API run:

```
deepaas-run --listen-ip 0.0.0.0
```

and go to <http://0.0.0.0:5000/ui>. You will see a nice UI with all the methods:

models		
GET	/v2/models/	Return loaded models and its information
GET	/v2/models/imgclas/	Return model's metadata
POST	/v2/models/imgclas/train/	Retrain model with available data
GET	/v2/models/imgclas/train/	Get a list of trainings (running or completed)
GET	/v2/models/imgclas/train/{uuid}	Get status of a training
DELETE	/v2/models/imgclas/train/{uuid}	Cancel a running training
POST	/v2/models/imgclas/predict/	Make a prediction given the input data

DEEP Dashboard



The DEEP dashboard allow users to access computing resources to deploy, perform inference, and train their modules. To be able to access the Dashboard you need [IAM credentials](#). There are two versions of the Dashboard:


- **Training dashboard** This dashboard allows you to interact with the modules hosted at the [DEEP Open Catalog](#), as well as deploying external Docker images hosted in Dockerhub. It simplifies the deployment and hides some of the technical parts that most users do not need to worry about. *Most of DEEP users would use this dashboard.*
- **General purpose dashboard** This dashboard allows you to interact with the underling [TOSCA templates](#) (which configure the job requirements) instead of modules and deploy more complex topologies (e.g. a kubernetes cluster). Modules can either use a [general template](#) or create a dedicated one based on the existing templates.

We will present the features of the Training Dashboard as it is the one that most users will interact with.

Selecting the modules

Once you log into the Dashboard, you can browse through all the Marketplace modules of you can deploy. It also possible to deploy external docker images that are not hosted in the [deepcdc organization](#) (that is from modules not available in the Marketplace) at “Run your own module”, as well as deploying a “Development module” to develop some code.

[TRAINING DASHBOARD](#)
[DEPLOYMENTS](#)
[ADVANCED ▾](#)
[EXTERNAL LINKS ▾](#)
 IGNACIO HEREDIA ▾


 Search...

Dogs breed detector

Identify a dogs breed on the image (133 known breeds)

[Train module](#)
[Marketplace](#)

DEEP OC Massive Online Data Streams

Deep learning for proactive network monitoring and security protection.

[Train module](#)
[Marketplace](#)

DEEP OC Retinopathy Test

A Tensorflow model to classify Retinopathy.

[Train module](#)
[Marketplace](#)

Train an image classifier

Train your own image classifier with your custom dataset. It comes also pretrained on the 1K ImageNet classes.

[Train module](#)
[Marketplace](#)

Plants species classifier

Classify plant images among 10K species from the iNaturalist dataset.

[Train module](#)
[Marketplace](#)

Conus species classifier

Classify conus images among 70 species.

[Train module](#)
[Marketplace](#)

Phytoplankton species classifier

Classify phytoplankton images among 60 classes.

[Train module](#)
[Marketplace](#)

Seed species classifier

Classify seeds images among 700K species.

[Train module](#)
[Marketplace](#)

Upscale multispectral satellites images

Upscale (superresolve) low resolution bands to high resolution in multispectral satellite imagery.

[Train module](#)
[Marketplace](#)

Making a deployment

Once you choose the module to deploy click in Train module. This allow a user to select:

- **The computing resources** to have available. A user can select multiple CPUs and GPUs, the machine RAM as well as optionally choosing the physical site where the machine must be deployed.
- **The service** to run. Currently, options include running the *DEEPaaS API* (recommended for fully developed modules than only need to be trained) and *JupyterLab* (recommended for developing code as well for cases where access to the bash console is needed).

TRAINING DASHBOARD
DEPLOYMENTS
ADVANCED
EXTERNAL LINKS
IGNACIO HEREDIA
deap

Module: deep-oc-image-classification-tf

Template
default
Hardware
CPU
Run
DEEPaaS

Configuration
Advanced

docker_image

deepfdc/deep-oc-image-classification-tf:cpu

docker image from Docker Hub to deploy

docker_privileged

False

equivalent of --privileged docker flag

Click **Submit** and you will be redirected to the page listing all the current deployments.

Note: It is also possible to make deployments via the *command line interface*

Managing the deployments

In the Deployments tab you have a view of all the deployments you have made:

TRAINING DASHBOARD
DEPLOYMENTS
ADVANCED
EXTERNAL LINKS
IGNACIO HEREDIA
deap

My deployments

Refresh New deployment

Show 10 entries
Search:

Deployment uid	Status	Container name	GPUs	Creation time	Deployed At	Actions
11ea473a-53b3-a9e5-a966-0242eca7ffec	CREATE COMPLETE	deepfdc/deep-oc-image-classification-tf:cpu	0	2020-02-04T10:37+0000	RECAS-BARI	Access Info Delete
11ea4742-7e3b-7259-a966-0242eca7ffec	CREATE COMPLETE	deepfdc/deep-oc-dogs_breed_det:cpu	0	2020-02-04T11:35+0000	RECAS-BARI	Access Info Delete
11ea4744-76fa-d64d-a966-0242eca7ffec	CREATE COMPLETE			2020-02-04T11:49+0000	IFCA-LCG2	Access Info Delete
11ea4744-0ceb-59a6-a966-0242eca7ffec	CREATE IN PROGRESS			2020-02-04T11:46+0000	RECAS-BARI	Access Info Delete
11ea4743-e5c3-648f-a966-0242eca7ffec	DELETE IN PROGRESS			2020-02-04T11:45+0000	RECAS-BARI	Access Info Delete

Showing 1 to 5 of 5 entries

Previous 1 Next

If you click on the UUID of the deployments you can access the deployments personal page where you can view things such as the training history, as long as your deployment is using DEEPaaS V2:

11ea3ddc-0dd2-7ff3-a966-0242eca7ffec

mods Train model Predict

Training uuid	Date	Status	Message	Actions
4473224736a24c869fce3da38c03f288	2020-01-28 12:19:31.943220	error	This is a module error.	Delete
55dc244cee04204a1dd8e79e59190a	2020-01-28 12:20:46.987462	done		Delete
691d7d55b52744c4bac8b5eff8de5b3	2020-01-28 12:08:00.437706	error	This is a module error.	Delete
b6686ab3de104563b0500e1a0caa4e3f	2020-01-28 12:22:52.640658	done		Delete
df2d13e3da13420496fbc3bd022c7f89	2020-01-28 12:22:58.749524	done		Delete
26a906a603a047859b5ad5aa79c35ca7	2020-01-28 12:23:07.996410	done		Delete
ac4dac0c06854913a964a97126f29163	2020-01-28 12:23:36.010563	done		Delete
0e012dea06f948f9b094b4b98d0d22996	2020-01-28 12:23:36.808146	running		Delete

Training history

A set of various HowTo examples.

1.1.3 HowTo's

Develop a model using DEEP Data Science template

1. Prepare DEEP Data Science environment

Install cookiecutter (if not yet done)

```
$ pip install cookiecutter
```

Run the DEEP DS cookiecutter template

```
$ cookiecutter https://github.com/indigo-dc/cookiecutter-data-science
```

Answer all questions from *Data Science template* with attentions to `repo_name` i.e. the name of your github repositories, etc. This creates two project directories:

```
~/DEEP-OC-your_project
~/your_project
```

Go to `github.com/your_account` and create corresponding repositories: `DEEP-OC-your_project` and `your_project`. Do `git push origin --all` in both created directories. This puts your initial code to github.

2. Improve the initial code of the model

The structure of `your_project` created using *Data Science template* contains the following core items needed to integrate your Machine learning model with DEEPaaS API:

```
requirements.txt
test-requirements.txt
data/
models/
```

(continues on next page)

(continued from previous page)

```
{{repo_name}}/dataset/make_dataset.py
{{repo_name}}/features/build_features.py
{{repo_name}}/models/deep_api.py
{{repo_name}}/tests/test_unit_model.py
```

2.1 Installing project requirements

Modify `requirements.txt` according to your needs (e.g. add more libraries) then run

```
$ pip install -r requirements.txt
```

You can modify and add more `source` files and put them accordingly into the directory structure.

2.2 Make datasets

Source files in this directory aim to manipulate raw datasets. The output of this step is also raw data, but cleaned and/or pre-formatted.

```
{{repo_name}}/dataset/make_dataset.py
{{repo_name}}/dataset/
```

2.3 Build features

This step takes the output from the previous step “2.2 Make datasets” and creates train, test as well as validation Machine Learning data from raw but cleaned and pre-formatted data. The realisation of this step depends on the concrete use case, the aim of the application as well as available technological backgrounds (e.g. high-performance supports for data processing).

```
{{repo_name}}/features/build_features.py
{{repo_name}}/features/
```

2.4 Develop models

This step deals with the most interesting phase in Machine Learning i.e. modelling. The most important thing is located in `deep_api.py` containing DEEP entry point implementations. DEEP entry points are defined using *API methods*. You don’t need to implement all of them, just the ones you need.

```
{{repo_name}}/models/deep_api.py
{{repo_name}}/models/
```

2.5 Code testing

Code testing, including unit testing, is a necessary part of modern application development. If your project is built based on *Data Science template*, you get an example on how to test `get_metadata()` method and can certainly add more tests. Also `test-requirements.txt` file requests installation of two python testing tools: `stestr` and `pytest`.

```
{{repo_name}}/tests/test_unit_model.py
{{repo_name}}/tests/
```

3. Connect with a remote storage

If you expect your model to use remotely hosted data, you can upload the data in [DEEP-Nextcloud](#) and later trigger copying of data to your container using [rclone](#) tool. The tool requires `rclone.conf` file to exist, even if it is an empty one. In the “[How to use rclone](#)” document you find an extended information and [examples on how to use it from python](#).

Tip: When developing an application with the [Data Science template](#), the Dockerfile already includes creation of an empty `rclone.conf`

Important: DO NOT save the rclone credentials in the CONTAINER

4. Create a python installable package

To create a python installable package the initial directory structure should look something like this:

```
your_model_package/
  your_model_package/
    __init__.py
  setup.py
  setup.cfg
  requirements.txt
  LICENSE
  README
```

- The top level directory will be the root of your repo, e.g. `your_model_package.git`. The subdir, also called `your_model_package`, is the actual python module.
- `setup.py` is the build script for `setuptools`. It tells `setuptools` about your package (such as the name and version) as well as which code files to include. You can find an example of a `setup.py` file [here](#). For the official documentation on how to write your setup script, you can go [here](#).
- `setup.cfg` can be used to get some information from the user, or from the user’s system in order to proceed. Configuration files also let you providedefault values for any command option. An example of a `setup.cfg` file can be found [here](#). The official python documentation on the setup configuration file can be found [here](#).
- `requirements.txt` contains any external requirement needed to run the package. An example of a requirements file can be found [here](#).
- The `README` file will contain information on how to run the package or anything else that you may find useful for someone running your package.
- `LICENSE` It’s important for every package uploaded to the Python Package Index to include a license. This tells users who install your package the terms under which they can use your package. For help choosing a license, go [here](#).

To see how to install your model package, check the Dockerfile in the next section.

5. Create a docker container for your model

Once your model is well in place, you can encapsulate it by creating a docker image. For this you need to modify the Dockerfile created during execution of the *Data Science template*. The Dockerfile is pre-populated with the information you provided while running the cookiecutter template. You may need, however, add packages you need installed to make your project run.

The simplest Dockerfile could look like this:

```
FROM tensorflow/tensorflow:1.14.0-py3

# Install ubuntu updates and python related stuff
# Remember: DEEP API V2 only works with python 3.6 [!]
RUN DEBIAN_FRONTEND=noninteractive apt-get update && \
    apt-get install -y --no-install-recommends \
        git \
        curl \
        wget \
        python3-setuptools \
        python3-pip \
        python3-wheel && \
    apt-get clean && \
    rm -rf /var/lib/apt/lists/* && \
    rm -rf /root/.cache/pip/* && \
    rm -rf /tmp/*

# Set LANG environment
ENV LANG C.UTF-8

WORKDIR /srv

# Install rclone
RUN wget https://downloads.rclone.org/rclone-current-linux-amd64.deb && \
    dpkg -i rclone-current-linux-amd64.deb && \
    apt install -f && \
    mkdir /srv/.rclone/ && touch /srv/.rclone/rclone.conf && \
    rm rclone-current-linux-amd64.deb && \
    apt-get clean && \
    rm -rf /var/lib/apt/lists/* && \
    rm -rf /root/.cache/pip/* && \
    rm -rf /tmp/*

# Install DEEPaaS and FLAAT
RUN pip install --no-cache-dir \
    deepaas \
    flaat

# Download and install your project
RUN git clone https://github.com/your_git/your_project && \
    cd your_project && \
    python -m pip install -e . && \
    cd ..

# Expose API on port 5000 and monitoring port 6006
EXPOSE 5000 6006

CMD ["deepaas-run", "--openwhisk-detect", "--listen-ip", "0.0.0.0", "--listen-port",
    ↪ "5000"]
```

Check the [rclone guide](#) and [DEEPaaS guide](#) for more details.

If you want to see examples of more complex Dockerfiles, you can check various applications [here](#) (look for DEEP-OC-xxx repositories), e.g. [this Dockerfile](#).

In order to compile the Dockerfile, you should choose a name for the docker image and use the docker build command:

```
docker build -t your_docker_image -f Dockerfile .
```

You can then upload it to Docker hub so that you can download the already compiled image directly. To do so, follow the instructions [here](#).

Train a model locally

This is a step by step guide on how to train a module from the Marketplace with your own dataset.

Requirements

- having [Docker](#) installed. For an up-to-date installation please follow the [official Docker installation guide](#).
 - **Optional:** having a [DEEP IAM](#) account if you want to use remote storage resources.
-

1. Choose your module

The first step is to choose a module from the [DEEP Open Catalog marketplace](#). For educational purposes we are going to use a [general model to identify images](#). This will allow us to see the general workflow.

Once we have chosen the model at the [DEEP Open Catalog marketplace](#) we will find that it has an associated docker container in [DockerHub](#). For example, in the example we are running here, the container would be `deephdc/deep-oc-image-classification-tf`. So let's pull the docker image from DockerHub:

```
$ docker pull deephdc/deep-oc-image-classification-tf
```

Docker images have usually tags depending on whether they are using master or test and whether they use cpu or gpu. Tags are usually:

- latest or cpu: master + cpu
- gpu: master + gpu
- cpu-test: test + cpu
- gpu-test: test + gpu

So if you wanted to use gpu and the test branch you could run:

```
$ docker pull deephdc/deep-oc-image-classification-tf:gpu-test
```

Instead of pulling from Dockerhub you can choose to build the image yourself:

```
$ git clone https://github.com/deephdc/deep-oc-image-classification-tf
$ cd deep-oc-image-classification-tf
$ docker build -t deephdc/deep-oc-image-classification-tf .
```

Tip: It's usually helpful to read the README in the source code of the module, in this case located at <https://github.com/deephdc/image-classification-tf>.

2. Store your data

To run locally you have two options:

- Have your data stored locally
- Have your data at a remote storage resource

Have your data stored locally

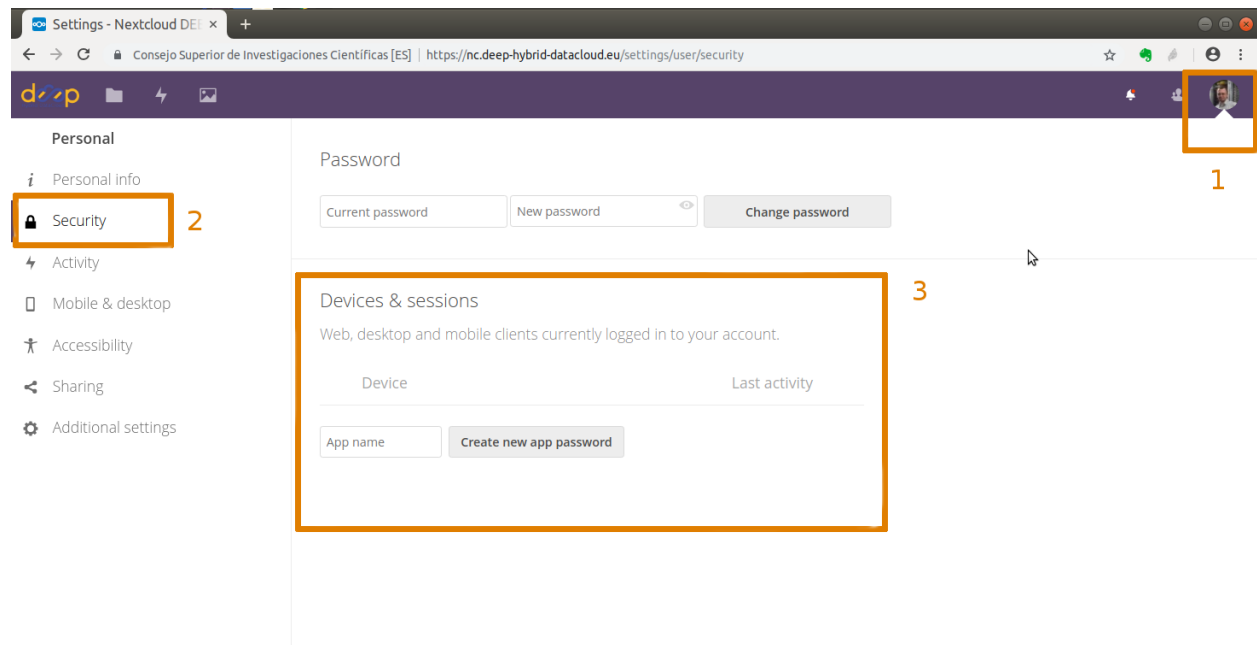
You should make sure that you export inside of the container all the folders you need for the training

```
$ docker run -ti -p 5000:5000 -p 6006:6006 -v path_to_local_folder:path_to_docker_
  ↳ folder deephdc/deep-oc-image-classification-tf
```

Have your data at a remote storage resource

For the time being we support using the DEEP Nextcloud for remote storage, although we plan to support addition platforms such as Google Drive, Dropbox and OneData. All of these platforms are supported through [rclone](#).

If you have the files you need for the training stored in Nextcloud you need first to login into [DEEP-Nextcloud](#) with your DEEP-IAM credentials. Then you have to go to: **(1) Settings (top right corner)** → **(2) Security** → **(3) Devices & sessions**



Set a name for your application (for this example it will be `rshare`) and click on **Create new app password**. This will generate `<your_nextcloud_username>` and `<your_nextcloud_password>` that you should include in

your `rclone.conf` file (see [more details](#)). Now you can create the folders that you need in order to data the inputs needed for the training.

