



VLAAMS
SUPERCOMPUTER
CENTRUM



Vlaanderen
is supercomputing

vscentrum.be



Possibilities of the use of VSC in Machine Learning Models

Kenneth Hoste & Lara Peeters (HPC-UGent)

Mon 26 Feb 2024

Transitioning ML Workflows from Academia to Practice

Flanders AI Academy (VAIA)

Who are we?



Kenneth Hoste: Computer scientist from Belgium



- Masters (2005) + PhD (2010) in Computer Science @ Ghent University
- Was using Machine Learning techniques in 2005-2010 during PhD research
- HPC system administrator + user support at Ghent University since Oct 2010
- Lead developer of **EasyBuild** - tool to install **scientific software on HPC clusters**
- Active contributor to **EESSI**, partner in MultiXscale EuroHPC CoE
- Fan of open source software, beers, and stickers


Lara Peeters: Digital Art Historian from Belgium



- Masters in Art History (2021) and Digital Humanities (2022) @ KU Leuven
- Hired on the MultiXscale project at Ghent University (Belgium) since May 2023
- Active contributor to **EasyBuild** & **EESSI**, partner in MultiXscale EuroHPC CoE
- Studying Artificial Intelligence in the field of art

What is High-Performance Computing (HPC)?

harnessing the power of multiple interconnected cores/nodes/processing units



Colossus - first digital computer (1944)



Cray-2, first supercomputer (1985)



"Rack" with PCs (~2010?)



HPC-UGent Tier-2 (2018)



IBM Blue Gene/P, ANL, US (2007)



MareNostrum 4 in Barcelona (2017)



Hortense VSC Tier-1 compute (since 2022)



Scientific Computing & Supercomputers

- HPC = Supercomputers \approx scientific computing
- Supercomputers are big, fast, expensive
- Used for scientific simulations + much more
- Multi-tenant: 1000s of scientists, different profiles



⇒ Wide variety of scientific software required...

- **Performance** is key: faster software means more science gets done (for roughly same cost)



Fastest supercomputers today can perform over 10^{18} 64-bit floating-point operations per second (1 exaFLOPS)




Cray-1 (1975)



Frontier @ Oak Ridge, US
(#1 in [Top500](#) of June '2023)


What are supercomputers used for?




Weather/climate prediction & modelling



Bioinformatics (DNA analysis, ...)




Molecular modelling, materials research, ...



*Computational Fluid Dynamics (CFD)
Aerodynamics, studying flow of gasses & liquids in volumes, etc.*



Simulation of atomic weapons, etc.



Training of machine learning & AI models

Training Artificial Intelligence models

Training good AI models requires *a lot* of compute!


GPT-3 175B:

- Trained on V100s (~125 TeraFLOPs)
- Requires ~ $3e23$ FLOP (Floating Point Operations), would take one A100 GPU ~30 years to calculate
- Took 34 days to train

GPT-4 (<https://medium.com/@daniellefranca96/gpt4-all-details-leaked-48fa20f9a4a>)

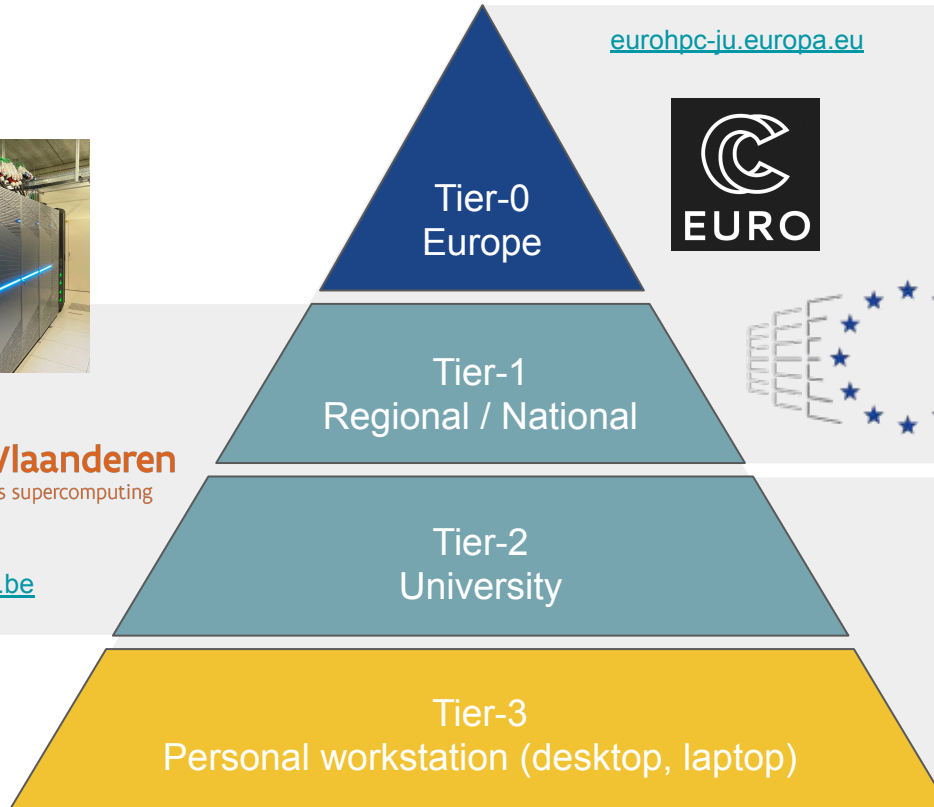

- Used 25,000 NVIDIA A100s (each ~312 TeraFLOPS)
- Took 90-100 days to train

To some extent, AI workloads are “just” another workload that can be run on supercomputers...



Source: Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, et al. *Language Models are few-shot Learner: OpenAI*, 2020.

