

Computing solutions at EP-NU



How computing infrastructure evolved at CERN?



Cloud computing

Cloud computing is the delivery of on-demand computer system resources, as data storage or computing power, without direct active management by the user.




Computing possibilities at CERN

Focused on EP-NU resources for high intensive computing.



Computing resources



1. CERN OpenStack

EP-NU Stack:

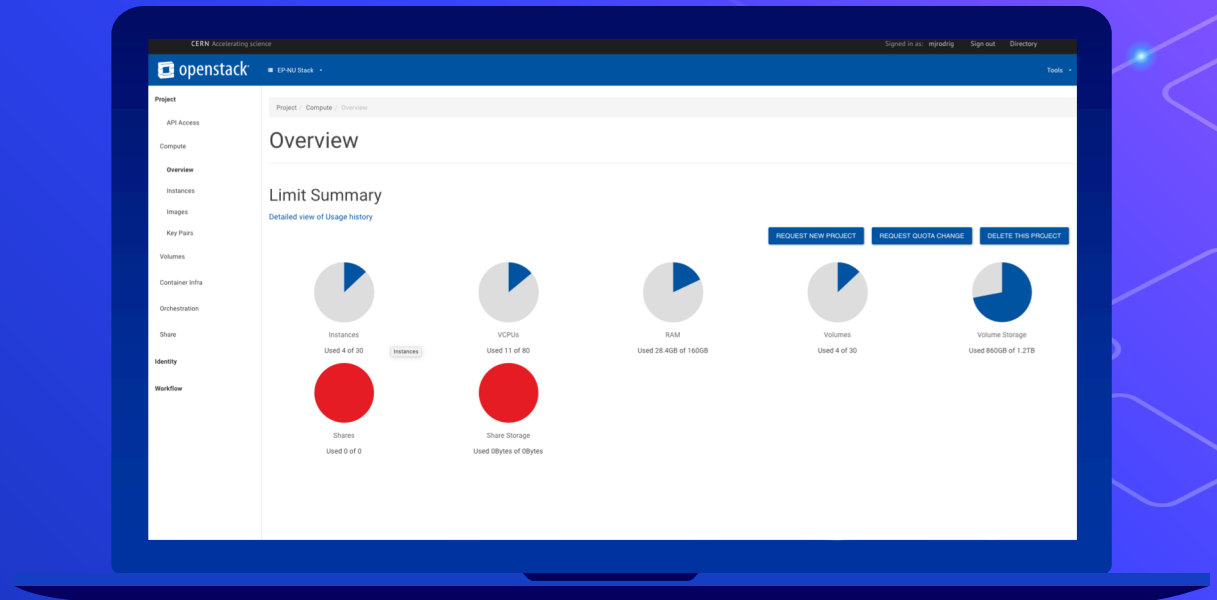
<https://openstack.cern.ch/project/>



CERN OpenStack



OpenStack allow us to create Virtual Machines (VM) on-demand using one of the CERN operative system images or your own image.



Tutorial:
<https://clouddocs.web.cern.ch/index.html>

CERN OpenStack



OpenStack allow us to create Virtual Machines (VM) on-demand using one of the CERN operative system images or your own image.

Tutorial:
<https://clouddocs.web.cern.ch/index.html>

The screenshot shows the OpenStack dashboard interface for the 'EP-NU Stack' project. The main content area is titled 'Instances' and displays a table of running virtual machines. A blue box highlights the 'LAUNCH INSTANCE' button in the top right corner of the table area.

<input type="checkbox"/>	Instance Name	Image Name	IP Address	Flavor	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
<input type="checkbox"/>	epnu-test-01	epnu-default-centos7	188.184.99.153 2001:1458:d00:3a::100:438	m2.small	mjrodrig-Mac	Active	cern-geneva-b	None	Running	6 days, 14 hours	CREATE SNAPSHOT -
<input type="checkbox"/>	epnu-radi	-	188.184.97.202 2001:1458:d00:39::100:4a3	m2.medium	mjrodrig-Mac	Active	cern-geneva-a	None	Running	6 days, 20 hours	CREATE SNAPSHOT -
<input type="checkbox"/>	epnu-gpu01	-	188.184.195.175 2001:1458:301:72::100:3ad	g112.xlarge	mjrodrig-Mac	Active	none	None	Running	1 week, 2 days	CREATE SNAPSHOT -
<input type="checkbox"/>	EPNU-OS-Wi nVM	-	188.185.112.246 2001:1458:d00:a::100:f0	m2.large	-	Active	cern-geneva-a	None	Running	2 months, 3 weeks	CREATE SNAPSHOT -

CERN OpenStack



OpenStack allow us to create Virtual Machines (VM) on-demand using one of the CERN operative system images or your own image.