Tip: When developing a model you should add some code to perform a sync to be able to see locally your remote data. If you are using a `trainable` module from the Marketplace that you have not developed yourself you can skip this tip as this will have been taken care of.

In order to be able to see your NextCloud folders from your docker, you should run `rclone` from your module's code, which will synchronize your NextCloud contents with your local contents (or the other way around). It is kind of an `rsync` but for remote storage.

So your module should run this synchronization before the `train()` function tries to access data. To run it from inside a python script you can use the following code:

```
import subprocess

def sync_nextcloud(frompath, topath):
    """
    Mount a NextCloud folder in your local machine or viceversa.
    """
    command = (['rclone', 'copy', frompath, topath])
    result = subprocess.Popen(command, stdout=subprocess.PIPE, stderr=subprocess.PIPE)
    output, error = result.communicate()
    if error:
        warnings.warn("Error while mounting NextCloud: {}".format(error))
    return output, error

sync_nextcloud('rshare:/your/dataset/folder', '/your/data/path/inside/the/container')
↪ # sync local with nextcloud
sync_nextcloud('/your/data/path/inside/the/container', 'rshare:/your/dataset/folder')
↪ # sync nextcloud with local
```

As you can see you can sync the local contents back to NextCloud, which is useful if you want to save your trained model back to NextCloud.

When running the container you should export the `rclone.conf` file so that it can be reached from within the docker. You can see an example on how to do this [here](#)

```
$ docker run -ti -p 5000:5000 -p 6006:6006 -v host_path_to_rclone.conf:/root/.config/
↪ rclone/rclone.conf deephdc/deep-oc-image-classification-tf
```

You can see this last step explained more in detail [here](#). We are using the port 5000 to deploy the API and the port 6006 to monitor the training (for example using [Tensorboard](#)).

3. Train the model

Now comes the fun! Go to <http://0.0.0.0:5000/ui> and look for the `train` method. Modify the training parameters you wish to change and execute. If some kind of monitorization tool is available for this model you will be able to follow the training progress from <http://0.0.0.0:6006>.

Once the training has finished, you can directly test it by clicking on the `predict` method. Upload the image your want to classify and check on the predicted classes.

models		
GET	/v2/models/	Return loaded models and its information
GET	/v2/models/imgclas/	Return model's metadata
POST	/v2/models/imgclas/train/	Retrain model with available data
GET	/v2/models/imgclas/train/	Get a list of trainings (running or completed)
GET	/v2/models/imgclas/train/{uuid}	Get status of a training
DELETE	/v2/models/imgclas/train/{uuid}	Cancel a running training
POST	/v2/models/imgclas/predict/	Make a prediction given the input data

Train a model remotely

This is a step by step guide on how to train a general model from the [DEEP Open Catalog marketplace](#) with your own dataset.

Requirements

Before being able to run your model at the Pilot Infrastructure you should first fulfill the following prerequisites:

- [DEEP-IAM](#) registration

For this example we are going to use [DEEP-Nextcloud](#) for storing data. This means you also have to:

- Register at [DEEP-Nextcloud](#) (i.e. login with your DEEP-IAM credentials)
- *Obtain Nextcloud security parameters for rclone.* This will provide you with `<your_nextcloud_username>` and `<your_nextcloud_password>`.


1. Choose a model

The first step is to choose a model from the [DEEP Open Catalog marketplace](#). For educational purposes we are going to use a [general model to identify images](#). Some of the model dependent details can change if using another model, but this tutorial will provide a general overview of the workflow to follow when using any of the models in the Marketplace. A demo showing the different steps in this HowTo has also be recorded and you can find it our [YouTube channel](#).

2. Upload your files to Nextcloud

We login to DEEP-Nextcloud with DEEP-IAM credentials and upload there the files needed for training. We create the folders that you need in order to store the inputs needed for the training and to retrieve the output. These folders

will be visible from within the container. In this example we just need two folders (`models` and `data`):



- A folder called **models** where the training weights will be stored after the training
- A folder called **data** that contains two different folders:
 - The folder **images** containing the input images needed for the training
 - The folder **dataset_files** containing a couple of scripts: *train.txt* indicating the relative path to the training images and *classes.txt* indicating which are the categories for the training

The folder structure and their content will of course depend on the model to be used. This structure is just an example in order to complete the workflow for this tutorial.

3. Deploy with the Training Dashboard

- Go to the [Training Dashboard](#)
- Login with DEEP-IAM credentials
- Find the **Train an image classifier** module, click **Train module**
- Fill the webform, e.g. provide `<your_nextcloud_username>` for `rclone_user` and `<your_nextcloud_password>` for `rclone_password` (see [Nextcloud configuration for rclone](#))
- Click **Submit**
- Look for **Access** and choose **DEEPaaS**, you will be redirected to `http://deepaas_endpoint`

If you prefer deploying using a more advanced deploying via the command-line-interface check the [deploy with CLI](#) guide.

See the [Dashboard guide](#) for more details.

4. Go to the API, train the model

Now comes the fun!

Go to `http://deepaas_endpoint/ui` and look for the `train` POST method. Modify the training parameters you wish to change and execute. If some kind of monitorization tool is available for the module, you will be able to follow the training progress at `http://monitor_endpoint`.

models		
GET	/v2/models/	Return loaded models and its information
GET	/v2/models/imgclas/	Return model's metadata
POST	/v2/models/imgclas/train/	Retrain model with available data
GET	/v2/models/imgclas/train/	Get a list of trainings (running or completed)
GET	/v2/models/imgclas/train/{uuid}	Get status of a training
DELETE	/v2/models/imgclas/train/{uuid}	Cancel a running training
POST	/v2/models/imgclas/predict/	Make a prediction given the input data

In the Dashboard you will be able to view the training history of that deployment:

TRAINING DASHBOARD				
DEPLOYMENTS	ADVANCED	EXTERNAL LINKS	IGNACIO HEREDIA	deep
11ea3ddc-0dd2-7ff3-a966-0242eca7ffec				
mods Train model Predict				
Training uuid	Date	Status	Message	Actions
4473224736a24c869fce3da38c031288	2020-01-28 12:19:31.943220	error	This is a module error.	Delete
55dc244cee04204a1dd8e7f9e59190a	2020-01-28 12:20:46.987462	done		Delete
691d7d55b52744c4bac86b5eff8de5b3	2020-01-28 12:08:00.437706	error	This is a module error.	Delete
b6686ab3de104563b0500e1a0caa4e3f	2020-01-28 12:22:52.640658	done		Delete
df2d13e3da13420496fbc3bd022c7f89	2020-01-28 12:22:58.749524	done		Delete
26a906a603a047859b5ad5aa79c35ca7	2020-01-28 12:23:07.996410	done		Delete
ac4dac0c06854913a964a9712629163	2020-01-28 12:23:36.010563	done		Delete
0e012dea06f948f9b094bb980d022996	2020-01-28 12:23:36.808146	running		Delete
Training history				

5. Testing the training

Once the training has finished, you can directly test it by clicking on the `predict` POST method. There you can either upload the image your want to classify or give an URL to it.

Try a service locally

Requirements

- having [Docker](#) installed. For an up-to-date installation please follow the [official Docker installation guide](#).

1. Choose your module

The first step is to choose a module from the [DEEP Open Catalog marketplace](#). For educational purposes we are going to use a [general model to identify images](#). This will allow us to see the general workflow.

Once we have chosen the model at the [DEEP Open Catalog marketplace](#) we will find that it has an associated docker container in [DockerHub](#). For example, in the example we are running here, the container would be `deephdc/deep-oc-image-classification-tf`. This means that to pull the docker image and run it you should:

```
$ docker pull deephdc/deep-oc-image-classification-tf
```

Docker images have usually tags depending on whether they are using master or test and whether they use cpu or gpu. Tags are usually:

- latest or cpu: master + cpu
- gpu: master + gpu
- cpu-test: test + cpu
- gpu-test: test + gpu

So if you wanted to use gpu and the test branch you could run:

```
$ docker pull deephdc/deep-oc-image-classification-tf:gpu-test
```

Instead of pulling from Dockerhub you can choose to build the image yourself:

```
$ git clone https://github.com/deephdc/deep-oc-image-classification-tf
$ cd deep-oc-image-classification-tf
$ docker build -t deephdc/deep-oc-image-classification-tf .
```

Tip: It's usually helpful to read the README in the source code of the module, in this case located at <https://github.com/deephdc/image-classification-tf>.

2. Launch the API and predict

Run the container with:

```
$ docker run -ti -p 5000:5000 deephdc/deep-oc-image-classification-tf
```

Once running, point your browser to <http://127.0.0.1:5000/ui> and you will see the API documentation, where you can test the module's functionality, as well as perform other actions.

models		
GET	/v2/models/	Return loaded models and its information
GET	/v2/models/imgclas/	Return model's metadata
POST	/v2/models/imgclas/train/	Retrain model with available data
GET	/v2/models/imgclas/train/	Get a list of trainings (running or completed)
GET	/v2/models/imgclas/train/{uuid}	Get status of a training
DELETE	/v2/models/imgclas/train/{uuid}	Cancel a running training
POST	/v2/models/imgclas/predict/	Make a prediction given the input data

Go to the `predict()` function and upload the file/data you want to predict (in the case of the image classifier this should be an image file). The appropriate data formats of the files you have to upload are often discussed in the module's Marketplace page.

The response from the `predict()` function will vary from module to module but usually consists on a JSON dict with the predictions. For example the image classifier return a list of predicted classes along with predicted accuracy. Other modules might return files instead of a JSON.

Add you model/service to the DEEP marketplace

This document describes how to add your trained service or model to the DEEP marketplace.

Creating the Github repositories

You model must have a repository to host the code and a repository to host the Dockerfiles. Both these repositories can be hosted under your personal Github account. Naming conventions enforce that the Docker repo name is the same as the code repo name with the prefix `DEEP-OC-`.

A typical example of this can be:

- `deephdc/image-classification-tf` - The Github repo hosting the code of an image classification model.
- `deephdc/DEEP-OC-image-classification-tf` - The Github repo hosting the Dockerfiles of the image classification model.

In case you are only developing a service based on an already existing module (like for example developing an animal classifier based on the `image-classification-tf` module) you only need to create the Docker repo.

The code repo

This is the repo containing the code of your model. If you are adding a service (ie. a trained model) the weights of the trained model must be stored in a location accessible over a network connection, so that your container can download

them upon creation.

A few MUSTs your code has to comply with in order to ensure compatibility and ease of use:

- your code must be packaged in order to be `pip` installable. This should be the default behaviour if you used the *DEEP Data Science template* to develop your code.
- your code must be integrated with the DEEPaaS API. Check *this guide* on how to do this.

The Docker repo

If you used the *DEEP Data Science template* to develop your code, a template of this repo should have been created alongside the template of your code.

This repo has to contain at least the following files (see the *Generic container* for a template):

- `Dockerfile`

This is the file used to build a container from your application. If you developed your application from the *DEEP Data Science template* you should have a draft of this file (although you might need to add additional code depending on the requirements of your model).

If you are adding a service derived from an existing module, it is good practice to draw inspiration from the Dockerfiles of the module or the services derived from that module (see for example the *plant classification Dockerfile* derived from the *image classification model*).

Some steps common to all Dockerfiles include cloning the model code repo, `pip` installing the DEEPaaS API, installing `rcclone` and downloading the trained weights if you are adding a service. For the details of all these steps please refer to this *Dockerfile example*.

- `Jenkinsfile`

This is the file that runs the Jenkins pipeline. You can copy this *Jenkinsfile example* replacing the repo names with your own Docker repo name.

- `metadata.json`

This file contains the information that is going to be displayed in the Marketplace. You can build your own starting from this *metadata.json example*. This metadata will be validated during integration tests when the PR is accepted but you can manually *validate the metadata* beforehand by running:

```
pip install git+https://github.com/deephdc/schema4apps
deep-app-schema-validator metadata.json
```

Making the Pull Request (PR)

Once your repos are set it's time to make a PR to add your model to the marketplace! For this you have to fork the code of the DEEP catalog repo (*deephdc/deep-oc*) and add your Docker repo name at the end of the `MODULES.yml`.

```
git clone https://github.com/[my-github-fork]
cd [my-github-fork]
echo '- module: https://github.com/[my-account-name]/DEEP-OC-[my-app-name]' >> MODULES.yml
git commit -a -m "adding new module to the catalogue"
git push
```

You can also make it *online on GitHub*.

Once the changes are done, make a PR of your fork to the original repo and wait for approval. Check the [GitHub Standard Fork & Pull Request Workflow](#) in case of doubt.

How to use rclone

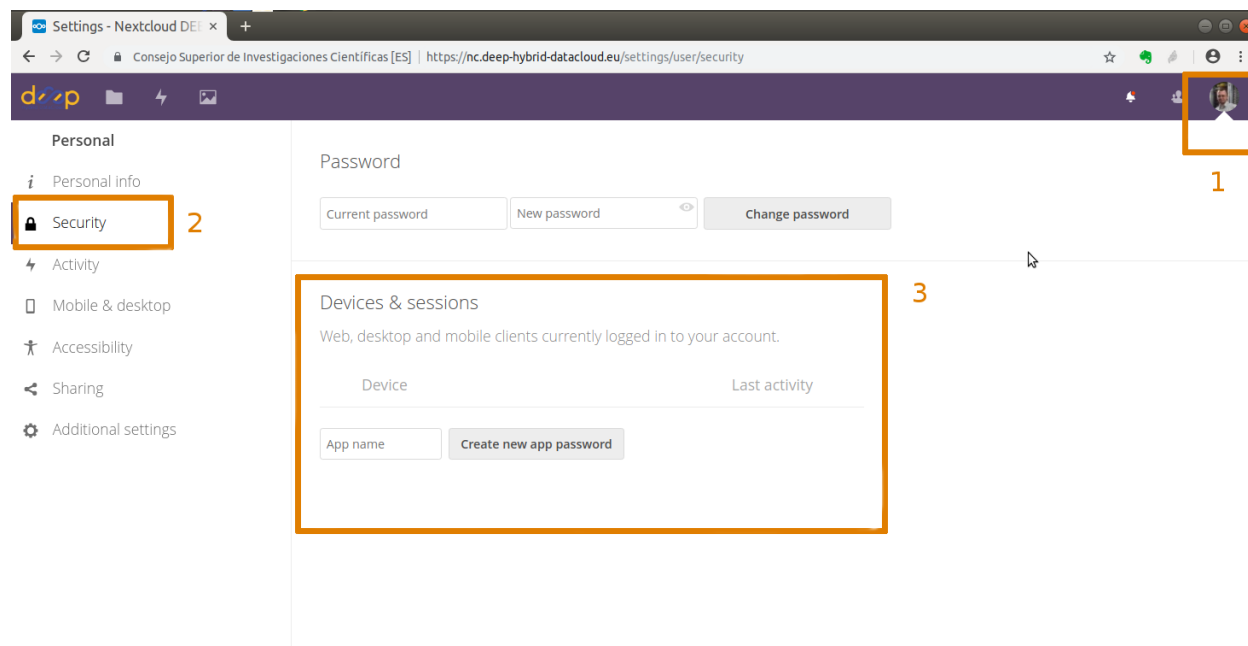
Installation of rclone in Docker image

All applications in the [DEEP Open Catalog](#) are packed in a Docker image and have [rclone](#) tool installed by default. If you want to create a Docker containing your own application, you should install rclone in the container to be able to access the data stored remotely. The following lines are an example of what has to be added in the Dockerfile when installation is based on Ubuntu. For other Linux flavors, please, refer to the [rclone official site](#)

```
# Install rclone
RUN wget https://downloads.rclone.org/rclone-current-linux-amd64.deb && \
    dpkg -i rclone-current-linux-amd64.deb && \
    apt install -f && \
    mkdir /srv/.rclone/ && touch /srv/.rclone/rclone.conf && \
    rm rclone-current-linux-amd64.deb && \
    apt-get clean && \
    rm -rf /var/lib/apt/lists/* && \
    rm -rf /root/.cache/pip/* && \
    rm -rf /tmp/*
```

Tip: When developing an application with the [Data Science template](#), the Dockerfile already includes installation of rclone.

Nextcloud configuration for rclone



After login into [DEEP-Nextcloud](#) with your DEEP-IAM credentials, go to (1) **Settings (top right corner)** → (2) **Security** → (3) **Devices & sessions**. Set a name for your application and click on **Create new app password**.

That user and password is what one needs to include in the rclone config file (`rclone.conf`) to run locally or as `rclone_user` and `rclone_password` either in the Dashboard webform or in the `orchent` script to generate the deployment when running remotely (see [here](#) and [here](#)).

Creating rclone.conf for your local host

You can install rclone at your host or run Docker image with rclone installed (see installation steps of rclone above). In order to create the configuration file (`rclone.conf`) for rclone:

```
$ rclone config
  choose "n"   for "New remote"
  choose name for DEEP-Nextcloud, e.g. deep-nextcloud
  choose "Type of Storage" \u2192 "Webdav" (24)
  provide DEEP-Nextcloud URL for webdav access: https://nc.deep-hybrid-datacloud.
  ↪ eu/remote.php/webdav/
  choose Vendor, Nextcloud (1)
  specify "user" (see "Nextcloud configuration for rclone" above). Your username_
  ↪ starts with "DEEP-IAM-..."
  specify password (see "Nextcloud configuration for rclone" above).
  by default rclone.conf is created in your $HOME/.config/rclone/rclone.conf
```

Important: The `rclone.conf` file should be in your host, i.e. outside of container. **DO NOT STORE IT IN THE CONTAINER**

Then one has two options:

If you know under what user you run your application in the container (e.g. if `docker` or `nvidia-docker` is used, most probably this is `'root'`) you can mount your host `rclone.conf` into the container as:

```
$ docker run -ti -v $HOSTDIR_WITH_RCLONE_CONF/rclone.conf:/root/.config/rclone/rclone.
  ↪ conf <your-docker-image>
```

i.e. you mount `rclone.conf` file itself directly as a volume.