Different “tiers” of supercomputers in Europe



eurohpc-ju.europa.eu



EuroHPC
Joint Undertaking

VLAAMS
SUPERCOMPUTER
CENTRUM




Vlaanderen
is supercomputing

vscentrum.be



**GHENT
UNIVERSITY**

ugent.be/hpc




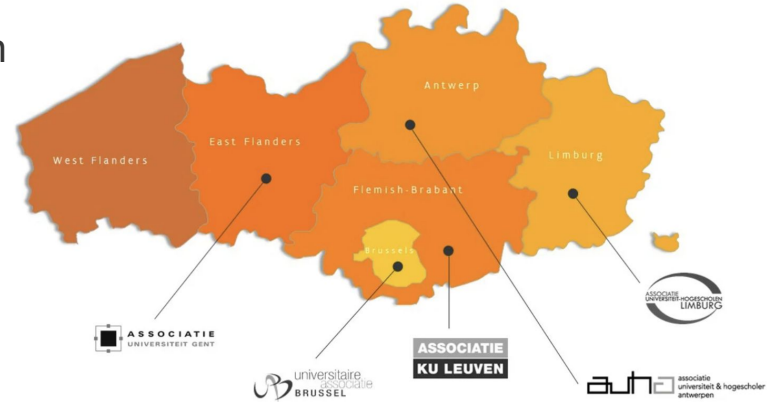
Flemish Supercomputer Centre (VSC)

vscentrum.be

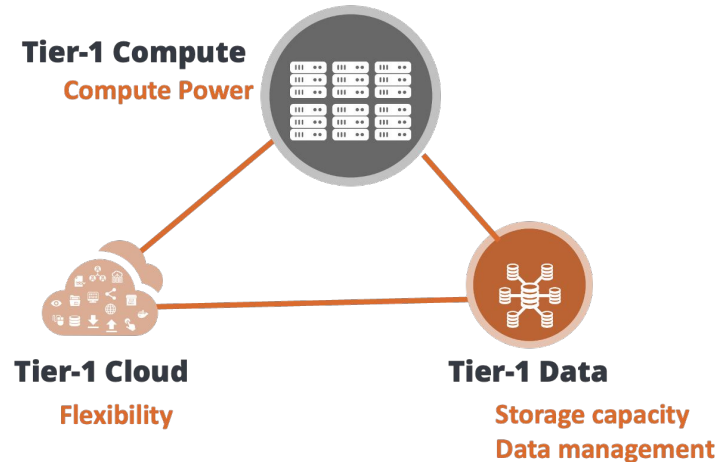
VLAAMS
SUPERCOMPUTER
CENTRUM



Virtual centre, collaboration between
Flemish Universities (since 2008),
funded via FWO + universities



*Tier-2 compute infrastructure,
operated by the VSC hubs*



VSC Tier-2 infrastructure

- You can use your VSC account to access HPC infrastructure provided by other VSC hubs
- Your `$VSC_HOME` and `$VSC_DATA` directories are available on each of these systems



Example: HPC-UGent Tier-2



- HPC-UGent Tier-2 infrastructure currently consists of **6 clusters** (+ login nodes, shared storage, ...)
- Different types of clusters:
 - 3 CPU-only compute clusters (+ 1 pilot cluster)
 - 2 GPU clusters + 1 CPU-only interactive + debug cluster
- Shared storage: home + apps (~90 TB), data (~3.900 TB), scratch (~1.700 TB)
- **Available for academic researchers free of charge**, funding through [FWO](#)
- All running **Red Hat Enterprise Linux 8 (RHEL8)** as operating system



Example: HPC-UGent Tier-2



- HPC-UGent Tier-2 infrastructure currently consists of **6 clusters**

(+ local)

Similar infrastructure available at other VSC hubs
(in Antwerp, Brussels, Leuven)

- Different

-
-

Small differences in user experience (software, credits, etc.)

- Shared

**Same VSC account to access all VSC infrastructure,
same home directory, \$VSC_DATA available everywhere**

- Available

docs.vscenrum.be

- All running



GPUs in VSC Tier-2 infrastructure



docs.vscenrum.be/hardware-tier2.html

UGent:

- Joltik cluster: 10 nodes with 4x NVIDIA V100 (32GB)
- Accelgor cluster: 9 nodes with 4x NVIDIA A100 (80GB)
- Debug cluster: 16 nodes with 1 shared NVIDIA A2

VUB:

- 4 nodes with 2x P100 (16GB)
+ 10 nodes with 2x A100 (40GB)

UA:

- Leibniz cluster: 2 nodes with 2x NVIDIA P100 (16GB)
- Vaughan cluster: 1 node with 4x NVIDIA A100 (40GB) + 2 nodes with 4x AMD MI100 (32GB)

KUL:

- Genius cluster: 20 nodes with 4x NVIDIA P100 (16GB) + 2 nodes with 8x NVIDIA V100 (32GB)
- wICE cluster: 4 nodes w/ 4x A100 (80GB) + 5 interactive nodes w/ 1x A100 (80GB)
(+ coming soon: 4 nodes w/ 4x NVIDIA H100 (80GB))



GPUs in VSC Tier-2 infrastructure



docs.vscentrum.be/hardware-tier2.html

UGent:

- Joltik cluster: 10 nodes with 4x NVIDIA V100 (32GB)
- Accelgor cluster: 9 nodes with 4x NVIDIA A100 (80GB)
- Debug cluster: 16 nodes with 1 shared NVIDIA A2

VUB:

- 4 nodes with 2x P100 (16GB)
+ 10 nodes with 2x A100 (40GB)

UA:

- Leibniz cluster: 2 nodes with 2x NVIDIA P100 (16GB)
- Vaughan cluster: 1 node with 4x NVIDIA A100 (40GB) + 2 nodes with 4x AMD MI100 (32GB)

KUL:

- Genius cluster: 20 nodes with 4x NVIDIA P100 (16GB) + 2 nodes with 8x NVIDIA V100 (32GB)
- wICE cluster: 4 nodes w/ 4x A100 (80GB) + 5 interactive nodes w/ 1x A100 (80GB)
(+ coming soon: 4 nodes w/ 4x NVIDIA H100 (80GB))

In total:

- 92x NVIDIA P100 (16GB)
- 56x NVIDIA V100 (32GB)
- 24x NVIDIA A100 (40GB)
- 57x NVIDIA A100 (80GB)
- 16x NVIDIA A2 (16GB) (*shared)
- 8x AMD MI100 (32GB)
- 16x NVIDIA H100 (80GB) [soon]

VSC Tier-1 compute cluster “Hortense” (a.k.a. dodrio)



- Hosted, operated, and supported by HPC-UGent team since 2021
- 2x 384 CPU-only nodes (2x 64-core AMD Rome/Milan CPUs)
+ 40 A100 GPU nodes (20x 40GB + 20x 80GB)
- **Over 100,000 CPU cores + 80 NVIDIA A100 GPUs in total!**
- High-speed Infiniband network (HDR-100) + 6PB of dedicated scratch storage

vscentrum.be/compute

docs.vscentrum.be

compute@vscentrum.be



- **Project-based access**
 - Free of charge for academics, funded by FWO
 - Pay-what-you-use for industry
- 3 cut-off dates per year for submitting project proposals
- Project duration is typically 8 months
- 500k - 5M core hours (CPU-only) or 1k - 25k GPU hours (or more)

Getting a VSC account

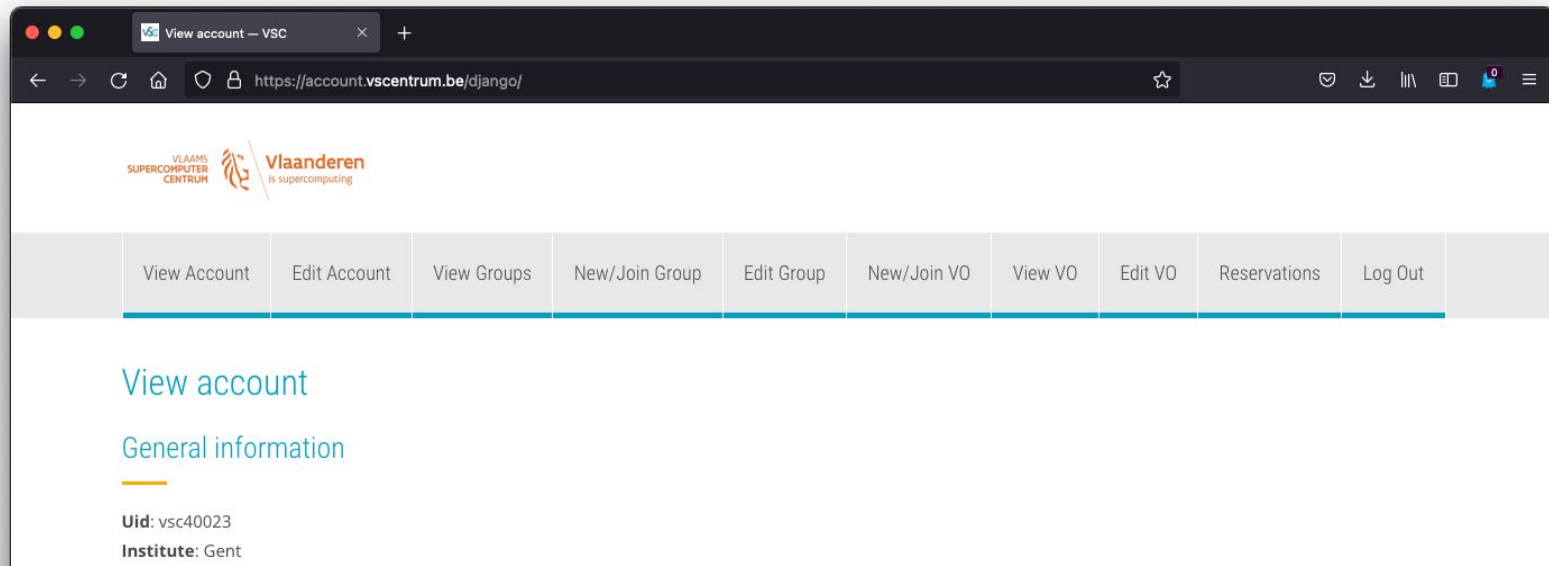
- All associated with one of the Flemish universities can request a VSC account
 - Researchers & staff
 - Master/Bachelor students
- **VSC account can be used to access HPC infrastructure on all VSC sites**
- Subscribed to `hpc-announce` and `hpc-users` mailing lists
- Beware of using HPC for teaching/exam purposes!
 - No guarantee on HPC availability (due unexpected power outage, maintenance, ...)
 - Have a backup plan at hand
 - Advisable teaching/exam formula: project work
- See also [VSCentrum documentation](#)

Managing your VSC account

You can manage your VSC account via the VSC account page

account.vscentrum.be

Can be used to join/leave user groups, consult storage usage, request more storage quota, ...
manage your Virtual Organisation (VO), ...



The screenshot shows a web browser window with the URL <https://account.vscentrum.be/django/>. The page header includes the VLAAMS SUPERCOMPUTER CENTRUM logo and the text "Vlaanderen is supercomputing". A navigation menu contains the following items: View Account, Edit Account, View Groups, New/Join Group, Edit Group, New/Join VO, View VO, Edit VO, Reservations, and Log Out. The main content area is titled "View account" and "General information". Below this, the user's details are displayed: **Uid:** vsc40023 and **Institute:** Gent.