Tutorial:
<https://clouddocs.web.cern.ch/index.html>

Launch Instance

Please provide the initial hostname for the instance, the availability zone where it will be deployed, and the instance count. Increase the Count to create multiple instances with the same settings.

Instance Name
my-test-vm

Description
This is a test VM on OpenStack

Availability Zone
Any Availability Zone

Count
1

Total Instances (30 Max)
17%
4 Current Usage
1 Added
25 Remaining

← BACK NEXT → LAUNCH INSTANCE

× CANCEL

CERN OpenStack

OpenStack allow us to create Virtual Machines (VM) on-demand using one of the CERN operative system images or your own image.

Tutorial:
<https://clouddocs.web.cern.ch/index.html>

EP-NU custom image with some common tools included

Launch Instance

Instance source is the template used to create an instance. You can use an image, a snapshot of an instance (image snapshot), a volume or a volume snapshot (if enabled). You can also choose to use persistent storage by creating a new volume.

Select Boot Source

Image

Create New Volume

Allocated

Name	Updated	Size	Type	Visibility
Select an item from Available items below				

Available 15 Select one

Click here for filters.

Name	Updated	Size	Type	Visibility
> CC7 TEST - x86_64 [2020-10-14]	10/14/20 7:39 PM	4.00 GB	raw	Public
> C8 TEST - x86_64 [2020-10-14]	10/14/20 7:39 PM	4.00 GB	raw	Public
epnu-default-centos7	10/12/20 6:31 PM	0 bytes	qcow2	Private
> epnu-gpu default	10/12/20 5:59 PM	0 bytes	qcow2	Private
> CC7 - x86_64 [2020-10-01]	10/1/20 12:11 PM	4.00 GB	raw	Public
> C8 - x86_64 [2020-10-01]	10/1/20 12:12 PM	4.00 GB	raw	Public
> CernVM 4 - Bootloader v2020.04-1 [2020-04-09]	4/14/20 2:07 PM	24.06 MB	raw	Public
> CernVM 3 - Bootloader v2020.04-1 [2020-04-09]	4/14/20 2:08 PM	24.06 MB	raw	Public

CERN OpenStack



OpenStack allow us to create Virtual Machines (VM) on-demand using one of the CERN operative system images or your own image.

Tutorial:
<https://clouddocs.web.cern.ch/index.html>

Launch Instance

Flavors manage the sizing for the compute, memory and storage capacity of the instance.

Allocated

Name	VCPUS	RAM	Total Disk	Root Disk	Ephemeral Disk	Public
Select an item from Available items below						

Available 4 Select one

Click here for filters.

Name	VCPUS	RAM	Total Disk	Root Disk	Ephemeral Disk	Public	
> m2.large	4	7.32 GB	40 GB	40 GB	0 GB	Yes	↑
> m2.small	1	1.83 GB	10 GB	10 GB	0 GB	Yes	↑
> m2.medium	2	3.66 GB	20 GB	20 GB	0 GB	Yes	↑
> g112.xlarge	4	15.63 GB	80 GB	80 GB	0 GB	No	↑

× CANCEL ← BACK NEXT → ▶ LAUNCH INSTANCE

CERN OpenStack



OpenStack allow us to create Virtual Machines (VM) on-demand using one of the CERN operative system images or your own image.

Tutorial:
<https://clouddocs.web.cern.ch/index.html>

Launch Instance

Details

Source

Flavor

Key Pair

Configuration

Metadata

A key pair allows you to SSH into your newly created instance. You may select an existing key pair, import a key pair, or generate a new key pair. ?