One can also mount rclone directory with the `rclone.conf` file:

```
$ docker run -ti -v $HOSTDIR_WITH_RCLONE_CONF:/root/.config/rclone <your-docker-image>
```

A more reliable way can be to mount either rclone directory or directly `rclone.conf` file into a pre-defined location and not (container) user-dependent place:

```
$ docker run -ti -v $HOSTDIR_WITH_RCLONE_CONF:/rclone <your-docker-image>
```

One has, however, to call rclone with `--config` option to point to the `rclone.conf` file, e.g:

```
$ rclone --config /rclone/rclone.conf ls deep-nextcloud:/Datasets/dogs_breed/models
```

Example code on usage of rclone from python

Simple example

A simple call of rclone from python is via `subprocess.Popen()`

```
import subprocess

# from "rshare" remote storage into the container
command = (['rclone', 'copy', 'rshare:/Datasets/dogs_breed/data', '/srv/dogs_breed_
↪det/data'])

result = subprocess.Popen(command, stdout=subprocess.PIPE, stderr=subprocess.PIPE)
output, error = result.communicate()
```

Important: When deploying a module on the DEEP Pilot testbed, you pass rclone parameters e.g. `rclone_user` and `rclone_password` during the deployment. If you use our [general template](#), the name of the remote storage has to be `rshare` as in the example above (`rshare:/Datasets/dogs_breed/data`). If you create your own TOSCA template, you need to pay attention on matching these names in your code and in the template (for example, see environment parameters in the [general template](#) like `RCLONE_CONFIG_RSHARE_USER` etc).

Advanced examples

More advanced usage includes calling rclone with various options (`ls`, `copy`, `check`) in order to check file existence at Source, check if after copying two versions match exactly.

- `rclone_call`

```
def rclone_call(src_path, dest_dir, cmd = 'copy', get_output=False):
    """ Function
        rclone calls
    """
    if cmd == 'copy':
        command = (['rclone', 'copy', '--progress', src_path, dest_dir])
    elif cmd == 'ls':
        command = (['rclone', 'ls', src_path])
    elif cmd == 'check':
        command = (['rclone', 'check', src_path, dest_dir])

    if get_output:
        result = subprocess.Popen(command, stdout=subprocess.PIPE, stderr=subprocess.
↪PIPE)
    else:
        result = subprocess.Popen(command, stderr=subprocess.PIPE)
    output, error = result.communicate()
    return output, error
```

- `rclone_copy`

```
def rclone_copy(src_path, dest_dir, src_type='file', verbose=False):
    """ Function for rclone call to copy data (sync?)
    :param src_path: full path to source (file or directory)
    :param dest_dir: full path to destination directory (not file!)
    :param src_type: if source is file (default) or directory
    :return: if destination was downloaded, and possible error
    """

    error_out = None

    if src_type == 'file':
```

(continues on next page)

(continued from previous page)

```

src_dir = os.path.dirname(src_path)
dest_file = src_path.split('/')[-1]
dest_path = os.path.join(dest_dir, dest_file)
else:
    src_dir = src_path
    dest_path = dest_dir

# check first if we find src_path
output, error = rclone_call(src_path, dest_dir, cmd='ls')
if error:
    print('[ERROR, rclone_copy()] %s (src):\n%s' % (src_path, error))
    error_out = error
    dest_exist = False
else:
    # if src_path exists, copy it
    output, error = rclone_call(src_path, dest_dir, cmd='copy')
    if not error:
        output, error = rclone_call(dest_path, dest_dir,
                                    cmd='ls', get_output=True)
        file_size = [elem for elem in output.split(' ') if elem.isdigit()][0]
        print('[INFO] Copied to %s %s bytes' % (dest_path, file_size))
        dest_exist = True
        if verbose:
            # compare two directories, if copied file appears in output
            # as not found or not matching -> Error
            print('[INFO] File %s copied. Check if (src) and (dest) really match..
→ ' % (dest_file))
            output, error = rclone_call(src_dir, dest_dir, cmd='check')
            if 'ERROR : ' + dest_file in error:
                print('[ERROR, rclone_copy()] %s (src) and %s (dest) do not match!
→ '
                        % (src_path, dest_path))
                error_out = 'Copy failed: ' + src_path + ' (src) and ' + \
                    dest_path + ' (dest) do not match'
                dest_exist = False
            else:
                print('[ERROR, rclone_copy()] %s (src):\n%s' % (dest_path, error))
                error_out = error
                dest_exist = False

return dest_exist, error_out

```

Configure oidc-agent

`oidc-agent` is a tool to manage OpenID Connect tokens and make them easily usable from the command line.

Requirements

- having a [DEEP IAM](#) account
- having the `oidc-agent` installed (follow the [official installation guide](#)).

Start `oidc-agent`:

```
$ eval $(oidc-agent)
$ oidc-gen
```

You will be asked for the name of the account to configure. Let's call it **deep-iam**. After that you will be asked for the additional *client-name-identifier*, you should choose the option:

```
[2] https://iam.deep-hybrid-datacloud.eu/
```

Then just click Enter to accept the default values for Space delimited list of scopes [openid profile offline_access]. After that, if everything has worked properly, you should see the following messages:

```
Registering Client ...
Generating account configuration ...
accepted
```

At this point you will be given a URL. You should visit it in the browser of your choice in order to continue and approve the registered client. For this you will have to login into your DEEP-IAM account and accept the permissions you are asked for.

Once you have done this you will see the following message:

```
The generated account config was successfully added to oidc-agent. You don't have to
↪run oidc-add
```

Next time you want to start oidc-agent from scratch, you will only have to do:

```
$ eval $(oidc-agent)
oidc-add deep-iam
Enter encryption password for account config deep-iam: *****
success
```

You can print the token:

```
$ oidc-token deep-iam
```

Usage with orchent

You should set OIDC_SOCK (this is not needed, if you did it before):

```
$ eval $(oidc-agent)
$ oidc-add deep-iam
```

Set the agent account to be used with orchent and the ORCHENT_URL:

```
$ export ORCHENT_AGENT_ACCOUNT=deep-iam
$ export ORCHENT_URL="https://paas.cloud.cnaf.infn.it/orchestrator"
```

Deployment with CLI (orchent)

This is a step by step guide on how to make a deployment using the command line interface (instead of the Training Dashboard).

Requirements

- `oidc-agent` installed and configured for **DEEP-IAM** (see *Configure oidc-agent*).

- [orchent](#) tool
-

Prepare your TOSCA file (optional)

The orchent tool needs TOSCA YAML file to configure and establish the deployment. One can generate an application specific TOSCA template or use a general one, [deep-oc-marathon-webdav.yml](#), while providing necessary inputs in the bash script (see next subsection).

If you create your own TOSCA YAML file, the following sections should be modified (TOSCA experts may modify the rest of the template to their will):

- Docker image to deploy. In this case we will be using `deephdc/deep-oc-image-classification-tf`:

```
docker_img:
  type: string
  description: docker image from Docker Hub to deploy
  required: yes
  default: deephdc/deep-oc-image-classification-tf
```

- Location of the `rclone.conf` (this file can be empty, but should be at the indicated location):

```
rclone_conf:
  type: string
  description: rclone.conf location
  required: yes
  default: "/srv/image-classification-tf/rclone.conf"
```

For further TOSCA templates examples you can go [here](#).

Important: DO NOT save the rclone credentials in the **CONTAINER** nor in the **TOSCA** file

Orchent submission script

You can use the general template, [deep-oc-mesos-webdav.yml](#), but provide necessary parameters in a bash script. Here is an example for such a script, e.g. `submit_orchent.sh` :

```
#!/bin/bash

orchent depcreate ./deep-oc-marathon-webdav.yml '{ "docker_image": "deephdc/deep-oc-
↪image-classification-tf"
                                                    "rclone_url": "https://nc.deep-hybrid-
↪datacloud.eu/remote.php/webdav/",
                                                    "rclone_vendor": "nextcloud",
                                                    "rclone_conf": "/srv/image-
↪classification-tf/rclone.conf"
                                                    "rclone_user": <your_nextcloud_
↪username>
                                                    "rclone_pass": <your_nextcloud_
↪password> }'
```

This script will be the **only place** where you will have to indicate `<your_nextcloud_username>` and `<your_nextcloud_password>`. This file should be stored locally and secured.

Important: **DO NOT** save the rclone credentials in the **CONTAINER** nor in the **TOSCA** file

Tip: When developing an application with the *Data Science template*, the DEEP-OC-<your_project> repository will contain an example script, named *submit_orchent_tmpl.sh*

Submit your deployment

The submission is then done by running the orchent submission script you generated in the previous step:

```
./submit_orchent.sh
```

This will give you a bunch of information including your deployment ID. To check status of your job:

```
$ orchent depshow <Deployment ID>
```

Once your deployment is in status **CREATED**, you will be given various endpoints:

```
http://deepaas_endpoint  
http://monitor_endpoint
```

N.B.: to check all your deployments:

```
$ orchent depls -c me
```

Video demos

Please refer to the Youtube channel [playlist](#) for a comprehensive list of video demos.

Available demos include:

- How to use the Cookiecutter template
- How to deploy/train a module locally and remotely
- How to use some of the modules available in the Marketplace
- How to develop a module in the Testbed
- ...

1.1.4 Modules

All modules are found at the [DEEP Marketplace](#), the source code is hosted under [Github's deephdc](#) organization and the corresponding Docker images are hosted under [DockerHub's deephdc](#) organization.

Github repositories follow the following convention:

- deephdc/some_module: source code of the module
- deephdc/DEEP-OC-some_module: Dockerfiles and metadata of that module.

Docker images have usually tags depending on whether they are using Github's `master` or `test` and whether they use `cpu` or `gpu`. Tags are usually:

- latest or cpu: master + cpu
- gpu: master + gpu
- cpu-test: test + cpu
- gpu-test: test + gpu

CHAPTER 2

Component documentation

Individual components' documentation can be found here:

- [DEEPaaS documentation](#)

Technical documentation

If you are searching for technical notes on various areas, please check the following section.

3.1 Technical documentation

These pages contain technical notes software documentations, guides, tutorials, logbooks and similar documents produced with DEEP Hybrid DataCloud project

3.1.1 Mesos

Introduction

Mesos 1.0.0 added first-class support for Nvidia GPUs. The minimum required Nvidia driver version is 340.29

Enabling GPU support in a Mesos cluster is really straightforward (as stated in the official project documentation and as documented in this page). It consists in the following steps:

1. configuring the agent nodes in order to expose the available gpus as resources to be advertised to the master nodes;
2. enabling the framework GPU_RESOURCES capability so that the master includes the GPUs in the resource offers sent to the frameworks.

Mesos exposes GPUs as a simple *SCALAR* resource in the same way it always has for CPUs, memory, and disk.

An important remark is that currently the GPU support is available for the Mesos containerizer and not for the Docker containerizer. Anyway the Mesos containerizer is now able to run docker images natively through the Universal Container Runtime (UCR).

The following limitations can, on the other hand, have impacts on the deployment of Long-Running services (Marathon) requiring GPUs:

- The UCR does not support the following: runtime privileges, Docker options, force pull, named ports, numbered ports, bridge networking, port mapping.

It is important to remember that the task definition must be properly written in order to specify the right containerizer (type=MESOS).

For Marathon:

```
{
  "id": "test",
  "cpus": 2,
  "mem": 2048,
  [...]
  "container": {
    "type": "MESOS",
    "docker": {
      "image": "tensorflow/tensorflow"
    }
  }
}
```

See also <https://mesosphere.github.io/marathon/docs/native-docker.html#provisioning-containers-with-the-ucr>

For Chronos:

```
{
  "name": "test-gpu",
  "command": "",
  "cpus": 1,
  "mem": 4096,
  [...]
  "container": {
    "type": "MESOS",
    "image": "tensorflow/tensorflow"
  },
  "schedule": "...
}
```

The GPU support is fully implemented and officially documented in Mesos and Marathon Framework whereas Chronos Framework does not support GPU resources yet. Anyway there is a pull request (still open) that seems in good shape and we have decided to give it a try.

Testbed Setup

Nodes characteristics

Tested Components Versions

Prepare the agent (slave) node

Download the driver repo from <http://www.nvidia.com/Download/index.aspx?lang=en-us> choosing the proper version.

Install the downloaded .deb file (repo), install the driver and reboot:

```
dpkg -i nvidia-diag-driver-local-repo-ubuntu1604-390.30_1.0-1_amd64.deb
apt-key add /var/nvidia-diag-driver-local-repo-390.30/7fa2af80.pub
apt-get update
apt-get install nvidia-390
reboot
```

Alternatively you can enable the graphics-drivers PPA. Currently, it supports Ubuntu 18.04 LTS, 17.10, 17.04, 16.04 LTS, and 14.04 LTS operating systems (still under testing phase):

```
add-apt-repository ppa:graphics-drivers/ppa
apt update
apt install nvidia-390
```

Verify the nvidia-driver installation

Launch the command:

```
nvidia-smi
```

Output:

```
+-----+
| NVIDIA-SMI 390.30                 Driver Version: 390.30                 |
+-----+-----+-----+-----+-----+-----+
| GPU  Name            Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+-----+-----+-----+-----+
|   0   Tesla K40m        Off      | 00000000:82:00:0 Off |             0      |
| N/A   21C    P8        18W / 235W |      0MiB / 11441MiB |      0%      Default |
+-----+-----+-----+-----+-----+-----+

+-----+
| Processes:                         GPU Memory |
|  GPU       PID    Type    Process name      Usage    |
+-----+-----+-----+-----+-----+
| No running processes found              |
+-----+
```

Mesos slave configuration

```
export MESOS_MASTER="zk://<master>:2181/mesos"
export MESOS_ZK="zk://<master>:2181/mesos"
export MESOS_EXECUTOR_REGISTRATION_TIMEOUT="10mins"
export MESOS_CONTAINERIZERS="mesos,docker"
export MESOS_LOG_DIR="/var/log/mesos"
export MESOS_IP="<agent-ip>"
export MESOS_WORK_DIR="/var/lib/mesos"
export MESOS_HOSTNAME="<agent-hostname>"
export MESOS_ISOLATION="docker/runtime,filesystem/linux,cgroups/devices,gpu/nvidia"
export MESOS_IMAGE_PROVIDERS="docker"
```

(Re)start the mesos agent. In the log you will see the available GPU as resource offered by the agent node:

```
Feb 26 23:10:07 hpc-09-02 mesos-slave[1713]: I0226 23:10:07.155503 1745 slave.cpp:533] Agent resources:
gpus(*):1; cpus(*):40; mem(*):256904; disk(*):3489713; ports(*):[31000-32000]
```

Testing GPU support in Mesos

Verify that Mesos is able to launch a task consuming GPUs:

```
mesos-execute --master=mesos-m0.recas.ba.infn.it:5050 --name=gpu-test --docker_
↳ image=nvidia/cuda --command="nvidia-smi" --framework_capabilities="GPU_
↳ RESOURCES" --resources="gpus:1"
I0305 15:22:38.346174 4443 scheduler.cpp:188] Version: 1.5.0
I0305 15:22:38.349104 4459 scheduler.cpp:311] Using default 'basic' HTTP_
↳ authenticatee
I0305 15:22:38.349442 4462 scheduler.cpp:494] New master detected at master@172.20.0.
↳ 38:5050
Subscribed with ID 6faa9a75-d48b-4dc6-96ee-73c35997706b-0017
Submitted task 'gpu-test' to agent 'd33d527c-8d1f-4e53-b65d-e2b2c67c0889-S2'
Received status update TASK_STARTING for task 'gpu-test'
  source: SOURCE_EXECUTOR
Received status update TASK_RUNNING for task 'gpu-test'
  source: SOURCE_EXECUTOR
Received status update TASK_FINISHED for task 'gpu-test'
  message: 'Command exited with status 0'
  source: SOURCE_EXECUTOR
```

Look into the task sandbox. The stdout should report the following:

```
Marked '/' as rslave
Prepared mount '{"flags":20480,"source":"\var\lib\mesos\slaves\d33d527c-8d1f-
↳ 4e53-b65d-e2b2c67c0889-S2\frameworks\6faa9a75-d48b-4dc6-96ee-73c35997706b-0017\
↳ executors\gpu-test\runs\5ebbfaf3-3b8b-4c32-9337-740a85feef75","target":"\var\
↳ lib\mesos\provisioner\containers\5ebbfaf3-3b8b-4c32-9337-740a85feef75\
↳ backends\overlay\rootfses\e56d62ea-4334-4582-a820-2b9406e2b7f8\mnt\mesos\
↳ sandbox"}'
Prepared mount '{"flags":20481,"source":"\var\run\mesos\isolators\gpu\nvidia_
↳ 390.30","target":"\var\lib\mesos\provisioner\containers\5ebbfaf3-3b8b-4c32-
↳ 9337-740a85feef75\backends\overlay\rootfses\e56d62ea-4334-4582-a820-
↳ 2b9406e2b7f8\usr\local\nvidia"}'
Prepared mount '{"flags":20513,"target":"\var\lib\mesos\provisioner\containers\
↳ 5ebbfaf3-3b8b-4c32-9337-740a85feef75\backends\overlay\rootfses\e56d62ea-4334-
↳ 4582-a820-2b9406e2b7f8\usr\local\nvidia"}'
Changing root to /var/lib/mesos/provisioner/containers/5ebbfaf3-3b8b-4c32-9337-
↳ 740a85feef75/backends/overlay/rootfses/e56d62ea-4334-4582-a820-2b9406e2b7f8
Mon Mar  5 14:23:41 2018
+-----+
| NVIDIA-SMI 390.30                Driver Version: 390.30                |
+-----+-----+-----+-----+-----+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+-----+-----+-----+-----+
|    0   Tesla K40m        Off      | 00000000:82:00:0 Off  |            0         |
| N/A   21C    P8         18W / 235W|    0MiB / 11441MiB |      0%      Default  |
+-----+-----+-----+-----+-----+-----+

+-----+
| Processes:                                     GPU Memory |
|  GPU       PID    Type    Process name                     Usage      |
+-----+-----+-----+-----+-----+-----+
| No running processes found                                |
+-----+
```

Testing Chronos patch for GPU support

Patch available on github: <https://github.com/mesos/chronos/pull/810>

Patch compilation

The following steps are needed to test the patch:

1. Get the patched code:

```
git clone https://github.com/mesos/chronos.git -b v3.0.2
cd chronos
git fetch origin pull/810/head:chronos
git checkout chronos
```

2. Compile:

```
docker run -v `pwd`::/chronos --entrypoint=/bin/sh maven:3-jdk-8 -c "\
curl -sL https://deb.nodesource.com/setup_7.x | bash - \
&& apt-get update && apt-get install -y --no-install-recommends nodejs \
&& ln -sf /usr/bin/nodejs /usr/bin/node \
&& cd /chronos \
&& mvn clean \
&& mvn versions:set -DnewVersion=3.0.2-1 \
&& mvn package -DskipTests"
```