Connecting via login nodes: SSH or web portal

```
$ ssh vsc40023@login.hpc.ugent.be

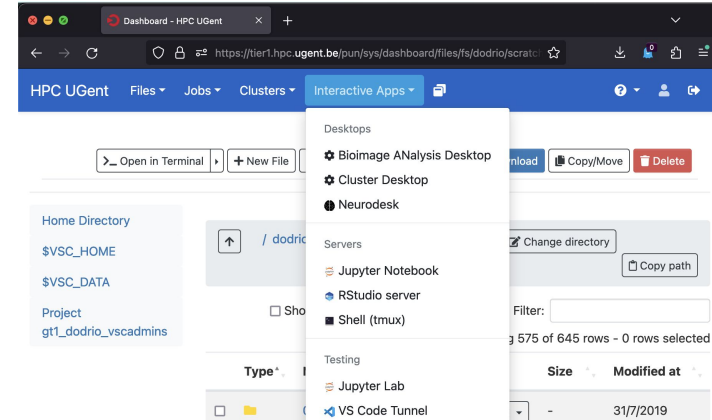
STEVIN HPC-UGent infrastructure status on Thu, 17 Nov 2022 20:30:01
cluster - full - free - part - total - running - queued
         nodes nodes free nodes jobs jobs
-----
slaking  0  6  2  10  3  4
swalot   19  0  94  118  1659  1172
skitty   33  0  33  72  1391  439
victini  7  0  87  96  588  30080
joltik   3  1  6  10  26  833
kirlia   7  0  9  16  23  40
doduo    50  0  67  128  1983  12695
accelgor  6  0  2  9  31  1267

Documentation is available at:
https://www.ugent.be/hpc/en/support/documentation.htm
For a full view of the current loads and queues see:
https://hpc.ugent.be/clusterstate/
Updates on current system status and maintenance can be found on:
https://www.ugent.be/hpc/en/infrastructure/status
To contact the support team, send an email to hpc@ugent.be

Last login: Thu Nov 17 20:30:34 2022 from 10.141.10.60
[vsc40023@gligar07 ~]$
[vsc40023@gligar07 ~]$ hostname
gligar07.gastly.os
```

Option 1: login via SSH

- Log in via terminal application
- Easy & powerful (if you can use Linux shell)
- Requires SSH public/private key pair.



Option 2: login via web portal




- Log in via standard internet browser (Firefox, Chrome, ...)
- Authentication via VSC account page & institute login.
- Includes file browser, desktop environment, GUI apps, Jupyter notebooks, in-browser terminal, etc.
- Currently only available for UGent/KUL Tier-2 + Tier-1

Rich central software stack on compute clusters



- All VSC compute cluster (Tier-2 + Tier-1) provide a central software stack
- *Optimized* installations of scientific software & required dependencies
- Available via intuitive “environment modules” interface (module command)
- Also supported: containers (via Apptainer runtime, supports Docker containers)

OpenFOAM®



TensorFlow



Questions?



Demo: Submitting a PyTorch job (CPU-only)

pytorch.org/tutorials/beginner/pytorch_with_examples.html#pytorch-nn









 PyTorch

Table of Contents

Learning PyTorch with Examples

 / user / gent / 400 / vsc40023 / VAIA / PyTorch-demo / 

Show On

	Type	Name	
<input type="checkbox"/>		job-gpu.sh	
<input type="checkbox"/>		job.sh	
<input type="checkbox"/>		pytorch-nn-demo.py	
<input type="checkbox"/>		README	

```
#!/bin/bash
#PBS -l nodes=1:ppn=8
#PBS -N pytorch-nn-demo

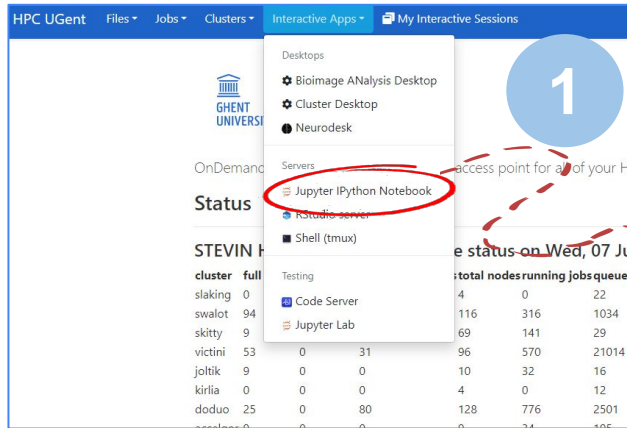
module load PyTorch/2.1.2-foss-2023a

cd $PBS_O_WORKDIR

python pytorch-nn-demo.py
```

```
[vsc40000@gligar08]$ qsub job.sh
61687984
```

Demo: Jupyter notebook via web portal



1

HPC UGent Files Jobs Clusters Interactive Apps My Interactive Sessions

Desktops

- Bioimage ANalysis Desktop
- Cluster Desktop
- Neurodesk

Servers

- Jupyter IPython Notebook**
- RStudio server
- Shell (tmux)

Testing

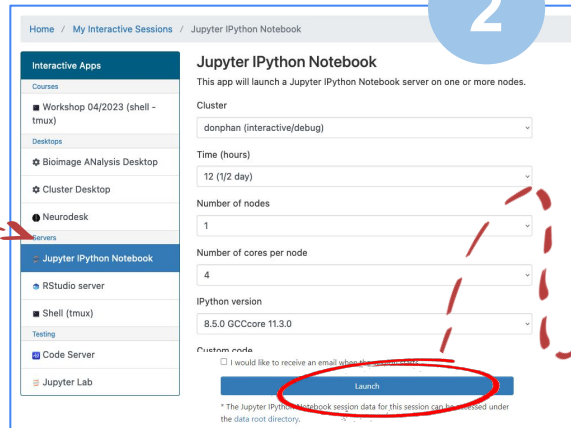
- Code Server
- Jupyter Lab

OnDemand

Status

STEVIN

cluster	full	total nodes	running	jobs queue
slaking	0	4	0	22
swalot	94	116	316	1034
skitty	9	69	141	29
victini	53	0	31	96
joltik	9	0	0	10
kirfia	0	0	0	4
doduo	25	0	80	128



2

Home / My Interactive Sessions / Jupyter IPython Notebook

Interactive Apps

Courses

- Workshop 04/2023 (shell - tmux)

Desktops

- Bioimage ANalysis Desktop
- Cluster Desktop
- Neurodesk

Servers

- Jupyter IPython Notebook**
- RStudio server
- Shell (tmux)

Testing

- Code Server
- Jupyter Lab

Jupyter IPython Notebook

This app will launch a Jupyter IPython Notebook server on one or more nodes.