+ CREATE KEY PAIR + IMPORT KEY PAIR

Allocated

Displaying 1 item

Name	Fingerprint
> mjrodrig-Mac	64:92:cb:75:4d:f6:56:88:95:e4:34:ba

Displaying 1 item

▼ Available 5 Select one

🔍 Click here for filters. ✕

Displaying 5 items

CERN OpenStack



OpenStack allow us to create Virtual Machines (VM) on-demand using one of the CERN operative system images or your own image.

Tutorial:
<https://clouddocs.web.cern.ch/index.html>

Launch Instance

This step allows you to add Metadata items to your instance. ?

You can specify resource metadata by moving items from the left column to the right column. In the left column there are metadata definitions from the Glance Metadata Catalog. Use the "Custom" option to add metadata with the key of your choice.

Available Metadata	Filter	Existing Metadata	Filter
▼ LanDB		landb-mainuser mjrodrig -	
LanDB Alias		landb-responsible mjrodrig -	
LanDB Description			
LanDB Internet Connectivity			
LanDB IPv6 Ready			
LanDB OS			
LanDB OS Version			

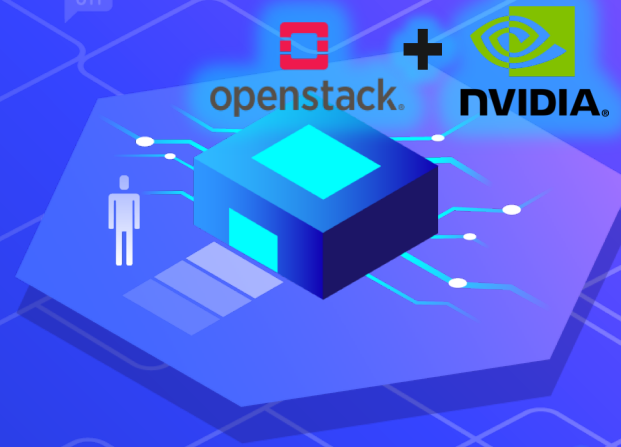
LanDB Main User (*landb-mainuser*)
Primary user of the virtual machine.

× CANCEL ← BACK NEXT → **LAUNCH INSTANCE**

2. GPU services on OpenStack

EP-NU Stack:

<https://openstack.cern.ch/project/>



CERN OpenStack



We have been granted with **ONE** GPU instance on OpenStack

<input type="checkbox"/> epnu-gpu01	-	188.184.195.175	g112.xlarge
		2001:1458:301:72::100:3ad	

No graphical environment.

Access is granted by e-group:

epnu-comp

Self-subscription policy

Access through ssh

Use virtual environments for python developments

CERN OpenStack



We have been granted with **ONE** GPU instance on OpenStack

<input type="checkbox"/>	epnu-gpu01	-	188.184.195.175	g112.xlarge
			2001:1458:301:72::100:3ad	