The jar **chronos-3.0.2-1.jar** will be created in the folder **./target/**

3. Create the docker image **Dockerfile**:

```
FROM ubuntu:16.04
ENV DEBIAN_FRONTEND noninteractive
RUN apt-key adv --keyserver keyserver.ubuntu.com --recv E56151BF && \
echo deb http://repos.mesosphere.io/ubuntu trusty main > /etc/apt/sources.list.d/\
↪mesosphere.list && \
apt-get update && \
apt-get -y install --no-install-recommends mesos openjdk-8-jre-headless && \
apt-get clean && \
rm -rf /var/lib/apt/lists/* /tmp/* /var/tmp/*
ADD chronos-3.0.2-1.jar /
COPY entrypoint.sh /
ENTRYPOINT ["/entrypoint.sh"]
```

File entrypoint.sh:

```
#!/bin/sh
CMD="java -Xmx512m -cp chronos-3.0.2-1.jar org.apache.mesos.chronos.scheduler.Main
↪"
# Parse environment variables
for k in `set | grep ^CHRONOS_ | cut -d= -f1`; do
    eval v=\${$k}
    CMD="$CMD --`echo $k | cut -d_ -f2- | tr '[:upper:]' '[:lower:]'` $v"
done
# authentication
PRINCIPAL=${PRINCIPAL:-root}
if [ -n "$SECRET" ]; then
    touch /tmp/secret
    chmod 600 /tmp/secret
    echo -n "$SECRET" > /tmp/secret
    CMD="$CMD --mesos_authentication_principal $PRINCIPAL --mesos_authentication_
↪secret_file /tmp/secret"
fi
```

(continues on next page)

(continued from previous page)

```
echo $CMD
if [ $# -gt 0 ]; then
    exec "$@"
fi
exec $CMD
```

Start the patched Chronos Framework:

Using the docker image described above you can run Chronos as follows:

```
docker run --name chronos -d --net host --env-file /etc/chronos/.chronosenv chronos:3.
↪0.2_gpu
```

with the following environment:

```
LIBPROCESS_IP=172.20.0.38
CHRONOS_HOSTNAME=172.20.0.38
CHRONOS_HTTP_PORT=4400
CHRONOS_MASTER=zk://172.20.0.38:2181/mesos
CHRONOS_ZK_HOSTS=zk://172.20.0.38:2181
CHRONOS_ZK_PATH=/chronos/state
CHRONOS_MESOS_FRAMEWORK_NAME=chronos
CHRONOS_HTTP_CREDENTIALS=admin:*****
CHRONOS_ENABLE_FEATURES=gpu_resources
```

Testing

Approach: submit a batch-like job that uses the tensorflow docker image, downloads the code available here and runs the convolutional network example

```
apt-get update; apt-get install -y git

git clone https://github.com/aymericdamien/TensorFlow-Examples

cd TensorFlow-Examples/examples/3_NeuralNetworks;

time python convolutional_network.py
```

The test is based on the tutorial provided by mesosphere DC/OS

Two different versions of the tensorflow docker image will be used in order to verify the correct execution of the job regardless of the version of CUDA and cuDNN used to build the binaries inside the docker image:

Test #1:

Job definition:

```
{
  "name": "test-gpu",
  "command": "cd $MESOS_SANDBOX && /bin/bash gpu_demo.sh",
  "shell": true,
  "retries": 2,
  "cpus": 4,
  "disk": 256,
```

(continues on next page)

(continued from previous page)

```

"mem": 4096,
"gpus": 1,
"uris": [
  "https://gist.githubusercontent.com/maricaantonacci/
↪1a7f02903513e7bba91f451e0f4f5ead/raw/78c737fd0e2a288a2040c192368f6c4ecf8eb88a/gpu_
↪demo.sh"
],
"environmentVariables": [],
"arguments": [],
"runAsUser": "root",
"container": {
  "type": "MESOS",
  "image": "tensorflow/tensorflow:latest-gpu"
},
"schedule": "R/2018-03-05T23:00:00.000Z/PT24H"
}

```

The job is correctly run. The following relevant info were retrieved from the stderr file in the job sandbox:

```

Cloning into 'TensorFlow-Examples'...
/usr/local/lib/python2.7/dist-packages/h5py/__init__.py:36: FutureWarning: Conversion
↪ of the second argument of issubdtype from `float` to `np.floating` is deprecated.
↪ In future, it will be treated as `np.float64 == np.dtype(float).type`.
  from ._conv import register_converters as _register_converters
WARNING:tensorflow:Using temporary folder as model directory: /tmp/tmpLLDDDs
2018-03-05 14:38:59.059890: I tensorflow/core/common_runtime/gpu/gpu_device.cc:1212]
↪ Found device 0 with properties:
name: Tesla K40m major: 3 minor: 5 memoryClockRate(GHz): 0.745
pciBusID: 0000:82:00.0
totalMemory: 11.17GiB freeMemory: 11.09GiB
2018-03-05 14:38:59.059989: I tensorflow/core/common_runtime/gpu/gpu_device.cc:1312]
↪ Adding visible gpu devices: 0
2018-03-05 14:38:59.496393: I tensorflow/core/common_runtime/gpu/gpu_device.cc:993]
↪ Creating TensorFlow device (/job:localhost/replica:0/task:0/device:GPU:0 with 10757
↪ MB memory) -> physical GPU (device: 0, name: Tesla K40m, pci bus id: 0000:82:00.0,
↪ compute capability: 3.5)
2018-03-05 14:39:23.210323: I tensorflow/core/common_runtime/gpu/gpu_device.cc:1312]
↪ Adding visible gpu devices: 0
2018-03-05 14:39:23.210672: I tensorflow/core/common_runtime/gpu/gpu_device.cc:993]
↪ Creating TensorFlow device (/job:localhost/replica:0/task:0/device:GPU:0 with 260
↪ MB memory) -> physical GPU (device: 0, name: Tesla K40m, pci bus id: 0000:82:00.0,
↪ compute capability: 3.5)

real    0m32.394s
user    0m35.192s
sys     0m12.204s
I0305 15:39:25.630180 4680 executor.cpp:938] Command exited with status 0 (pid: 4716)

```

Test #2

Job definition:

```

{
  "name": "test2-gpu",
  "command": "cd $MESOS_SANDBOX && /bin/bash gpu_demo.sh",

```

(continues on next page)

(continued from previous page)

```

"shell": true,
"retries": 2,
"cpus": 4,
"disk": 256,
"mem": 4096,
"gpus": 1,
"uris": [
  "https://gist.githubusercontent.com/maricaantonacci/
↪1a7f02903513e7bba91f451e0f4f5ead/raw/78c737fd0e2a288a2040c192368f6c4ecf8eb88a/gpu_
↪demo.sh"
],
"environmentVariables": [],
"arguments": [],
"runAsUser": "root",
"container": {
  "type": "MESOS",
  "image": "tensorflow/tensorflow:1.4.0-gpu"
},
"schedule": "R/2018-03-05T23:00:00.000Z/PT24H"
}

```

As you can see, the only difference wrt Test#1 is the docker image: here we are using the tag 1.4.0-gpu of the tensorflow docker image that has been built using a different CUDA and cuDNN version.

Also in this case the job is correctly run:

```

Cloning into 'TensorFlow-Examples'...
WARNING:tensorflow:Using temporary folder as model directory: /tmp/tmpomIcq9
2018-03-05 16:36:24.518455: I tensorflow/core/platform/cpu_feature_guard.cc:137] Your
↪CPU supports instructions that this TensorFlow binary was not compiled to use: SSE4.
↪1 SSE4.2 AVX
2018-03-05 16:36:25.261578: I tensorflow/core/common_runtime/gpu/gpu_device.cc:1030]
↪Found device 0 with properties:
name: Tesla K40m major: 3 minor: 5 memoryClockRate(GHz): 0.745
pciBusID: 0000:82:00.0
totalMemory: 11.17GiB freeMemory: 11.09GiB
2018-03-05 16:36:25.261658: I tensorflow/core/common_runtime/gpu/gpu_device.cc:1120]
↪Creating TensorFlow device (/device:GPU:0) -> (device: 0, name: Tesla K40m, pci bus
↪id: 0000:82:00.0, compute capability: 3.5)
2018-03-05 16:36:52.299346: I tensorflow/core/common_runtime/gpu/gpu_device.cc:1120]
↪Creating TensorFlow device (/device:GPU:0) -> (device: 0, name: Tesla K40m, pci bus
↪id: 0000:82:00.0, compute capability: 3.5)

real    0m35.273s
user    0m37.828s
sys 0m13.164s
I0305 17:36:54.803642 7405 executor.cpp:938] Command exited with status 0 (pid: 7439)

```

Additional note:

Running the job without gpu (using the image tensorflow/tensorflow:latest and “gpus”: 0) we got for the same script:

```


real    2m15.647s
user    22m33.384s
sys 15m51.164s

```

Testing GPU support in Marathon

In order to enable GPU support in Marathon, you need to start the framework with the commandline option `--enable_features=gpu_resources` (or using the env variable `MARATHON_ENABLE_FEATURES`):

Start Marathon Framework:

```
docker run -d --name marathon --net host --env-file /etc/marathon/.marathonenv  indigodatacloud/marathon:1.5.6
```

with the following enviroment:

```
LIBPROCESS_IP=<mesos-master-ip>
MARATHON_HOSTNAME=<mesos-master-fqdn/ip>
MARATHON_HTTP_ADDRESS=<mesos-master-ip>
MARATHON_HTTP_PORT=8080
MARATHON_MASTER=zk://<mesos-master>:2181/mesos
MARATHON_ZK=zk://<mesos-master>:2181/marathon
MARATHON_FRAMEWORK_NAME=marathon
MARATHON_HTTP_CREDENTIALS=admin:*****
MARATHON_ENABLE_FEATURES=gpu_resources
```

Test:

The following application has been submitted to Marathon:

```
{
  "id": "tensorflow-gpus",
  "cpus": 4,
  "gpus": 1,
  "mem": 2048,
  "disk": 0,
  "instances": 1,
  "container": {
    "type": "MESOS",
    "docker": {
      "image": "tensorflow/tensorflow:latest-gpu"
    }
  }
}
```

Running tensorflow docker container

1) Using “DOCKER” containerizer on a Mesos cluster without GPUs

Submit to Marathon the following application:

```
{
  "id": "/tensorflow-app",
  "cmd": "PORT=8888 /run_jupyter.sh --allow-root",
  "cpus": 2,
  "mem": 4096,
  "instances": 1,
  "container": {
    "type": "DOCKER",
    "docker": {
      "image": "tensorflow/tensorflow:latest",

```

(continues on next page)

(continued from previous page)

```

    "network": "BRIDGE",
    "portMappings": [
      {
        "containerPort": 8888,
        "hostPort": 0,
        "servicePort": 10000,
        "protocol": "tcp"
      }
    ],
    "privileged": false,
    "forcePullImage": true
  }
},
"env": {
  "PASSWORD": "s3cret"
},
"labels": {
  "HAPROXY_GROUP": "external"
}
}

```

Then you can access the service through the cluster LB on port 10000 (servicePort)

2) Using “MESOS” containerizer on a Mesos cluster with GPUs

```

{
  "id": "tensorflow-gpus",
  "cpus": 4,
  "gpus": 1,
  "mem": 2048,
  "disk": 0,
  "instances": 1,
  "container": {
    "type": "MESOS",
    "docker": {
      "image": "tensorflow/tensorflow:latest-gpu"
    }
  },
  "portDefinitions": [
    { "port": 10000, "name": "http" }
  ],
  "networks": [ { "mode": "host" } ],
  "labels": {
    "HAPROXY_GROUP": "external"
  },
  "env": {
    "PASSWORD": "s3cret"
  }
}

```

Then you can access the service through the cluster LB on port 10000.

If the “port” field in portDefinitions is set to 0 then Marathon will assign a random service port (that you can know with a GET request to /v2/apps/app-name)

References

Enabling open-id connect authentication

Mesos/Marathon/Chronos do not support open-id connect authentication natively.

A very simple solution is to front the mesos cluster with an Apache server that itself is capable of negotiating authentication for users.

The following configuration can be used to setup a reverse proxy that uses the module *mod_auth_openidc*:

```

ServerName mesos.example.com

<VirtualHost *:443>
  ServerName mesos.example.com

  LoadModule auth_openidc_module /usr/lib/apache2/modules/mod_auth_openidc.so

  OIDCClaimPrefix          "OIDC-"
  OIDCResponseType          "code"
  OIDCScope                 "openid email profile"
  OIDCProviderMetadataURL   https://iam.deep-hybrid-datacloud.eu/.well-known/
  ↪openid-configuration
  OIDCClientID              332e618b-d3bf-440d-aea1-6da2823aaece # replace with_
  ↪your client ID
  OIDCClientSecret          **** # replace with_
  ↪your client secret
  OIDCProviderTokenEndpointAuth client_secret_basic
  OIDCCryptoPassphrase      **** # replace with_
  ↪your passphrase
  OIDCRedirectURI           https://mesos.example.com/mesos/redirect_uri

  OIDCOAuthVerifyJwksUri "https://iam.deep-hybrid-datacloud.eu/jwk"

  <Location /mesos>
    AuthType openid-connect
    Require valid-user
    LogLevel debug
  </Location>

  <Location /marathon>
    AuthType oauth20
    Require valid-user
    LogLevel debug
    RequestHeader set Authorization "Basic YWRtaC46bTNzb3NNLjIwMTY="
  </Location>

  <Location /chronos>
    AuthType oauth20
    Require valid-user
    LogLevel debug
    RequestHeader set Authorization "Basic YWRtaC46bTNzb3NDLjIwMTY="
  </Location>

  ProxyTimeout 1200
  ProxyRequests Off
  ProxyPreserveHost Off

  ProxyPass /mesos/ http://172.20.30.40:5050/
  ProxyPassReverse /mesos/ http://172.20.30.40:5050/

```

(continues on next page)

(continued from previous page)

```

ProxyPass /marathon/ http://172.20.30.40:8080/
ProxyPassReverse /marathon/ http://172.20.30.40:8080/

ProxyPass /chronos/ http://172.20.30.40:4400/
ProxyPassReverse /chronos/ http://172.20.30.40:4400/

RemoteIPHeader X-Forwarded-For

## Logging
ErrorLog "/var/log/apache2/proxy_mesos_error_ssl.log"
ServerSignature Off
CustomLog "/var/log/apache2/proxy_mesos_access_ssl.log" combined

## SSL directives

SSLProxyEngine on
SSLEngine on
SSLCertificateFile      "/etc/letsencrypt/live/mesos.example.com/fullchain.pem"
SSLCertificateKeyFile    "/etc/letsencrypt/live/mesos.example.com/privkey.pem"
</VirtualHost>

```

Note that Line 30 is needed if you have enabled basic HTTP authentication to protect your endpoints (in the example above, username/password authentication has been enable for Marathon).

In this case you need to add the Authorization header in the request to the backend. The hash can be computed with the following python script:

```

import base64
hash = base64.b64encode(b'user:password')

```

Once the proxy is up and running you can contact the cluster API endpoints using the IAM (open-id connect) token:

Marathon API endpoint: <https://mesos.example.com/marathon>

Chronos API endpoint: <https://mesos.example.com/chronos>

For example:

```

curl -H "Authorization: bearer $IAM_ACCESS_TOKEN" -X GET https://mesos.example.com/
↪marathon/v2/apps

```

If you want to allow users to access also the Web interfaces of Marathon and Chronos, then add the following configuration:

```

<Location /marathon-web>
  AuthType openid-connect
  Require valid-user
  LogLevel debug
  RequestHeader set Authorization "Basic YWRtaC46bTNzb3NNLjIwMTY="
</Location>

<Location /chronos-web>
  AuthType openid-connect
  Require valid-user
  LogLevel debug
  RequestHeader set Authorization "Basic YWRtaZ46bTNzb3NDLjIwMTY="
</Location>

```

(continues on next page)

(continued from previous page)

```
ProxyPass /marathon-web/ http://172.20.30.40:8080/
ProxyPassReverse /marathon-web/ http://172.20.30.40:8080/

ProxyPass /chronos-web/ http://172.20.30.40:4400/
ProxyPassReverse /chronos-web/ http://172.20.30.40:4400/
```

The Web UIs will be accessible at the following urls:

Marathon Web UI: <https://mesos.example.com/marathon-web/>

Chronos Web UI: <https://mesos.example.com/chronos-web/>

3.1.2 Kubernetes

DEEP : Installing and testing GPU Node in Kubernetes - CentOS7

- *Introduction*
- *Cluster Status*
- *Tests*
 - *Test #1 - Simple vector-add (CUDA8)*
 - *Test #2 - Simple vector-add with different CUDA driver*
 - *Test #3 - Simple BinomialOption with different CUDA driver*
 - *Test #4 - Job Scheduling features*
- *Access PODs from outside the cluster*
 - *Test #5 - Long running service with NodePort*
- *References*
 - *Kubernetes Installation/Configuration docs*
 - *GPU Node - CentOS7.4*
 - * *Install docker*
 - * *Driver nvidia-install repo*
 - * *Driver Nvidia-install driver*
 - * *Install NVIDIA-docker repo*
 - * *Install-NVIDIA docker*
- *Related articles*

Introduction

The **manual** procedure for installation and configuration of a Kubernetes cluster is provided. The cluster is composed by a Master node and one Worker node

Cluster Status

Cluster Name: KubeDeep

Kubectl components

```
# kubectl get componentstatuses
NAME                STATUS    MESSAGE              ERROR
controller-manager  Healthy   ok
scheduler           Healthy   ok
etcd-0              Healthy   {"health": "true"}
```

```
# kubectl get nodes
NAME        STATUS    ROLES    AGE      VERSION
Node1_GPU   Ready     <none>    13d      v1.10.0
```

Worker GPU specifications

```
# nvidia-smi
Mon Apr  2 23:13:37 2018

+-----+
| NVIDIA-SMI 390.30                 Driver Version: 390.30          |
+-----+-----+
| GPU   Name           Persistence-M | Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap |      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|    0  Tesla K40m        On         | 00000000:02:00.0 Off |             0         |
| N/A   30C    P8        20W / 235W |      0MiB / 11441MiB |      0%      Default |
+-----+-----+
|    1  Tesla K40m        On         | 00000000:84:00.0 Off |             0         |
| N/A   32C    P8        20W / 235W |      0MiB / 11441MiB |      0%      Default |
+-----+-----+

+-----+
| Processes:                                                                  GPU Memory |
|  GPU       PID    Type Process name                               Usage      |
+-----+-----+
| No running processes found                                                  |
+-----+
```