Cluster: donphan (interactive/debug)

Time (hours): 12 (1/2 day)

Number of nodes: 1

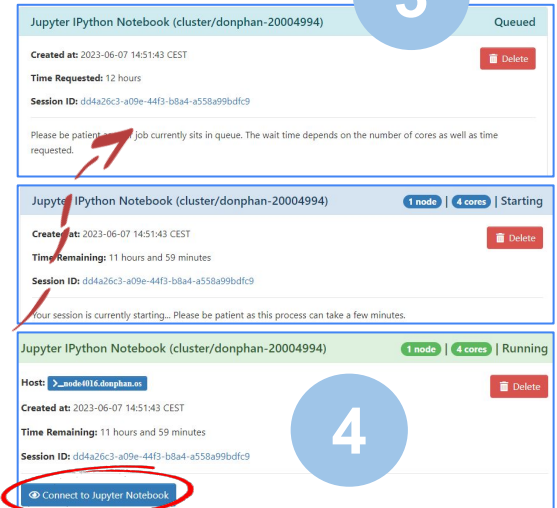
Number of cores per node: 4

IPython version: 8.5.0 GCCcore 11.3.0

Custom nodes

I would like to receive an email when the session is cancelled under the data root directory.

Launch



3

Jupyter IPython Notebook (cluster/donphan-20004994) Queued

Created at: 2023-06-07 14:51:43 CEST

Time Requested: 12 hours

Session ID: dd4a26c3-a09e-44f3-b8a4-a558a99bdfc9

Please be patient... job currently sits in queue. The wait time depends on the number of cores as well as time requested.

Jupyter IPython Notebook (cluster/donphan-20004994) 1 node | 4 cores | Starting

Created at: 2023-06-07 14:51:43 CEST

Time Remaining: 11 hours and 59 minutes

Session ID: dd4a26c3-a09e-44f3-b8a4-a558a99bdfc9

Your session is currently starting... Please be patient as this process can take a few minutes.

Jupyter IPython Notebook (cluster/donphan-20004994) 1 node | 4 cores | Running

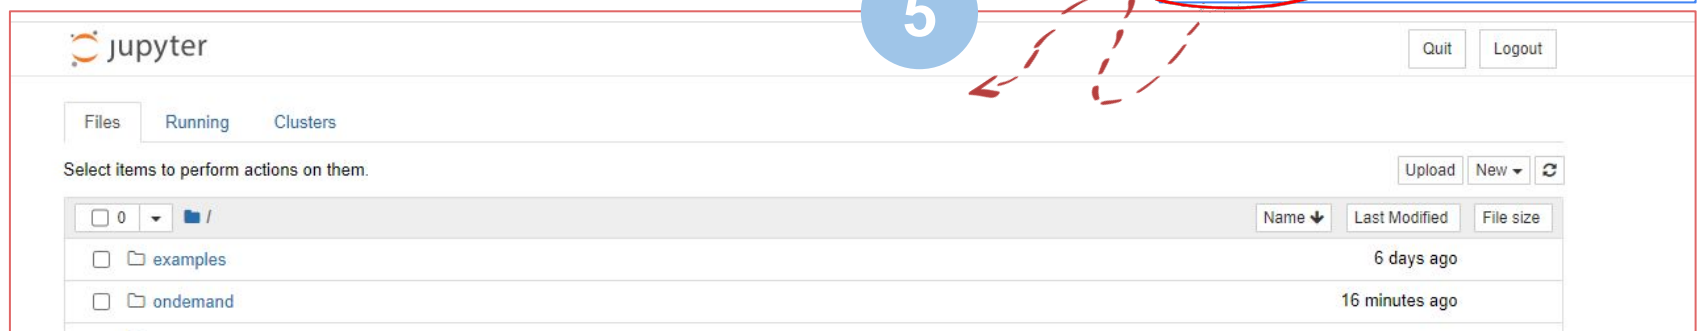
Host: node0116.donphan.us

Created at: 2023-06-07 14:51:43 CEST

Time Remaining: 11 hours and 59 minutes

Session ID: dd4a26c3-a09e-44f3-b8a4-a558a99bdfc9

Connect to Jupyter Notebook



5

jupyter

Quit Logout

Files Running Clusters

Select items to perform actions on them.

Upload New ↕

	Name	Last Modified	File size
<input type="checkbox"/>	0		
<input type="checkbox"/>	examples	6 days ago	
<input type="checkbox"/>	ondemand	16 minutes ago	

Beyond VSC: EuroHPC



EuroHPC
Joint Undertaking

- There is more compute power available beyond what VSC provides...
- Belgian researchers can also get access to EuroHPC JU supercomputers
- Overview: https://eurohpc-ju.europa.eu/supercomputers/our-supercomputers_en
- Project-based access, more competitive than VSC systems (but way more compute power)