- Reserved to the most *exotic* computing tasks
- Shared resource. Pay attention to others' work!

```
[mjrodrig@epnu-gpu01 ~]$ nvidia-smi
Mon Oct 19 11:01:52 2020

+-----+
| NVIDIA-SMI 455.23.05   Driver Version: 455.23.05   CUDA Version: 11.1   |
+-----+-----+
| GPU  Name            Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|                                           |                  |     MIG M.     |
+-----+-----+
|  0  Tesla T4             Off          | 00000000:00:05:0 Off |   0          0      |
| N/A   64C    P0         31W / 70W |  0MiB / 15109MiB |      4%      Default |
+-----+-----+

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


3. Kubeflow

CERN Machine Learning platform:
<https://ml.cern.ch>*

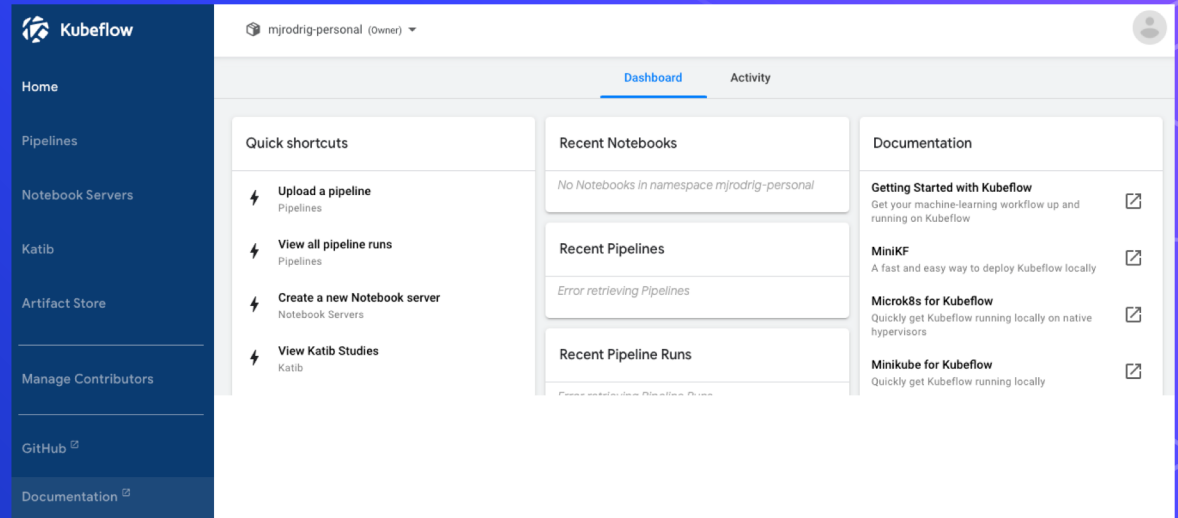


*The service is still in an experimental stage.

Kubeflow



- Kubeflow is a machine learning framework that easy the common machine learning workflow in a kubernetes environment.



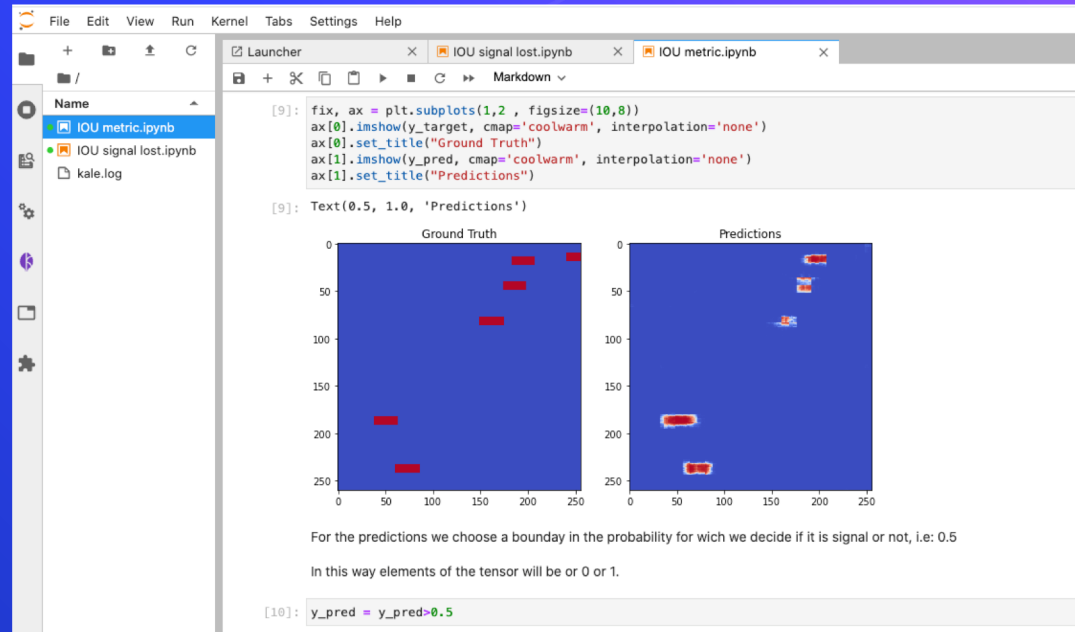
More details on Dejan's talk: <https://indico.cern.ch/event/961276>

Kubeflow



Advance features:

- Various Frameworks
 - Tensorflow, PyTorch, MPI
- Jupyter Notebooks
- Machine Learning Pipelines
- Katib - Hyper-parameter Optimization
- KALE - Notebooks to Pipelines or Katib
- Fairing - High level API



4. EP-NU techlab server



EP-NU Techlab server

The server was hosted at CERN Computing center. The server is working however it reached its end of life (in terms of warranty).

We will place it in a rack borrowed from EP-DT-DI



Thanks!

Any questions? 🤖