Tests

Test #1 - Simple vector-add (CUDA8)

This CUDA Runtime API sample is a very basic sample that implements element by element vector addition. The examples uses CUDA8 driver.

```
#cat vector-add.yaml
apiVersion: v1
kind: Pod
metadata:
  name: vector-add
spec:
  restartPolicy: OnFailure
  containers:
    - name: cuda-vector-add
```

(continues on next page)

(continued from previous page)

```
# https://github.com/kubernetes/kubernetes/blob/v1.7.11/test/images/nvidia-cuda/
↪Dockerfile
image: "k8s.gcr.io/cuda-vector-add:v0.1"
resources:
  limits:
    nvidia.com/gpu: 1
```

```
# kubectl apply -f vector-add.yaml
pod "vector-add" created

# kubectl get pods --show-all
NAME                READY    STATUS    RESTARTS   AGE
vector-add          0/1      Completed 0           4s
```

Test #2 - Simple vector-add with different CUDA driver

This CUDA Runtime API sample is a very basic sample that implements element by element vector addition. The examples uses two Docker images with different version of CUDA driver. To complete the test, a new Docker image with CUDA driver version 9 has been built and uploaded in a private repo.

```
# cat cuda8-vector-add.yaml
apiVersion: v1
kind: Pod
metadata:
  name: cuda8-vector-add
spec:
  restartPolicy: OnFailure
  containers:
    - name: cuda-vector-add
      # https://github.com/kubernetes/kubernetes/blob/v1.7.11/test/images/nvidia-cuda/
      ↪Dockerfile
      image: "k8s.gcr.io/cuda-vector-add:v0.1"
      resources:
        limits:
          nvidia.com/gpu: 1

# cat cuda9-vector-add.yaml
apiVersion: v1
kind: Pod
metadata:
  name: cuda9-vector-add
spec:
  restartPolicy: OnFailure
  containers:
    - name: cuda-vector-add
      image: <private repo>/deep/cuda-vector-add:v0.2
      resources:
        limits:
          nvidia.com/gpu: 1
```

```
# kubectl apply -f cuda9-vector-add.yaml -f cuda8-vector-add.yaml
pod "cuda9-vector-add" created
pod "cuda8-vector-add" created
```

(continues on next page)

(continued from previous page)

```
# kubectl get pods --show-all
NAME                READY    STATUS    RESTARTS   AGE
cuda8-vector-add    0/1      Completed 0          2s
cuda9-vector-add    0/1      Completed 0          2s
```

Test #3 - Simple BinomialOption with different CUDA driver

This sample evaluates fair call price for a given set of European options under binomial model. To complete the test, two new Docker images with CUDA8 and CUDA9 has been built and uploaded in a private repo. The test will take some seconds and GPU engage can be shown

```
# cat cuda8-binomialoption.yaml
apiVersion: v1
kind: Pod
metadata:
  name: cuda8-binomialoption
spec:
  restartPolicy: OnFailure
  containers:
  - name: cuda8-binomilaoption
    image: <private_repo>/deep/cuda-binomialoption:v0.1
    resources:
      limits:
        nvidia.com/gpu: 1

# cat cuda9-binomialoption.yaml
apiVersion: v1
kind: Pod
metadata:
  name: cuda9-binomialoption
spec:
  restartPolicy: OnFailure
  containers:
  - name: cuda9-binomialoption
    image: <private_repo>/deep/cuda-binomialoption:v0.2
    resources:
      limits:
        nvidia.com/gpu: 1
```

```
# kubectl apply -f cuda8-binomialoption.yaml -f cuda9-binomialoption.yaml
pod "cuda8-binomialoption" created
pod "cuda9-binomialoption" created

# kubectl get pods --show-all
NAME                READY    STATUS    RESTARTS   AGE
cuda8-binomialoption 1/1      Running   0          2s
cuda9-binomialoption 1/1      Running   0          2s

# kubectl get pods --show-all
NAME                READY    STATUS    RESTARTS   AGE
cuda8-binomialoption 1/1      Running   0         22s
cuda9-binomialoption 1/1      Running   0         22s

# kubectl get pods --show-all
```

(continues on next page)

(continued from previous page)

NAME	READY	STATUS	RESTARTS	AGE
cuda8-binomialoption	0/1	Completed	0	1m
cuda9-binomialoption	0/1	Completed	0	1m

```
# nvidia-smi
Mon Apr  2 23:35:17 2018
```

```

+-----+
| NVIDIA-SMI 390.30                 Driver Version: 390.30           |
+-----+-----+
| GPU  Name      Persistence-M | Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap |      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|   0   Tesla K40m        On   | 00000000:02:00.0 Off |                    0 |
| N/A   31C    P0      63W / 235W |      80MiB / 11441MiB |      0%      Default |
+-----+-----+
|   1   Tesla K40m        On   | 00000000:84:00.0 Off |                    0 |
| N/A   33C    P0      63W / 235W |      80MiB / 11441MiB |      0%      Default |
+-----+-----+

```

```

+-----+
| Processes:                                                       GPU Memory |
|  GPU       PID    Type  Process name                               Usage      |
+-----+-----+
|      0      3385    C      ./binomialOptions                             69MiB      |
|      1      3369    C      ./binomialOptions                             69MiB      |
+-----+

```

Test #4 - Job Scheduling features

Tests highlighting the features of the Kubernetes scheduler. Default schedule policies are used (FIFO).

Submission of a bunch of different cuda jobs with different running time.

- Parrec (1h)
- Cuda8-binomialoption.yaml (5 min)
- Cuda9-binomialoption.yaml (5 min)
- Cuda8-vector-add.yaml (few sec)
- Cuda9-vector-add.yaml (few sec)

The parrec job has been launched as first job. One GPU has been engaged by the job; the other is still available for other jobs.

```
# kubectl apply -f parrec.yaml
pod "parrec" created

# kubectl get pods -o wide
NAME      READY   STATUS    RESTARTS   AGE   IP           NODE
parrec    1/1     Running   0           22s   172.30.0.52  gpu-node-01
```

Other jobs have been submitted in the following order:

```
# kubectl apply -f cuda8-binomialoption.yaml -f cuda9-binomialoption.yaml -f cuda8-
↪vector-add.yaml -f cuda9-vector-add.yaml
pod "cuda8-binomialoption" created
```

(continues on next page)

(continued from previous page)

```

pod "cuda9-binomialoption" created
pod "cuda8-vector-add" created
pod "cuda9-vector-add" created

# kubectl get pods -o wide
NAME                                READY    STATUS    RESTARTS   AGE    IP             NODE
cuda8-binomialoption                1/1     Running   0           4s     172.30.0.53    gpu-node-
↳01
cuda8-vector-add                    0/1     Pending   0           4s     <none>         <none>
cuda9-binomialoption                0/1     Pending   0           4s     <none>         <none>
cuda9-vector-add                    0/1     Pending   0           4s     <none>         <none>
parrec                              1/1     Running   0           1m     172.30.0.52    gpu-node-
↳01

```

The “cuda8-binomialoption” is running, the other are in the FIFO queue in pending state. After completion, the other job will be running in the same order they have been submitted.

```

# kubectl get pods -o wide

NAME                                READY    STATUS    RESTARTS   AGE    IP             NODE
cuda8-binomialoption                1/1     Running   0           31s    172.30.0.53    gpu-node-
↳01
cuda8-vector-add                    0/1     Pending   0           31s    <none>         <none>
cuda9-binomialoption                0/1     Pending   0           31s    <none>         <none>
cuda9-vector-add                    0/1     Pending   0           31s    <none>         <none>
parrec                              1/1     Running   0           2m     172.30.0.52    gpu-node-
↳01

# kubectl get pods -o wide
NAME                                READY    STATUS    RESTARTS   AGE    IP             NODE
cuda8-binomialoption                0/1     Completed 0           49s    172.30.0.53    gpu-
↳node-01
cuda8-vector-add                    0/1     Pending   0           49s    <none>         <none>
cuda9-binomialoption                0/1     Pending   0           49s    <none>         <none>
cuda9-vector-add                    0/1     Pending   0           49s    <none>         <none>
parrec                              1/1     Running   0           2m     172.30.0.52    gpu-
↳node-01

# kubectl get pods -o wide
NAME                                READY    STATUS    RESTARTS   AGE    IP             NODE
↳NODE
cuda8-binomialoption                0/1     Completed 0           1m     172.30.0.53    ↳
↳gpu-node-01
cuda8-vector-add                    0/1     Pending   0           1m     <none>
↳<none>
cuda9-binomialoption                0/1     ContainerCreating 0           1m     <none>    ↳
↳gpu-node-01
cuda9-vector-add                    0/1     Pending   0           1m     <none>
↳<none>
parrec                              1/1     Running   0           2m     172.30.0.52    ↳
↳gpu-node-01

# kubectl get pods -o wide
NAME                                READY    STATUS    RESTARTS   AGE    IP             NODE
cuda8-binomialoption                0/1     Completed 0           1m     172.30.0.53    gpu-
↳node-01

```

(continues on next page)

(continued from previous page)

cuda8-vector-add	0/1	Pending	0	1m	<none>	<none>
cuda9-binomialoption	1/1	Running	0	1m	172.30.0.54	gpu-
↪node-01						
cuda9-vector-add	0/1	Pending	0	1m	<none>	<none>
parrec	1/1	Running	0	2m	172.30.0.52	gpu-
↪node-01						
# kubectl get pods -o wide						
NAME	READY	STATUS	RESTARTS	AGE	IP	NODE
cuda8-binomialoption	0/1	Completed	0	2m	172.30.0.53	gpu-
↪node-01						
cuda8-vector-add	0/1	Completed	0	2m	172.30.0.55	gpu-
↪node-01						
cuda9-binomialoption	0/1	Completed	0	2m	172.30.0.54	gpu-
↪node-01						
cuda9-vector-add	0/1	Pending	0	2m	<none>	<none>
parrec	1/1	Running	0	3m	172.30.0.52	gpu-
↪node-01						
# kubectl get pods -o wide						
NAME	READY	STATUS	RESTARTS	AGE	IP	NODE
cuda8-binomialoption	0/1	Completed	0	2m	172.30.0.53	gpu-
↪node-01						
cuda8-vector-add	0/1	Completed	0	2m	172.30.0.55	gpu-
↪node-01						
cuda9-binomialoption	0/1	Completed	0	2m	172.30.0.54	gpu-
↪node-01						
cuda9-vector-add	0/1	Completed	0	2m	172.30.0.56	gpu-
↪node-01						
parrec	1/1	Running	0	4m	172.30.0.52	gpu-
↪node-01						

Access PODs from outside the cluster

To access PODs from outside the cluster it can be possible following different procedures that strictly depend on the usecase and (cloud) providers.

NodePort, **hostNetwork**, **hostPort**, **LoadBalancer** and **Ingress** features of Kubernetes can be adopted as described in the following Reference:

<http://alesnosek.com/blog/2017/02/14/accessing-kubernetes-pods-from-outside-of-the-cluster/>

For example purposes, the **Test #5** example will describe and use the **NodePort** procedure as the cluster is defined as 1 Master and 1 Worker both with routable IPs.

Test #5 - Long running service with NodePort

Prerequisites

1. Kubernetes Node with routable IP
2. Port range dynamically selected as from the “kubernetes-apiservice.service” configuration file
3. Nginx replica 2; V. 1.13.12 - latest (as from the YAML files)

Yaml files related to nginx deployment and nginx service

```
# cat nginx.deploy.yaml
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: nginx-example
  namespace: default
  labels:
    app: nginx
spec:
  replicas: 2
  strategy:
    type: RollingUpdate
    rollingUpdate:
      maxSurge: 1
      maxUnavailable: 0
  template:
    metadata:
      labels:
        app: nginx
        name: nginx
    spec:
      containers:
      - image: nginx:latest
        name: ingress-example
        ports:
        - name: http
          containerPort: 80
        readinessProbe:
          httpGet:
            path: /
            port: 80
            scheme: HTTP
        livenessProbe:
          httpGet:
            path: /
            port: 80
            scheme: HTTP
          initialDelaySeconds: 5
          timeoutSeconds: 1
```

```
# cat nginx.svc.yaml
apiVersion: v1
kind: Service
metadata:
  name: nginx
  namespace: default
spec:
  type: NodePort
  ports:
  - name: http
    port: 80
    targetPort: 80
    protocol: TCP
  selector:
    app: nginx
```

Creation of nginx POD and nginx service. The following commands will return the Node hostname and the port associated to the nginx.

```
# kubectl apply -f nginx.deploy.yaml -f nginx.svc.yaml
deployment.extensions "nginx-example" created
service "nginx" created
# kubectl get pods
```

NAME	READY	STATUS	RESTARTS	AGE
nginx-example-78847794b7-8nm8t	0/1	Running	0	11s
nginx-example-78847794b7-n8nxs	0/1	Running	0	11s

```
# kubectl get pods
```

NAME	READY	STATUS	RESTARTS	AGE
nginx-example-78847794b7-8nm8t	1/1	Running	0	30s
nginx-example-78847794b7-n8nxs	1/1	Running	0	30s

```
# kubectl get svc
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
nginx	NodePort	192.168.0.130	<none>	80:30916/TCP	51s

Test of nginx

```
# curl http://gpu-node-01:30916
<!DOCTYPE html>
<html>
<head>
<title>Welcome to nginx!</title>
<style>
...
```

Delete one of the two PODs from nginx

```
# kubectl delete pod nginx-example-78847794b7-8nm8t
pod "nginx-example-78847794b7-8nm8t" deleted
```

A new POD is creating while the old POD is getting deleted. No service downtime is registered from the user

```
# kubectl get pods
```

NAME	READY	STATUS	RESTARTS	AGE
nginx-example-78847794b7-6gvnn	0/1	Running	0	4s
nginx-example-78847794b7-8nm8t	0/1	Terminating	0	12m
nginx-example-78847794b7-n8nxs	1/1	Running	0	12m

```
# kubectl get pods
```

NAME	READY	STATUS	RESTARTS	AGE
nginx-example-78847794b7-6gvnn	0/1	Running	0	12s
nginx-example-78847794b7-n8nxs	1/1	Running	0	12m

```
# kubectl get pods -o wide
```

NAME	READY	STATUS	RESTARTS	AGE	IP
↪ NODE					
nginx-example-78847794b7-6gvnn	1/1	Running	0	28s	172.30.0.61
↪ gpu-node-01					
nginx-example-78847794b7-n8nxs	1/1	Running	0	12m	172.30.0.59
↪ gpu-node-01					

Delete the seconf POD. The internal POD IP is changing, but the public endpoint is the same.

```
# kubectl delete pod nginx-example-78847794b7-n8nxs
pod "nginx-example-78847794b7-n8nxs" deleted
```

(continues on next page)

(continued from previous page)

```
# kubectl get pods -o wide
NAME                                READY   STATUS    RESTARTS   AGE   IP
↪   NODE
nginx-example-78847794b7-2szlv      0/1     Running   0           4s    172.30.0.
↪62   gpu-node-01
nginx-example-78847794b7-6gvnn      1/1     Running   0           50s    172.30.0.
↪61   gpu-node-01
nginx-example-78847794b7-n8nxs      0/1     Terminating 0           13m    172.30.0.
↪59   gpu-node-01
# kubectl get pods -o wide
NAME                                READY   STATUS    RESTARTS   AGE   IP
↪   NODE
nginx-example-78847794b7-2szlv      1/1     Running   0           13s    172.30.0.62
↪   gpu-node-01
nginx-example-78847794b7-6gvnn      1/1     Running   0           59s    172.30.0.61
↪   gpu-node-01
```

Changing the version of nginx with an older version (v. 1.12).

```
# cat nginx.deploy.yaml
apiVersion: extensions/v1beta1
...
  containers:
  - image: nginx:1.12
```

Apply changes. New PODs are created while old PODs are deleted. No nginx downtime is registered from the user. Public endpoint and port remain unchanged.

```
# kubectl apply -f nginx.deploy.yaml -f nginx.svc.yaml
deployment.extensions "nginx-example" configured
service "nginx" unchanged

# kubectl get pods
NAME                                READY   STATUS    RESTARTS   AGE
nginx-example-5d9764f848-8kc8b      0/1     Running   0           9s
nginx-example-78847794b7-2szlv      1/1     Running   0           2m
nginx-example-78847794b7-6gvnn      1/1     Running   0           2m

# kubectl get pods -o wide
NAME                                READY   STATUS    RESTARTS   AGE   IP
↪   NODE
nginx-example-5d9764f848-8kc8b      1/1     Running   0           23s    172.30.0.
↪63   gpu-node-01
nginx-example-5d9764f848-xwr77      1/1     Running   0           7s     172.30.0.
↪64   gpu-node-01
nginx-example-78847794b7-6gvnn      0/1     Terminating 0           3m     172.30.0.
↪61   gpu-node-01

# kubectl get pods -o wide
NAME                                READY   STATUS    RESTARTS   AGE   IP
↪   NODE
nginx-example-5d9764f848-8kc8b      1/1     Running   0           54s    172.30.0.63
↪   gpu-node-01
nginx-example-5d9764f848-xwr77      1/1     Running   0           38s    172.30.0.64
↪   gpu-node-01
```


References

The following guides have been followed for the Installation and Configuration of Kubernetes cluster - a detailed step-by-step guide will be provided soon

Kubernetes Installation/Configuration docs

<https://github.com/kelseyhightower/kubernetes-the-hard-way>

GPU Node - CentOS7.4

Install docker

<https://docs.docker.com/install/linux/docker-ce/centos/#set-up-the-repository>

Driver nvidia-install repo

https://developer.nvidia.com/cuda-downloads?target_os=Linux

Driver Nvidia-install driver

<http://docs.nvidia.com/cuda/cuda-installation-guide-linux/index.html#post-installation-actions>

Install NVIDIA-docker repo

<https://nvidia.github.io/nvidia-docker/>

Install-NVIDIA docker

[https://github.com/NVIDIA/nvidia-docker/wiki/Installation-\(version-2.0\)#prerequisites](https://github.com/NVIDIA/nvidia-docker/wiki/Installation-(version-2.0)#prerequisites)

Installing GPU node and adding it to Kubernetes cluster

This is a guide on how to install a GPU node and join it in a running Kubernetes cluster deployed with kubeadm. The guide was tested on a Kubernetes cluster v1.9.4 installed with kubeadm. The cluster nodes are KVM virtual machines deployed by OpenStack. VMs are running Ubuntu 16.04.4 LTS. The node with GPU has a single NVIDIA K20m GPU card.

Step-by-step guide

1. We start with a blank node with a GPU. This is the node, we would like to join in our Kubernetes cluster. First, update the node and install graphic drivers. The version of the drivers has to be at least 361.93. We have installed version 387.26 and CUDA Version 8.0.61. Drivers and CUDA installation is not a part of this guide.

NVIDIA drivers information Expand source

```
ubuntu@virtual-kubernetes-gpu-2:~$ nvidia-smi
Wed Mar 14 08:52:53 2018

+-----+
| NVIDIA-SMI 387.26                  Driver Version: 387.26                  |
+-----+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+-----+
|    0   Tesla K20m        Off      | 00000000:00:07.0 Off |             0         |
| N/A   30C    P0       53W / 225W |      0MiB /  4742MiB |    100%    Default   |
+-----+-----+-----+

+-----+
| Processes:                                                       GPU Memory |
|  GPU       PID    Type    Process name                       Usage    |
+-----+-----+
| No running processes found                                     |
+-----+
```

CUDA information Expand source

```
ubuntu@virtual-kubernetes-gpu-2:~$ cat /usr/local/cuda-8.0/version.txt
CUDA Version 8.0.61
```

2. The next step is to install Docker on the GPU node. Install Docker CE 17.03 from Docker's repositories for Ubuntu. Proceed with the following commands as a root user.