LUMI, the Queen of the North: hardware



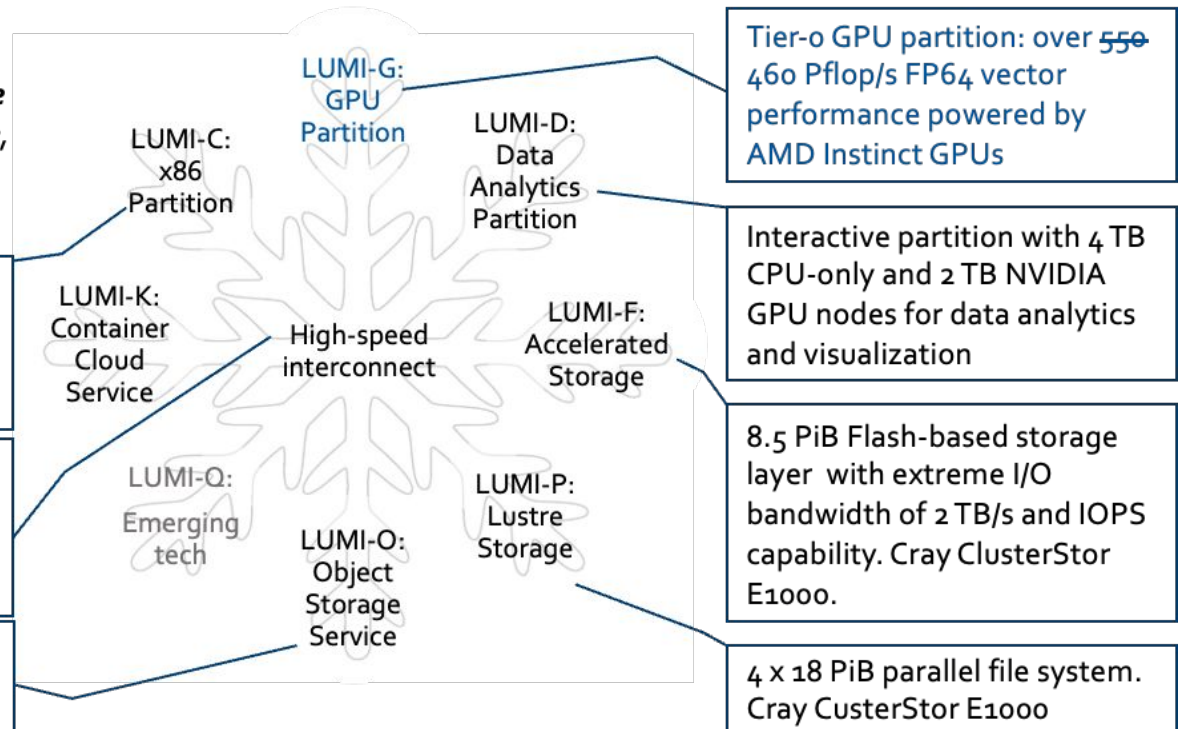
docs.lumi-supercomputer.eu

LUMI is a Tier-0 GPU-accelerated supercomputer that enables the convergence of high-performance computing, artificial intelligence, and high-performance data analytics.

- Supplementary CPU partition
- ~260,000 AMD EPYC CPU cores

Possibility for combining different resources within a single run. HPE Slingshot technology.

30 PB encrypted object storage (Ceph) for storing, sharing and staging data



Tier-0 GPU partition: over ~~550~~ 460 Pflop/s FP64 vector performance powered by AMD Instinct GPUs

Interactive partition with 4 TB CPU-only and 2 TB NVIDIA GPU nodes for data analytics and visualization

8.5 PiB Flash-based storage layer with extreme I/O bandwidth of 2 TB/s and IOPS capability. Cray ClusterStor E1000.

4 x 18 PiB parallel file system. Cray ClusterStor E1000

LUMI, the Queen of the North: hardware



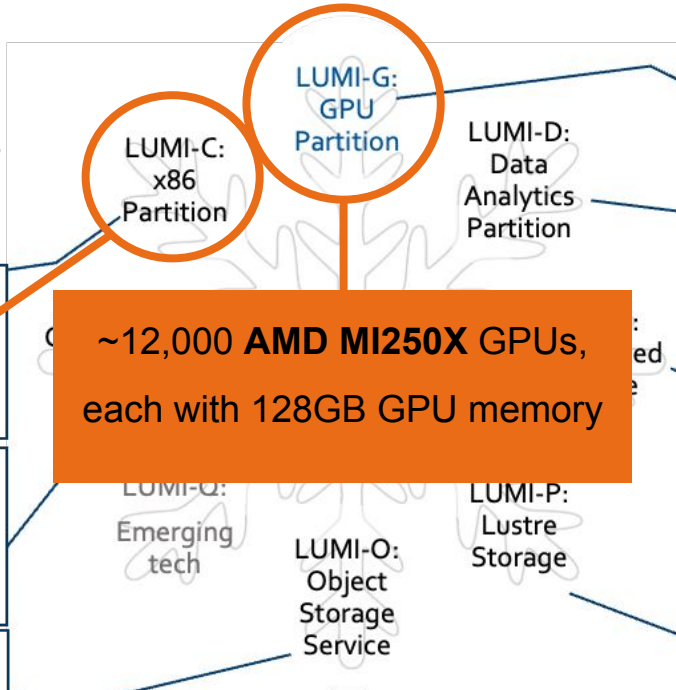
docs.lumi-supercomputer.eu

LUMI is a Tier-0 GPU-accelerated supercomputer that enables the convergence of high-performance computing, artificial intelligence, and high-performance data analytics.

- Supplementary CPU partition
- ~260,000 AMD EPYC CPU cores

Possibility for combining different resources within a single run. HPE Slingshot technology.

30 PB encrypted object storage (Ceph) for storing, sharing and staging data



~12,000 AMD MI250X GPUs, each with 128GB GPU memory

Tier-0 GPU partition: over ~~550~~ 460 Pflop/s FP64 vector performance powered by AMD Instinct GPUs


Interactive partition with 4 TB CPU-only and 2 TB NVIDIA GPU nodes for data analytics and visualization

8.5 PiB Flash-based storage layer with extreme I/O bandwidth of 2 TB/s and IOPS capability. Cray ClusterStor E1000.

4 x 18 PiB parallel file system. Cray ClusterStor E1000

LUMI, the Queen of the North: access

docs.lumi-supercomputer.eu



EuroHPC
Joint Undertaking

50% of LUMI reserved for consortium members, 50% for EuroHPC

Belgian share of LUMI: ~7.4%

Access for Belgian users: via LUMI-BE or EuroHPC


- LUMI-BE tries to sit between the Tier-1 systems and EuroHPC allocations
- Repeat users are expected to try the EuroHPC channel also

Separate channels for:

- Benchmarking to prepare a regular project application (4 months in most cases)
- Software development (1 year), Regular projects (1 year)
- EuroHPC extreme scale projects (1 year) aimed primarily at projects with very scalable codes and very large simulations



LUMI, the Queen of the North: access



EuroHPC
Joint Undertaking

docs.lumi-supercomputer.eu

CPU-only resources available for Belgium are limited compared to Tier-1

- Belgian share roughly equivalent of a 150-node cluster

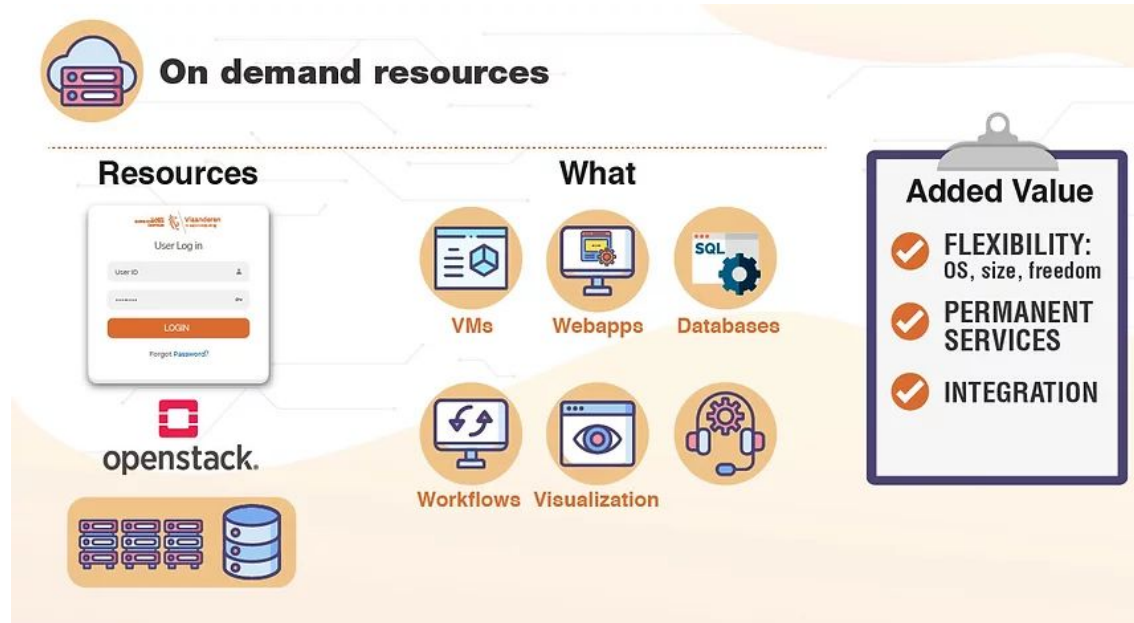
Access via project applications

- Scientific and technical review for EuroHPC regular and extreme scale projects
- Only technical review for LUMI-BE projects in case of corresponding reviewed scientific project
- For open research only, otherwise buying compute time is possible
- Participation in the promotion expected as LUMI needs to show to funding agencies that access to large machines is important

Questions related to LUMI, contact lumi-be-support@enccb.be

VSC Tier-1 cloud

- **Project-based access**
- Free of charge
- **Self-managed virtual machines**
- For use cases that are not a good fit for compute clusters
- More info: vscentrum.be/cloud
- Contact: cloud@vscentrum.be



VSC Tier-1 cloud: instance types

UPSV1

- AMD EPYC 7542
- Oversubscribed vCPU (2:1)
- Max. 16 vCPUs
- Max. ~120GB RAM
- 25Gbit Ethernet
- No GPU
- Uninterruptible power supply (UPS)

CPUv1

- Intel Xeon E5-2670
- 10Gbit Ethernet
- Max. 20 vCPUs
- Max. ~360GB RAM
- No GPU

GPUv2

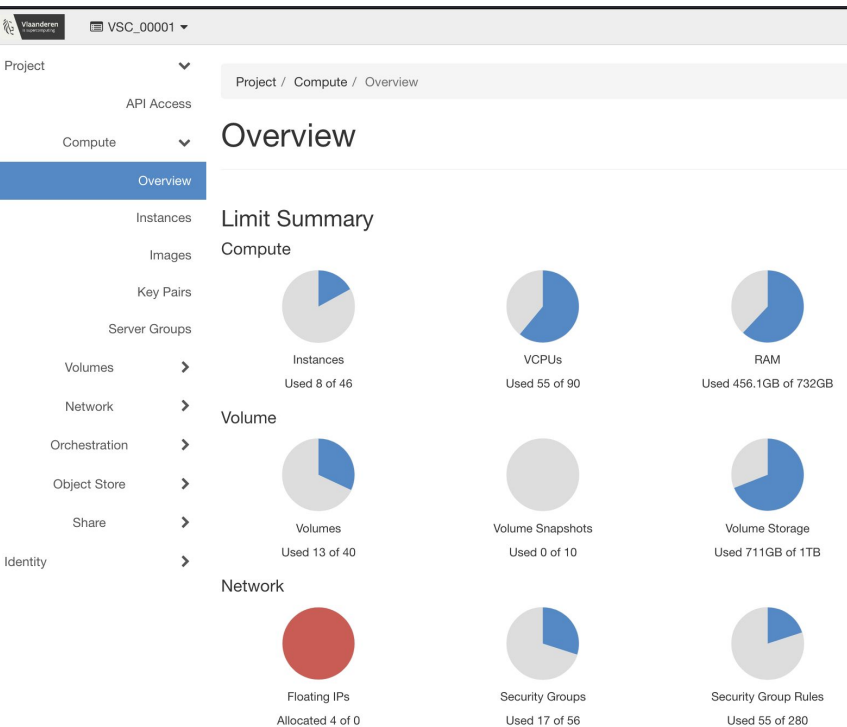
- AMD EPYC 7542
- Max. 16 vCPUs
- Max. ~60GB RAM
- 25Gbit Ethernet
- 1 NVIDIA T4 GPU (16GB, dedicated)

GPUv3

- AMD EPYC 7542
- Max. 16 vCPUs
- Max. ~60GB RAM
- 25Gbit Ethernet
- 1 NVIDIA A2 GPU (16GB, dedicated)
- Uninterruptible power supply (UPS)