```
sudo apt-get update
sudo apt-get install -y \
    apt-transport-https \
    ca-certificates \
    curl \
    software-properties-common
sudo curl -fsSL https://download.docker.com/linux/ubuntu/gpg | apt-key add -
sudo add-apt-repository \
    "deb https://download.docker.com/linux/$(. /etc/os-release; echo "$ID") \
    $(lsb_release -cs) \
    stable"
sudo apt-get update && apt-get install -y docker-ce=$(apt-cache madison docker-ce_
→| grep 17.03 | head -1 | awk '{print $3}')
```

Docker installation test Expand source

```
root@virtual-kubernetes-gpu-2:~# docker --version
Docker version 17.03.2-ce, build f5ec1e2
root@virtual-kubernetes-gpu-2:~# docker run hello-world
Unable to find image 'hello-world:latest' locally
latest: Pulling from library/hello-world
ca4f61b1923c: Pull complete
Digest: sha256:97ce6fa4b6cdc0790cda65fe7290b74cfebd9fa0c9b8c38e979330d547d22ce1
Status: Downloaded newer image for hello-world:latest

Hello from Docker!
This message shows that your installation appears to be working correctly.

To generate this message, Docker took the following steps:
 1. The Docker client contacted the Docker daemon.
```

(continues on next page)

(continued from previous page)

2. The Docker daemon pulled the "hello-world" image from the Docker Hub. (amd64)
3. The Docker daemon created a new container from that image which runs the executable that produces the output you are currently reading.
4. The Docker daemon streamed that output to the Docker client, which sent it to your terminal.

To try something more ambitious, you can run an Ubuntu container with:

```
$ docker run -it ubuntu bash
```

Share images, automate workflows, and more with a free Docker ID:

<https://cloud.docker.com/>

For more examples and ideas, visit:

<https://docs.docker.com/engine/userguide/>

3. On the GPU node, add nvidia-docker2 package repositories, install it and reload Docker daemon configuration, which might be altered by nvidia-docker2 installation.

```
sudo curl -s -L https://nvidia.github.io/nvidia-docker/gpgkey | \
  sudo apt-key add -
sudo curl -s -L https://nvidia.github.io/nvidia-docker/ubuntu16.04/amd64/nvidia-
  docker.list | \
  sudo tee /etc/apt/sources.list.d/nvidia-docker.list
sudo apt-get update
sudo apt-get install -y nvidia-docker2
sudo pkill -SIGHUP dockerd
```

nvidia-docker2 GPU test Expand source

```
root@virtual-kubernetes-gpu-2:~# docker run --runtime=nvidia --rm nvidia/cuda_
  nvidia-smi
Unable to find image 'nvidia/cuda:latest' locally
latest: Pulling from nvidia/cuda
22dc81ace0ea: Pull complete
1a8b3c87dba3: Pull complete
91390a1c435a: Pull complete
07844b14977e: Pull complete
b78396653dae: Pull complete
95e837069dfa: Pull complete
fef4aadda783: Pull complete
343234bd5cf3: Pull complete
64e8786fc8c1: Pull complete
d6a4723d353c: Pull complete
Digest: sha256:3524adf9b563c27d9a0f6d0584355c1f4f4b38e90b66289b8f8de026a9162eee
Status: Downloaded newer image for nvidia/cuda:latest
Wed Mar 14 10:14:51 2018
```

```
+-----+
| NVIDIA-SMI 387.26                  Driver Version: 387.26                  |
+-----+-----+-----+-----+-----+
| GPU  Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+-----+-----+-----+
|   0   Tesla K20m        Off      | 00000000:00:07.0 Off |                    0 |
| N/A   30C    P0        52W / 225W |      0MiB /  4742MiB |      100%    Default |
+-----+-----+-----+-----+-----+
```

(continues on next page)

(continued from previous page)

```

+-----+
| Processes:                                     GPU Memory |
| GPU      PID    Type    Process name                      Usage          |
+=====+
| No running processes found                      |
+-----+

```

4. Set `nvidia-runtime` as the default runtime for Docker on the GPU node. Edit the `/etc/docker/daemon.json` configuration file and set the "default-runtime" parameter to `nvidia`. This also allows us to omit the `--runtime=nvidia` parameter for Docker.

```

{
  "default-runtime": "nvidia",
  "runtimes": {
    "nvidia": {
      "path": "/usr/bin/nvidia-container-runtime",
      "runtimeArgs": []
    }
  }
}

```

5. As a root user on the GPU node, add Kubernetes package repositories and install `kubeadm`, `kubect1` and `kubelet`. Then turn the swap off as it is not supported by Kubernetes.

```

apt-get update && apt-get install -y apt-transport-https
curl -s https://packages.cloud.google.com/apt/doc/apt-key.gpg | apt-key add -
cat <<EOF >/etc/apt/sources.list.d/kubernetes.list
deb http://apt.kubernetes.io/ kubernetes-xenial main
EOF
apt-get update
apt-get install -y kubelet kubeadm kubectl
# turn off swap or comment the swap line in /etc/fstab
sudo swapoff -a

```

Specific version installation; e.g., 1.9.3-00 Expand source

```

# install aptitude, an interface to package manager
root@virtual-kubernetes-gpu-2:~# apt install aptitude -y

# show available kubeadm versions in the repositories
root@virtual-kubernetes-gpu-2:~# aptitude versions kubeadm
Package kubeadm:
p 1.5.7-00 kubernetes-xenial 500
p 1.6.1-00 kubernetes-xenial 500
p 1.6.2-00 kubernetes-xenial 500
...
p 1.9.3-00 kubernetes-xenial 500
p 1.9.4-00 kubernetes-xenial 500

# install specific version of kubelet, kubeadm and kubectl
root@virtual-kubernetes-gpu-2:~# apt-get install -y kubelet=1.9.3-00 kubeadm=1.9.
↪3-00 kubectl=1.9.3-00

```

6. On the GPU node, edit the `/etc/systemd/system/kubelet.service.d/10-kubeadm.conf` file add the following environment argument to enable `DevicePlugins` feature gate. If there is already `Accelerators` feature gate set, remove it.

```
Environment="KUBELET_EXTRA_ARGS=--feature-gates=DevicePlugins=true"
```

/etc/systemd/system/kubelet.service.d/10-kubeadm.conf Expand source

```
[Service]
Environment="KUBELET_KUBECONFIG_ARGS=--bootstrap-kubeconfig=/etc/kubernetes/
↳bootstrap-kubelet.conf --kubeconfig=/etc/kubernetes/kubelet.conf"
Environment="KUBELET_SYSTEM_PODS_ARGS=--pod-manifest-path=/etc/kubernetes/
↳manifests --allow-privileged=true"
Environment="KUBELET_NETWORK_ARGS=--network-plugin=cni --cni-conf-dir=/etc/cni/
↳net.d --cni-bin-dir=/opt/cni/bin"
Environment="KUBELET_DNS_ARGS=--cluster-dns=10.96.0.10 --cluster-domain=cluster.
↳local"
Environment="KUBELET_AUTHZ_ARGS=--authorization-mode=Webhook --client-ca-file=/
↳etc/kubernetes/pki/ca.crt"
Environment="KUBELET_CADVISOR_ARGS=--cadvisor-port=0"
Environment="KUBELET_CERTIFICATE_ARGS=--rotate-certificates=true --cert-dir=/var/
↳lib/kubelet/pki"
Environment="KUBELET_EXTRA_ARGS=--feature-gates=DevicePlugins=true"
ExecStart=
ExecStart=/usr/bin/kubelet $KUBELET_KUBECONFIG_ARGS $KUBELET_SYSTEM_PODS_ARGS
↳$KUBELET_NETWORK_ARGS $KUBELET_DNS_ARGS $KUBELET_AUTHZ_ARGS $KUBELET_CADVISOR_
↳ARGS $KUBELET_CERTIFICATE_ARGS $KUBELET_EXTRA_ARGS
```

7. On the GPU node, reload and restart kubelet to apply previous changes to the configuration.

```
sudo systemctl daemon-reload
sudo systemctl restart kubelet
```

8. If not already done, enable GPU support on the Kubernetes master by deploying following Daemonset.

```
kubect1 create -f https://raw.githubusercontent.com/NVIDIA/k8s-device-plugin/v1.9/
↳nvidia-device-plugin.yml
```

9. For the simplicity, generate a new token on the Kubernetes master and print the join command.

```
ubuntu@virtual-kubernetes-1:~$ sudo kubeadm token create --print-join-command
kubeadm join --token 6e112b.a598ccc2e90671a6 KUBERNETES_MASTER_IP:6443 --
↳discovery-token-ca-cert-hash
↳sha256:863250f81355e64074cedf5e3486af32253e394e939f4b03562e4ec87707de0a
```

10. Go back to the GPU node and use the printed join command to add GPU node into the cluster.

```
ubuntu@virtual-kubernetes-gpu-2:~$ sudo kubeadm join --token 6e112b.
↳a598ccc2e90671a6 KUBERNETES_MASTER_IP:6443 --discovery-token-ca-cert-hash
↳sha256:863250f81355e64074cedf5e3486af32253e394e939f4b03562e4ec87707de0a
[preflight] Running pre-flight checks.
[WARNING FileExisting-crictl]: crictl not found in system path
[discovery] Trying to connect to API Server "KUBERNETES_MASTER_IP:6443"
[discovery] Created cluster-info discovery client, requesting info from "https://
↳KUBERNETES_MASTER_IP:6443"
[discovery] Requesting info from "https://KUBERNETES_MASTER_IP:6443" again to
↳validate TLS against the pinned public key
[discovery] Cluster info signature and contents are valid and TLS certificate
↳validates against pinned roots, will use API Server "KUBERNETES_MASTER_IP:6443"
[discovery] Successfully established connection with API Server "KUBERNETES_
↳MASTER_IP:6443"
```

(continues on next page)

(continued from previous page)

```

This node has joined the cluster:
* Certificate signing request was sent to master and a response
  was received.
* The Kubelet was informed of the new secure connection details.

Run 'kubectl get nodes' on the master to see this node join the cluster.

```

11. Run following command to see the GPU node (virtual-kubernetes-gpu-2) status on the cluster.

```

ubuntu@virtual-kubernetes-1:~$ kubectl get nodes
NAME                                STATUS    ROLES    AGE      VERSION
virtual-kubernetes-gpu             Ready     <none>    1d       v1.9.4
virtual-kubernetes-gpu-2          NotReady  <none>    13s      v1.9.4
virtual-kubernetes-1              Ready     master    5d       v1.9.4
virtual-kubernetes-2              Ready     <none>    5d       v1.9.4
virtual-kubernetes-3              Ready     <none>    5d       v1.9.4

```

12. After a while, the node is ready.

```

virtual-kubernetes-gpu-2    Ready     <none>    7m       v1.9.4

```

13. Now we have 2 GPU nodes ready in our Kubernetes cluster. We can label the recently added node (virtual-kubernetes-gpu-2) with the accelerator type by running following command on the master.

```

kubectl label nodes virtual-kubernetes-gpu-2 accelerator=nvidia-tesla-k20m

```

14. To check nodes for accelerator label, run `kubectl get nodes -L accelerator` on Kubernetes master.

```

ubuntu@virtual-kubernetes-1:~/kubernetes$ kubectl get nodes -L accelerator
NAME                                STATUS    ROLES    AGE      VERSION    ACCELERATOR
virtual-kubernetes-gpu             Ready     <none>    1d       v1.9.4     nvidia-tesla-
↳k20m
virtual-kubernetes-gpu-2          Ready     <none>    24m      v1.9.4     nvidia-tesla-
↳k20m
virtual-kubernetes-1              Ready     master    5d       v1.9.4
virtual-kubernetes-2              Ready     <none>    5d       v1.9.4
virtual-kubernetes-3              Ready     <none>    5d       v1.9.4

```

15. To test the GPU nodes, go to the master and create a file with the following content and execute it.

gpu-test.yml

```

apiVersion: v1
kind: Pod
metadata:
  name: cuda-vector-add
spec:
  restartPolicy: OnFailure
  containers:
    - name: cuda-vector-add
      # https://github.com/kubernetes/kubernetes/blob/v1.7.11/test/images/nvidia-
↳cuda/Dockerfile
      image: "k8s.gcr.io/cuda-vector-add:v0.1"
      resources:
        limits:
          nvidia.com/gpu: 1 # requesting 1 GPU per container

```

(continues on next page)

(continued from previous page)

```
nodeSelector:
  accelerator: nvidia-tesla-k20m # or nvidia-tesla-k80 etc.
```

```
ubuntu@virutal-kubernetes-1:~/kubernetes$ kubectl create -f gpu-test.yml
pod "cuda-vector-add" created
ubuntu@virutal-kubernetes-1:~/kubernetes$ kubectl get pods -a
NAME                                READY    STATUS    RESTARTS   AGE
cuda-vector-add                     0/1      Completed 0           19s
```

3.1.3 OpenStack nova-lxd

OpenStack nova-lxd installation via Ansible

This document provide step by step deployment of Openstack site with nova-lxd via Openstack Ansible. That would allows site admins to avoid obstacles and pitfalls during deployment and create a Openstack site with nova-lxd for testing and development in a short time without studying extensive documentation of Openstack Ansible.

This documentation also show that nova-lxd is supported (although not perfectly, see steps 10-11 for fixing LXD configuration) in mainstream Openstack deployment tool beside Ubuntu-specific Juju charms. Ubuntu 18.04 and LXD 3.0 are also supported (instead of 16.04 / LXD 2.0 used in Juju).

Comparison between Openstack Ansible and Juju/conjure-up

Juju is Ubuntu- specific, and using Ubuntu distro packages for installing Openstack. Openstack Ansible is distro-neutral and by default it uses source code from github for Openstack installation (configurable for distro packages, however). For testing and development, installation via source code has advantages using latest codes, however, for production sites, distro packages are more stable, especially when Ubuntu offers up to 5-year software maintenance in comparison with 18-month from Openstack.

Other differences: Juju uses LXD containers for all Openstack services and has better integration with Ubuntu ecosystem (MAAS, LXD). Openstack Ansible use LXC containers for services on master (like nova-api, keystone, glance, cinder) or directly baremetal for services on workers (like nova-compute, cinder-volume).

Openstack Ansible offers wide possibilities of customization, e.g. method of installation (distro package vs source), based Linux distro (RedHat/CenOS, openSUSE and Ubuntu), selection of services to be installed. Beside Openstack installation, it also does many other tasks for security hardening and high availability (e.g. haproxy/Galera). As the result, it is more complex and need more time for deployment. (Un)fortunately, Openstack Ansible has very extensive documentation that is useful but may require time to study. See [1] and [2] from references for more information.

Installing a All-in-One Openstack site with nova-lxd via Openstack Ansible

This installation takes rather long time (totally around ~2h according to disk performance, network connection and list of installed services). Fast disk and network connection is strongly recommended because of intensive disk operation during installation and large amount of files downloaded from repositories

1. Install Ubuntu Bionic/ Create VM with vanilla Ubuntu Bionic (at least 16GB RAM, 80GB disk). So far do not init/configure the LXD service, Ansible script may report error like “storage already exist”.
2. Optional: Update all packages, then reboot
3. Cloning Openstack-Ansible

```
# git clone https://git.openstack.org/openstack/openstack-ansible \
/opt/openstack-ansible
# cd /opt/openstack-ansible
```

1. Optional: Choose version/branch if needed

```
# git tag -l
# git checkout master
# git describe --abbrev=0 --tags
```

1. Bootstrap Ansible (about 6min in the test)

```
# scripts/bootstrap-ansible.sh
```

1. Default scenario is without CEPH. Select CEPH scenario if you want and bootstrap AOI (about 6min in the test)

```
# export SCENARIO='ceph'
# scripts/bootstrap-aio.sh
```

1. Setting hypervisor to LXD:

Edit file “/etc/openstack_deploy/user_variables.yml”. Add a line “nova_virt_type: lxd” into it.

```
nova_virt_type: lxd
```

1. Examine list of services to be installed in “/etc/openstack_deploy/conf.d/” Copy more services from “/opt/openstack-ansible/etc/openstack_deploy/conf.d/” to “/etc/openstack_deploy/conf.d/” as needed. Remember to change the filename extension from aio to yml. For example, CEPH scenario is without Horizon and we want to add it.

```
# cp /opt/openstack-ansible/etc/openstack_deploy/conf.d/horizon.yml.aio \
/etc/openstack_deploy/conf.d/horizon.yml
```

1. Start the core installation (in 3 big steps). It takes long time (about 15min + 30min + 30min in the test), so you may want to use tmux or screen terminal if using SSH on unreliable network.

```
# cd /opt/openstack-ansible/playbooks
# openstack-ansible setup-hosts.yml
# openstack-ansible setup-infrastructure.yml
# openstack-ansible setup-openstack.yml
```

During execution of the first command “openstack-ansible setup-hosts.yml”, it is possible that you have timeout error during LXC caching if you have slow disk and/or network connection. In this case, please increase the value of “lxc_cache_prep_timeout” in “/etc/ansible/roles/lxc_hosts/defaults/main.yml” and re-execute the command.

The test “TASK [os_tempest : Execute tempest tests]” in the last command “openstack-ansible setup-openstack.yml” will fail. Ignore it and continue to next steps.

1. The Ansible script create a storage with the name “default” and driver “dir” for LXD, however, it does not work (btrfs or zfs required). Now create a new storage for LXD with btrfs with other name, e.g. “nova”. A simplest way is to run “lxd init” and configure storage via it. Do not touch networks/bridges or other configuration.

```
# lxd init
```

1. Setting LXD storage for nova: Edit file /etc/nova/nova.conf, add the following section into the file. For the pool option, use the name of storage created in the previous step, e.g. “nova”


```
[lxd]
allow_live_migration = True
pool = "nova"
```

Restart nova-compute service:

```
# systemctl restart nova-compute
# systemctl status nova-compute
```

1. Installation is now completed, however, some post-installation configurations are needed before starting the first VM. Refer to [3] for more information. The post-installation configuration can be done via CLI or via Horizon dashboard. The following steps show the configuration via Horizon.
2. Get the IP address of load balancer from “external_lb_vip_address” in “/etc/openstack_deploy/openstack_user_config.yml” file. Use iptables for IP forwarding to get the dashboard from your PC, e.g. :

```
# iptables -t nat -A PREROUTING -p tcp -m tcp --dport 8080 -j DNAT --to-
↪ destination external_lb_vip_address:443
```

Also remember do open firewall for the chosen port (8080 in the example).

1. Open Horizon dashboard in your browser as https://ip_address_of_host:8080. Log in using “admin” user and password stored in “keystone_auth_admin_password” item in “/etc/openstack_deploy/user_secrets.yml” file.
2. In Horizon, do the steps for configuring Openstack network: create a new private network, create a private subnet for the private network, create a router to connect private subnet to existing public network, open ports in security groups (e.g. port 22 for SSH). Also import SSH public key from “~/.ssh/id_rsa.pub” on the host.
3. Creating images does not work in default Horizon installed by Ansible, you must change Horizon setting or use command-line to create image. Use “lxc-attach” command to get into “aio1_utility_container_XXXXXX” container, load Openstack credential, download a LXD image from repository and add it to glance:

```
# lxc-ls
# lxc-attach aio1_utility_container_XXXXXXXX
# cd
# source openrc
# wget http://cloud-images.ubuntu.com/xenial/current/xenial-server-cloudimg-amd64-
↪ root.tar.gz
# glance image-create --name xenial-lxd --disk-format raw --container-format bare_
↪ --file xenial-server-cloudimg-amd64-root.tar.gz
```

1. Return to Horizon, create a new VM with the newly added “xenial-lxd” image. Remember to no create a new volume. Allocating a floating IP and assign it to the VM. From command line on the host, try to connect to the VM via ssh.
2. Success.

Notes:

- CEPH volume is still not attachable to VM by defaults, some additional work required.

References

1. <https://docs.openstack.org/openstack-ansible/latest/user/aio/quickstart.html>

2. <https://docs.openstack.org/project-deploy-guide/openstack-ansible/latest/>
3. <https://docs.openstack.org/openstack-ansible/latest/admin/openstack-firststrun.html>

Deploying OpenStack environment with nova-lxd via DevStack

This document is a step-by-step guide of OpenStack with nova-lxd deployment via DevStack on a single machine (All-in-One). The guide was tested on a host running Xenial or Bionic. The host has pre-installed following libraries:

- Python 2.7
- PyLXD 2.2.7

Installation steps

1. Create a new host running Xenial or Bionic with Python 2.7
2. Install the newest version of a library PyLXD 2.2.7

1. Install pip

```
$ sudo apt update
$ sudo apt install python-pip
```

1. Set up your environment variable `bash LC_ALL`

```
$ export LC_ALL="en_US.UTF-8"
$ export LC_CTYPE="en_US.UTF-8"
$ sudo dpkg-reconfigure locales
```

1. Download and install the library PyLXD

```
$ git clone https://github.com/lxc/pylxd.git
$ cd pylxd/
$ sudo pip install .
```

3. *Optional step:* Install ZFS

```
$ sudo apt install lxd zfsutils-linux
```

1. *Optional step:* If you need to install a different LXD/LXC version, execute the following steps:

1. Uninstall LXD and LXC (delete configuration and data files of LXD/LXC and it's dependencies)

```
$ sudo apt-get purge --auto-remove lxd lxc
```

1. Install the wanted version of LXD/LXC:

- LXD/LXC 3.0.1 on a host running Xenial

```
$ sudo apt-get purge --auto-remove lxd lxc
```

- the newest version

```
$ sudo snap install lxd
```

if you wish to install LXD 3.0 and then only get bugfixes and security updates. If running staging systems, you may want to run those on the candidate channels, using `bash--channel=candidate` and `bash--channel=3.0/candidate` respectively.

```
$ sudo snap install lxd --channel=3.0
```

2. Configure LXD:

1. In order to use LXD, the system user must be a member of the 'lxd' user group.

```
$ sudo adduser martin lxd
$ newgrp lxd
$ groups
```

2. LXD initialisation

```
$ sudo lxd init
```

The session below (LXD 3.0.1 with a zfs storage backend) is what was used to write this guide. Note that pressing Enter (a null answer) will accept the default answer (provided in square brackets).

```
Would you like to use LXD clustering? (yes/no) [default=no]:
Do you want to configure a new storage pool? (yes/no) [default=yes]:
Name of the new storage pool [default=default]: lxd
Name of the storage backend to use (btrfs, dir, lvm, zfs) [default=zfs]:
Create a new ZFS pool? (yes/no) [default=yes]:
Would you like to use an existing block device? (yes/no) [default=no]:
Size in GB of the new loop device (1GB minimum) [default=15GB]:
Would you like to connect to a MAAS server? (yes/no) [default=no]:
Would you like to create a new local network bridge? (yes/no) [default=yes]:
What should the new bridge be called? [default=lxdbr0]:
What IPv4 address should be used? (CIDR subnet notation, "auto" or "none")
↪[default=auto]:
What IPv6 address should be used? (CIDR subnet notation, "auto" or "none")
↪[default=auto]: none
Would you like LXD to be available over the network? (yes/no) [default=no]:
Would you like stale cached images to be updated automatically? (yes/no)
↪[default=yes]
Would you like a YAML "lxd init" preseed to be printed? (yes/no) [default=no]:
```

3. Optional step: Remove old LXD version to avoid conflict

```
$ sudo /snap/bin/lxd.migrate
```

1. *Optional step:* Increase the limit of number open files (only needed for larger tests). See <https://lxd.readthedocs.io/en/latest/production-setup/>
2. *Optional step:* Check configuration (as a user), if your configuration is correct

```
$ sudo /snap/bin/lxc storage list
$ sudo /snap/bin/lxc storage show default
$ sudo /snap/bin/lxc network show lxdbr0
$ sudo /snap/bin/lxc profile show default
```

9. Optional step: Run a test container in LXD (as a user), if LXD work correctly

```
$ sudo lxc launch ubuntu:16.04 u1
$ sudo lxc exec u1 ping www.ubuntu.com
```

(continues on next page)

(continued from previous page)

```
$ sudo lxc stop u1
$ sudo lxc delete u1
```

1. Create a user “Stack” and add it to the ‘lxd’ user group

```
$ sudo useradd -s /bin/bash -d /opt/stack -m stack
$ echo "stack ALL=(ALL) NOPASSWD: ALL" | sudo tee /etc/sudoers.d/stack
$ sudo usermod -G lxd -a stack
$ sudo su - stack
```

1. Download OpenStack installation scripts from DevStack repository

```
$ git clone https://git.openstack.org/openstack-dev/devstack
$ cd devstack
```

1. Create a **local.conf** file (a branch of a nova-lxd plugin is `bash stable/rocky`) as follows:

```
[[local|localrc]]

HOST_IP=127.0.0.1 # set this to your IP
FLAT_INTERFACE=ens2 # change this to your eth0

ADMIN_PASSWORD=devstack
DATABASE_PASSWORD=$ADMIN_PASSWORD
RABBIT_PASSWORD=$ADMIN_PASSWORD
SERVICE_PASSWORD=$ADMIN_PASSWORD
SERVICE_TOKEN=$ADMIN_PASSWORD

## run the services you want to use
ENABLED_SERVICES=rabbit,mysql,key
ENABLED_SERVICES+=,g-api,g-reg
ENABLED_SERVICES+=,n-cpu,n-api,n-crt,n-obj,n-cond,n-sch,n-novnc,n-cauth,placement-api,
→placement-client
ENABLED_SERVICES+=,neutron,q-svc,q-agt,q-dhcp,q-meta,q-l3
ENABLED_SERVICES+=,cinder,c-sch,c-api,c-vol
ENABLED_SERVICES+=,horizon

## disabled services
disable_service n-net

## enable nova-lxd
enable_plugin nova-lxd https://git.openstack.org/openstack/nova-lxd stable/rocky
```

1. Start installation of an OpenStack environment (it will take a 15 - 20 minutes, largely depending on the speed of your internet connection. Many git trees and packages will be installed during this process.)

```
$ ./stack.sh
```

1. Configuration of `bash nova-compute`: In order for a lxd storage pool to be recognized in nova, the `bash/etc/nova/nova-cpu.conf` file needs to have `bash [lxd]` section containing the following lines:

```
[lxd]
allow_live_migration = True
pool = {{ storage_pool }}
```

Restart nova-compute service:

```
$ systemctl restart devstack@n-cpu.service
$ systemctl status devstack@n-cpu.service
```

1. *Optional step:* if your OpenStack installation is set to incorrect repository, execute the following commands (adding a Rocky cloud archive):

```
$ sudo rm /etc/apt/sources.list.d/{{Bad_archive}}.list
$ sudo add-apt-repository cloud-archive:rocky
```

1. *Optional step:* Use IP forwarding to get access to the dashboard from outside by executing the following command on the host where the whole OpenStack environment with nova-lxd is installed:

```
$ sudo iptables -t nat -A PREROUTING -p tcp --dport 8080 -j DNAT --to "{{IP_of_
↪horizon}}:80"
```

where IP_of_horizon is the IP address of the dashboard that is given when the installation finishes (in the format of 10.x.x.x, e.g. 10.110.236.154)

1. Log into the dashboard (http://{{IP_address_of_host}}:8080/horizon), with “admin_domain” domain, “admin” user and “devstack” password, where ip_address_of_host is the IP of the host machine, where the whole OpenStack environment is installed (e.g. <http://147.213.76.100:8080/horizon>)
2. Create a new VM (Launch instance) in Compute->Instances. **Do not create a new volume (choose NO for new volume)**, and add only internal network.
3. In Network->Security group, add new ingress rules for ICMP (ping) and TCP port 22 (SSH) to default security group.
4. Allocate a new floating IP from Network -> Floating IPs and assign to the VM.
5. From host machine, try ssh to floating IP of VM

```
$ ssh ubuntu@{{Floating_ip}}
```

Handy commands:

Notes:

- CEPH volume is still not attachable to VM by defaults, some additional work required.

References

1. <https://discuss.linuxcontainers.org/t/lxd-3-0-0-has-been-released/1491>
2. <https://docs.jujucharms.com/devel/en/tut-lxd>
3. <https://docs.openstack.org/devstack/latest/>
4. <https://github.com/openstack/nova-lxd/blob/master/devstack/local.conf.sample>
5. <https://wiki.ubuntu.com/OpenStack/CloudArchive>

Installing nova-lxd with Juju

This is instruction of deploying Openstack with nova-lxd on single machine (All-in-One) for testing and deployment. Tested on Ubuntu Xenial and Bionic. Whole process of installation would take around 2h. Be careful with the LXD setting, leave default values when possible to avoid later problems. (“none” to IPv6, “lxdbr0” for bridge and “default” for name of LXD storage).

Installation

1. Create/Install a new machine with Ubuntu 16/18.04 and adequate performance (16GB RAM needed)
2. Optional: Update all packages (as root)

```
# apt update && apt dist-upgrade -y && apt autoremove -y
```

1. Install LXD (version 3.x required):

```
# sudo snap install lxd
```

1. Configure LXD:

```
# /snap/bin/lxd init
```

Say none to IPv6. If there a empty disk volume (block storage) attached, use it as existing block storage for LXD without mounting/formatting, otherwise use dir (with name of the storage as “default” for both cases). For other options, leave default value.

1. Optional: Add current user to LXD group (for using LXD as user later)

```
# sudo usermod -a -G lxd $USER  
# newgrp lxd
```

1. Optional: remove old LXD version to avoid conflict

```
# sudo /snap/bin/lxd.migrate
```

1. Optional: increase the limit of number open files (only needed for larger tests). See <https://lxd.readthedocs.io/en/latest/production-setup/>

```
# echo fs.inotify.max_queued_events=1048576 | sudo tee -a /etc/sysctl.conf  
# echo fs.inotify.max_user_instances=1048576 | sudo tee -a /etc/sysctl.conf  
# echo fs.inotify.max_user_watches=1048576 | sudo tee -a /etc/sysctl.conf  
# echo vm.max_map_count=262144 | sudo tee -a /etc/sysctl.conf  
# echo vm.swappiness=1 | sudo tee -a /etc/sysctl.conf  
# sudo sysctl -p
```

Also update default profile for improving network connection

```
# lxc profile device set default eth0 mtu 9000
```

1. Optional: Check configuration (as user), if your configuration is correct

```
# /snap/bin/lxc storage list  
# /snap/bin/lxc storage show default  
# /snap/bin/lxc network show lxdbr0  
# /snap/bin/lxc profile show default
```

1. Optional: Run a test container in LXD (as user), if LXD work correctly

```
# lxc launch ubuntu:16.04 u1
# lxc exec u1 ping www.ubuntu.com
# lxc stop u1
# lxc delete u1
```

1. Install juju:

```
# sudo snap install juju --classic
```

1. Install conjure-up

```
# sudo snap install conjure-up --classic
```

1. Optional: Start tmux terminal (to avoid unwanted termination in the case of network disruption)

```
# tmux
```

1. Start conjure-up (in tmux terminal if tmux is used):

```
# conjure-up
```

Choose Openstack with Nova-LXD, localhost. Other options can be left as default (lxdbr0 network bridge, “default” storage pool, ~/.ssh for SSH key), then deploy

1. Go to have a coffee for about 45-90 min (depending on performance of host machine). The installation with deploy nova-lxd with relevant services (keystone, glance, cinder, horizon). Remember the IP address of horizon dashboard from the output (e.g. <http://10.110.236.154/horizon>).
2. Use IP forwarding to get access to dashboard from outside by executing the following command on the host where whole openstack with nova-lxd is installed:

```
# sudo iptables -t nat -A PREROUTING -p tcp --dport 8080 -j DNAT --to "IP_of_
↪horizon:80"
```

where IP_of_horizon is the IP address of the dashboard that is given when the installation finishes (in format of 10.x.x.x, e.g. 10.110.236.154)

1. Log into the dashboard (http://ip_address_of_host:8080/horizon), with “admin_domain” domain, “admin” user and “openstack” password, where ip_address_of_host is the IP of host machine, where whole Openstack is installed (e.g. <http://147.213.76.100:8080/horizon>)
2. Create a new VM (Launch instance) in Compute->Instances. Do not create new volume (choose NO for new volume), and add only internal network.
3. In Network->Security group, add new ingress rules for ICMP (ping) and TCP port 22 (SSH) to default security group.
4. Allocate a new floating IP from Network -> Floating IPs and assign to the VM.
5. From host machine, try ssh to floating IP of VM

```
# ssh ubuntu@floating_ip
```

Notes

- Ceph and Cinder are installed together with other Openstack services, however attaching block storage does not work. According to <https://lists.gt.net/openstack/dev/64776>, it should require some additional work.

- Although Ubuntu Bionic with LXD 3.0 was used as base OS on the host, in LXD containers are Ubuntu Xenial with LXD 2.0

OpenStack nova-lxd testing configuration

As the nova-lxd plugin in OpenStack is still experimental, we need to deploy and test its current status to find out the working configuration before implementing new features

Testing of nova-lxd with different software configurations

OpenStack with nova-lxd has been deployed by different methods: OpenStack DevStack, JuJu, and OpenStack Ansible. Various combinations with OpenStack/LXD and base OS version have been tested to find out the working configuration of OpenStack nova-lxd for production.

Since nova-lxd plugin and OpenStack DevStack (deployment of OpenStack Cloud by Python scripts) have extremely shallow documentation, the first step was finding out the best starting configuration for development. We reached the following results (without advanced post-configuration efforts):

The hosts have the following static parameters:

There are two other deployment options for OpenStack alongside DevStack which are the following: Juju, and Ansible. The main advantage of these two approaches is automated deployment (with a post-configuration) of an OpenStack Cloud environment. The main differences between the deployment approaches are the configuration of an OpenStack environment, and LXD/LXC support. Juju supports LXD/LXC 2.0.11. It creates inside a Xenial/Bionic host another virtual layer by an LXD daemon from the host, and so Xenial with LXD/LXC 2.0.11 is installed in a container. All in all, it doesn't support higher versions of LXD/LXC which supports GPUs.

On the other hand, Ansible supports LXD/LXC 3.0.1. However, the situation about a nova-lxd plugin integration is the same, but the configuration of an OpenStack environment is different. According to Alex Kavanagh (one of the maintainers for the nova-lxd plugin) suggestions, the plugin has to be post-configured within a nova-compute configuration that is not documented in any official sources. The configuration file has to contain the following lines:

```
[DEFAULT]
compute_driver = nova_lxd.nova.virt.lxd.LXDDriver

[lxd]
allow_live_migration = True
pool = {{ storage_pool }}
```

A problem with the post-configuration is that it has to be performed in a different way within an OpenStack Cloud environment deployed by DevStack. The environment has a dedicated configuration file nova-cpu.conf. The other approaches create the standard deployment of an OpenStack Cloud environment, and so the configuration of the nova-lxd plugin is performed by editing of a nova.conf file.

Working configuration

According to the performed test mentioned above, we chose OpenStack Ansible repository for deployment of an OpenStack Cloud. The main reason is that it deploys a standard OpenStack Cloud with LXD/LXC 3.0.1 which supports GPUs.

3.1.4 uDocker

uDocker new GPU implementation

The use of NVIDIA GPUs for scientific computing requires the deployment of proprietary drivers and libraries. In order to use the GPUs inside containers, the devices associated to the GPU have to be visible inside the container. Furthermore, the driver has to be installed in the image and the version has to match the one deployed on the guest host. This turns the Docker images un-shareable and the image must be built locally for each host. The alternative is to have an image for each version of the driver, which is un-manageable since at each update, many images would have to be built. The *uDocker* released at the end of the Indigo-Datacloud project does not have such features, as such in order to use GPUs, the image has to have the NVIDIA drivers and libraries matching the host system. On the other hand, it is not necessary to pass the NVIDIA devices to the *uDocker* container since they are visible inside the container, in the same way a non-privileged user can use those devices in the host system.

The work performed during the first months of the DEEP-HybridDataCloud project, was to implement such automation. The development is available in the “devel” branch of the official GitHub repository [1], and scheduled for the first release of DEEP-HybridDataCloud software stack at the end of October 2018. The libraries and drivers deployed in the host system are made available to the containers. This version has been tested under several conditions and by several users and use cases. The tests performed in the framework of DEEP-HybridDataCloud project WP3 and WP4 are described in the following.


Test and evaluation of new implementation

The following tests and benchmark tools were developed to test performance of the python code packed in a Docker container and executed by means of *uDocker* [2]. We compare *uDocker* performance with baremetal execution and via Singularity 2.5.2-dist [3]. The benchmark tools are based on official Tensorflow Docker images from the Docker Hub [4] and deep learning python scripts publicly available at GitHub [5]. All scripts implement convolutional neural networks (CNN) but of different architecture: AlexNet, GoogLeNet, Overfeat, VGG [6], and one based on Tensorflow example for the MNIST dataset [7]. The latter processes MNIST data placed inside the Docker container while others process synthetic data generated on-the-fly. We adapted the scripts for more recent Tensorflow versions and homogenized the scripts to have a ‘burn-in’ phase, measure total runtime, mean time per batch and its standard deviation. In all tests the same version of the python scripts and corresponding Docker images, tagged as ‘181004’ at both GitHub and Docker Hub [7] are used. The tests comprise the following:

1. They are executed on GPU nodes of ForHLR II cluster [8], where each of the nodes contains four 12-core Intel Xeon processors E7-4830 v3 (Haswell), 1 TB of main memory, 4x960 GB local SSDs, 4 NVIDIA GeForce GTX980 Ti graphic cards. Operating System is RedHat Enterprise Linux 7.5, CUDA Toolkit 9.0.176 and NVIDIA Driver 384.81 are installed system-wide, cudnn 7.0.5 is installed in the user’s \$HOME directory. The test are performed with Python version 2.7.5.
2. For baremetal performance tests, two Tensorflow GPU versions 1.5.0 and 1.8.0 are installed in separate virtual environments via pip installation tool.
3. For *uDocker* tests, we build two Docker images based on the same Tensorflow versions, 1.5.0 and 1.8.0. The python scripts and MNIST data are stored inside the images. In both Docker images CUDA Toolkit is 9.0.176, cudnn version is 7.0.5, and Python is 2.7.12. We use *uDocker* (devel branch) to pull images from the Docker Hub. To run the containers, F3 (Fakechroot) mode is set with `-nvidia` flag (to add NVIDIA libraries and binaries).
4. For Singularity tests, Docker images built for *uDocker* tests were converted to Singularity images and uploaded to ForHLR II. Singularity version 2.5.2-dist was used. The containers are executed with `-nv` flag for NVIDIA support.
5. In all tests, the CNN scripts with synthetic data are executed for 1000 batches, therefore the mean time per batch is averaged over 1000 steps, the MNIST script is run for 20000 steps.

The results of the tests are shown in Figure 1, where we normalize the mean runtime per batch to the baremetal case and Tensorflow 1.8.0. Error bars are scaled by the mean time per batch for baremetal and Tensorflow 1.8.0. The tests do not indicate any penalty in executing the CNN scripts in either container technology *uDocker* or Singularity in comparison

with baremetal within the statistical uncertainty. They may even suggest that running the scripts inside containers is slightly faster than in baremetal, which could be connected to caching of data locally at the node in case of containers but needs to be better understood. Tensorflow 1.5.0 tends to be a bit faster than Tensorflow 1.8.0 for synthetic data but slower when real MNIST dataset is processed. This might be interpreted as improved I/O performance in Tensorflow 1.8.0 comparing to 1.5.0.



GPU

implementation tests Figure 1: *uDocker* performance tests using Tensorflow 1.5.0 and 1.8.0 in comparison with Baremetal installation in a user directory and Singularity. Lower values indicate better performance.

One job per node is executed in these tests, i.e. only one GPU card of the node is used and three other cards are not involved. *uDocker* however allows to pass environment settings inside containers, therefore making it possible to define which GPU card to use through providing the `CUDA_VISIBLE_DEVICES` parameter. It was possible to submit a job to our batch system such that in one job we can instantiate 4 *uDocker* containers where each container runs on a separate GPU card of the same node. This test does not indicate a degradation of performance in terms of total runtime and mean runtime per batch but shows significant increase in the uncertainty of the runtime per batch, especially in the case of the MNIST dataset.

As the tests suggest, it is of interest to study the behavior of processing large datasets from containers. Therefore we plan to extend our scripts to establish training on e.g. CIFAR datasets [9]. Such training may also require multi-GPU training so that we can preform tests to access all GPU cards on one node from single container. If our use-cases show interest, we may add other neural network architectures to the tests, such as Long Short-Term Memory (LSTM) or Generative Adversarial Networks (GAN).

9-Oct-2018

References

1. *uDocker* official GitHub repository: <https://github.com/indigo-dc/udocker>

2. Jorge Gomes, Emanuele Bagnaschi, Isabel Campos, MarioDavid, Luís Alves, João Martins, João Pina, Alvaro López-García, PabloOrviz, “Enabling rootless Linux Containers in multi-user environments: The udocker tool”, Computer Physics Communications, Volume 232, 2018, Pages 84-97, ISSN 0010-4655, <https://doi.org/10.1016/j.cpc.2018.05.021>.
3. G. M. Kurtzer, V. Sochat, and M. W. Bauer, “Singularity: Scientific containers for mobility of compute,” PLoS ONE, 2017.
4. Docker Hub tensorflow/tensorflow, <https://hub.docker.com/r/tensorflow/tensorflow/>
5. convnet-benchmarks: <https://github.com/soumith/convnet-benchmarks/tree/master/tensorflow> , Tensorflow MNIST example: https://github.com/tensorflow/tensorflow/blob/r1.2/tensorflow/examples/tutorials/mnist/mnist_deep.py
6.
 - AlexNet: Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks (2012), <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>
 - GoogLeNet: C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich: Going Deeper with Convolutions (2015), In Computer Vision and Pattern Recognition (CVPR) [<http://arxiv.org/abs/1409.4842>]
 - Overfeat: Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., & Lecun, Y. (2014). Overfeat: Integrated recognition, localization and detection using convolutional networks. In International Conference on Learning Representations (ICLR2014), CBLIS, April 2014 [<http://openreview.net/document/d332e77d-459a-4af8-b3ed-55ba>, <http://arxiv.org/abs/1312.6229>]
 - VGG: Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. CoRR abs/1409.1556 (2014), <http://arxiv.org/abs/1409.1556>
7. tf-benchmarks GitHub: <https://github.com/vykozlov/tf-benchmarks/tree/181004>, Docker Hub: <https://hub.docker.com/r/vykozlov/tf-benchmarks/tags/> , tags 181004-tf150-gpu and 181004-tf180-gpu
8. Computational resource ForHLR II available at Karlsruhe Institute of Technology, <https://wiki.scc.kit.edu/hpc/index.php/Category:ForHLR>
9. CIFAR-10 and CIFAR-100 datasets, <https://www.cs.toronto.edu/~kriz/cifar.html>

3.1.5 Miscellaneous

GPU sharing with MPS

From [1], MPS is a runtime service designed to let multiple MPI processes using CUDA to run concurrently on a single GPU. A CUDA program runs in MPS mode if the MPS control daemon is running on the system.

When CUDA is first initialized in a program, the CUDA driver attempts to connect to the MPS control daemon. If the connection attempt fails, the program continues to run as it normally would without MPS. If however, the connection attempt to the control daemon succeeds, the CUDA driver then requests the daemon to start an MPS server on its behalf.

If there’s an MPS server already running, and the user id of that server process matches that of the requesting client process, the control daemon simply notifies the client process of it, which then proceeds to connect to the server. If there’s no MPS server already running on the system, the control daemon launches an MPS server with the same user id (UID) as that of the requesting client process. If there’s an MPS server already running, but with a different user id than that of the client process, the control daemon requests the existing server to shutdown as soon as all its clients are done. Once the existing server has terminated, the control daemon launches a new server with the user id same as that of the queued client process.

The MPS server creates the shared GPU context, manages its clients, and issues work to the GPU on behalf of its clients. An MPS server can support upto 16 client CUDA contexts at a time. MPS is transparent to CUDA programs,

with all the complexity of communication between the client process, the server and the control daemon hidden within the driver binaries.

From [2], the Volta architecture introduced new MPS capabilities. Compared to MPS on pre-Volta GPUs, Volta MPS provides a few key improvements:

- Volta MPS clients submit work directly to the GPU without passing through the MPS server.
- Each Volta MPS client owns its own GPU address space instead of sharing GPU address space with all other MPS clients.
- Volta MPS supports limited execution resource provisioning for Quality of Service (QoS).

How to use MPS service

Start MPS service:

```
# nvidia-cuda-mps-control -d
```

Stop MPS service:

```
# echo quit | nvidia-cuda-mps-control
```

Testing environment

- HW: Virtual machine on IISAS-GPU cloud, flavor gpu1cpu6 (6 cores, 24GB RAM, 1 GPU Tesla K20m)
- SW: Ubuntu 16.04, latest nvidia driver and CUDA (nvidia driver version 410.48, CUDA 10.0.130)

Test 1. Test with CUDA native sample nbody, without nvidia-cuda-mps service

a. Single CUDA process

```
#!/nbody -benchmark -numbodies=512000  
  
number of bodies = 512000  
512000 bodies, total time for 10 iterations: 29438.994 ms  
= 89.047 billion interactions per second  
= 1780.930 single-precision GFLOP/s at 20 flops per interaction
```

b. Two processes at the same time:

```
# ./nbody -benchmark -numbodies=512000  
  
512000 bodies, total time for 10 iterations: 52418.652 ms  
= 50.010 billion interactions per second  
= 1000.194 single-precision GFLOP/s at 20 flops per interaction
```

Performance of each process reduced to about 1/2 due to parallel execution (overall GPU performance is the same). nvidia-smi shows both processes using GPU at the same time. No GPU conflicts detected.

Test 2. Test with CUDA native sample nboddy, with nvidia-cuda-mps service

a. Single CUDA process:

Same performance as without MPS server.

b. Two processes with different user IDs:

The second process is blocked (waiting without termination), it starts computation only the first process is terminated. Performance is the same as without MPS server.

In both case, nvidia-smi indicates nvidia-cuda-mps-server is using GPU, not the nboddy process.

c. Two processes with the same user ID:

Both processes will run in parallel. Performance will be evenly divided between processes, like without MPS service (Test 1b).

Test 2 has been repeated with nboddy commands placed inside Docker containers instead of baremetal, the same behavior. Note that Docker set user ID at root by default.

Test 3. Test with Docker using mariojmdavid/tensorflow-1.5.0-gpu image, without nvidia-cuda-mps service

Command used in the test:

```
# sudo docker run --runtime=nvidia --rm -ti mariojmdavid/tensorflow-1.5.0-gpu /home/
↪tf-benchmarks/run-bench.sh all
```

a. Single container:

all tests passed.

b. Two containers in parallel:

the second container shows different error messages according to the running benchmarks. Some error message is rather clear like that

```
2018-10-03 13:58:33.135064: W tensorflow/core/framework/op_kernel.cc:1198] Resource_
↪exhausted: OOM when allocating tensor with shape[1000] and type float on /
↪job:localhost/replica:0/task:0/device:GPU:0 by allocator GPU_0_bfc
```

Some error messages are rather internal

```
2018-10-03 13:51:00.160626: W tensorflow/stream_executor/stream.cc:1901] attempting_
↪to perform BLAS operation using StreamExecutor without BLAS support
Traceback (most recent call last):
  File "/home/tf-benchmarks/benchmark_alexnet.py", line 221, in <module>
    tf.app.run()
```

(continues on next page)

(continued from previous page)

```

File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/platform/app.py",
↳line 124, in run
    _sys.exit(main(argv))
File "/home/tf-benchmarks/benchmark_alexnet.py", line 217, in main
    run_benchmark()
File "/home/tf-benchmarks/benchmark_alexnet.py", line 210, in run_benchmark
    timing_entries.append(time_tensorflow_run(sess, grad, "Forward-backward"))
File "/home/tf-benchmarks/benchmark_alexnet.py", line 136, in time_tensorflow_run
    _ = session.run(target_op)
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/client/session.py",
↳line 895, in run
    run_metadata_ptr)
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/client/session.py",
↳line 1128, in _run
    feed_dict_tensor, options, run_metadata)
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/client/session.py",
↳line 1344, in _do_run
    options, run_metadata)
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/client/session.py",
↳line 1363, in _do_call
    raise type(e)(node_def, op, message)
tensorflow.python.framework.errors_impl.InternalError: Blas GEMM launch failed : a.
↳shape=(128, 9216), b.shape=(9216, 4096), m=128, n=4096, k=9216
    [[Node: affine1/affine1/MatMul = MatMul[T=DT_FLOAT, transpose_a=false,
↳transpose_b=false, _device="/job:localhost/replica:0/task:0/device:GPU:0"] (Reshape,
↳affine1/weights/read)]]
Caused by op u'affiliate/affine1/MatMul', defined at:
File "/home/tf-benchmarks/benchmark_alexnet.py", line 221, in <module>
    tf.app.run()
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/platform/app.py",
↳line 124, in run
    _sys.exit(main(argv))
File "/home/tf-benchmarks/benchmark_alexnet.py", line 217, in main
    run_benchmark()
File "/home/tf-benchmarks/benchmark_alexnet.py", line 181, in run_benchmark
    last_layer = inference(images)
File "/home/tf-benchmarks/benchmark_alexnet.py", line 120, in inference
    affn1 = _affine(resh1, 256 * 6 * 6, 4096)
File "/home/tf-benchmarks/benchmark_alexnet.py", line 76, in _affine
    affine1 = tf.nn.relu_layer(inpOp, kernel, biases, name=name)
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/ops/nn_impl.py",
↳line 272, in relu_layer
    xw_plus_b = nn_ops.bias_add(math_ops.matmul(x, weights), biases)
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/ops/math_ops.py",
↳line 2022, in matmul
    a, b, transpose_a=transpose_a, transpose_b=transpose_b, name=name)
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/ops/gen_math_ops.py",
↳line 2516, in _mat_mul
    name=name)
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/framework/op_def_
↳library.py", line 787, in _apply_op_helper
    op_def=op_def)
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/framework/ops.py",
↳line 3160, in create_op
    op_def=op_def)
File "/usr/local/lib/python2.7/dist-packages/tensorflow/python/framework/ops.py",
↳line 1625, in __init__

```

(continues on next page)

(continued from previous page)

```

self._traceback = self._graph._extract_stack() # pylint: disable=protected-access
InternalError (see above for traceback): Blas GEMM launch failed : a.shape=(128,
↪9216), b.shape=(9216, 4096), m=128, n=4096, k=9216
[[Node: affine1/affine1/MatMul = MatMul[T=DT_FLOAT, transpose_a=false,
↪transpose_b=false, _device="/job:localhost/replica:0/task:0/device:GPU:0"] (Reshape,
↪affine1/weights/read)]]

```

Test 4. Test with Docker using mariojmdavid/tensorflow-1.5.0-gpu image, with nvidia-cuda-mps service

Option “-ipc=host” required for connecting MPS service (full command “sudo docker run --runtime=nvidia --rm -ipc=host -ti mariojmdavid/tensorflow-1.5.0-gpu /home/tf-benchmarks/run-bench.sh all”), see <https://github.com/NVIDIA/nvidia-docker/issues/419>

Some tests passed but not all

```

/home/tf-benchmarks/run-bench.sh: line 78: 137 Aborted (core
↪dumped) python ${TFTest[$i]}

```

According to [3], only Tensorflow with version 1.6 and higher can support MPS.

Test 5. Test with Docker using vykozlov/tf-benchmarks:181004-tf180-gpu image, without and with nvidia-cuda-mps service

Tensorflow 1.8.0, GPU version, python 2, command:

```

sudo docker run --ipc=host --runtime=nvidia --rm -ti vykozlov/tf-benchmarks:181004-
↪tf180-gpu ./tf-benchmarks.sh all

```

a. without MPS service

all tests passed.

b. with MPS service

Only first part (forward) of each test passed, then the execution terminated (core dumped).

Identified reasons why Tensorflow does not work correctly with MPS

The reasons have been discussed in [3]:

- stream callbacks are not supported on pre-Volta MPS clients. Calling any stream callback APIs will return an error. (from MPS official document [4])
- But CUDA streams are used everywhere in Tensorflow

So Tensorflow will not work with MPS on old (pre-Volta) GPU.

Final remarks:

- Without MPS service, native CUDA samples can be executed in parallel and the GPU performance is divided among processes
- With MPS service, CUDA executions with different user IDs are serialized, one needs to wait until other finishes.
- CUDA processes with the same user ID can be executed in parallel.
- Tensorflow will not work with MPS on old (pre-Volta) GPU.
- Need to test on newer GPU cards (Volta)

References

1. <http://manpages.ubuntu.com/manpages/xenial/man1/alt-nvidia-340-cuda-mps-control.1.html>
2. <https://docs.nvidia.com/deploy/mps/index.html>
3. <https://github.com/tensorflow/tensorflow/issues/9080>
4. <https://docs.nvidia.com/deploy/mps/index.html>

Useful project links

- [DEEP IAM](#)
 - [DEEP Open Catalog - Marketplace](#)
 - [DEEP Dashboard \(Training, Advanced\)](#)
 - [DEEP Nextcloud](#)
 - [Official GitHub](#)
 - [Official DockerHub](#)
-

CHAPTER 4

Indices and tables

- `genindex`
- `search`