VSC Tier-1 cloud: dashboard

<https://cloud.vscentrum.be>



Instances

Displaying 8 items

Instance Name	Image Name	IP Address	Flavor
demo-vaia	AlmaLinux-8	10.113.1.56	GPUv3.2xlarge
MyVM-window-s-Cintia	Windows 10	10.113.1.90	CPUv1.4xlarge
MyVM-window-s-Jan	Windows 10	10.113.1.167	CPUv1.medium
wp-test-4	Windows 10	10.113.1.163, 10.113.1.207	CPUv1.xlarge

Instance ID =

Flavor Details: GPUv3.2xlarge

ID	Value
ID	23
VCPUs	16
RAM	60GB
Size	40GB

VSC Tier-1 cloud: using Terraform

- VMs can be created manually via dashboard, but **network configuration is very tricky**
- It is recommended to use **Terraform** for creating VMs
- `terraform` command is available on HPC-UGent Tier-2 login nodes
- Easy to install on your own system (see terraform.io/downloads)
- Documentation + examples are available, support available via cloud@vscentrum.be

Demo: creating a virtual machine in Tier-1 cloud




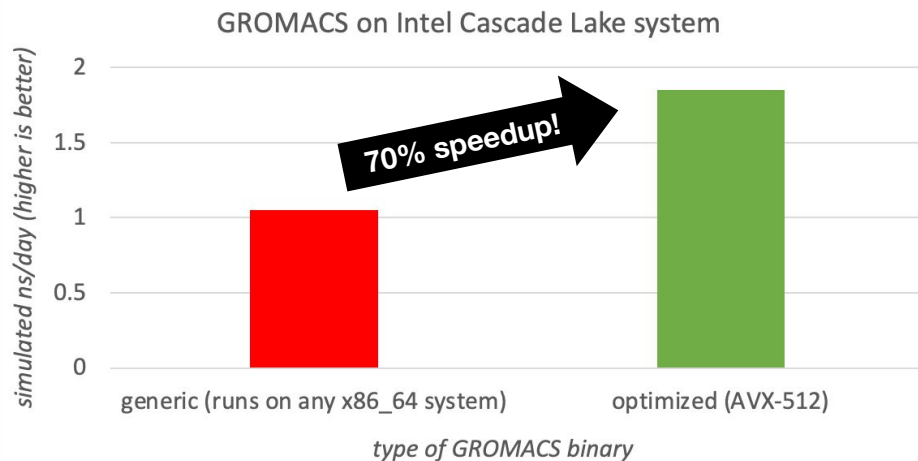
```
# first create application credentials for Terraform!  
git clone https://github.com/hpcugent/openstack-templates  
cd openstack-templates  
  
# edit terraform/scripts/modify_variable.config (change flavor + image)  
# edit terraform/environment/main.tf.template  
  
cd terraform/scripts; ./modify_variable.sh  
  
cd ../environment  
  
terraform init  
  
terraform plan  
  
terraform apply
```

Questions?



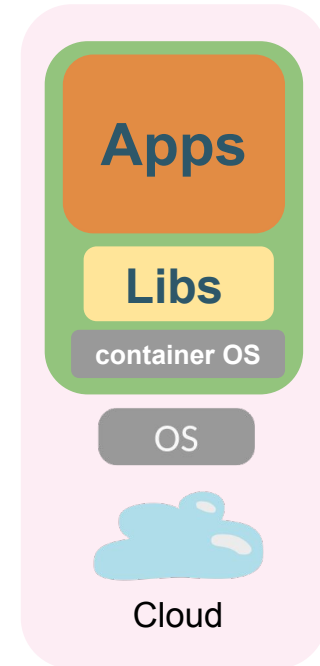
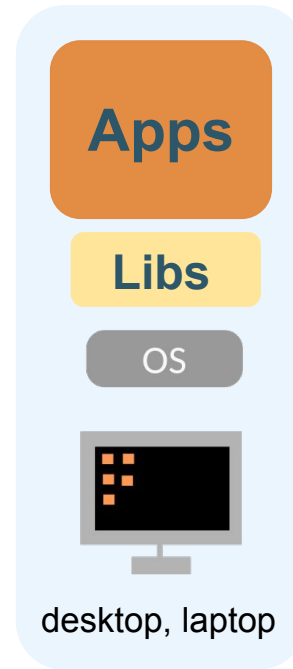
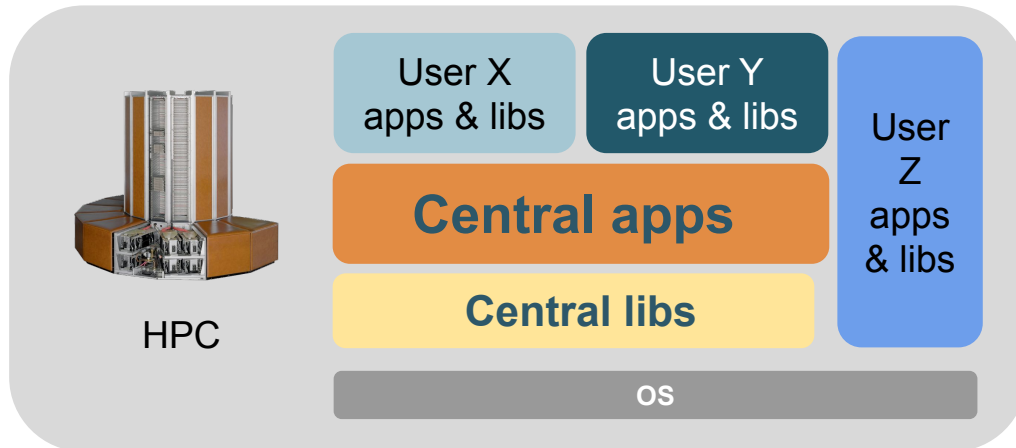
Scientific software is a different breed

- Scientists are often not trained software engineers (and that's OK)
- They often need help to get the software installed properly (especially on HPC systems)
- **Standard packaging tools do not suffice, we want need to build from source**
- Scientific software should be compiled for system on which it will be used (performance!)

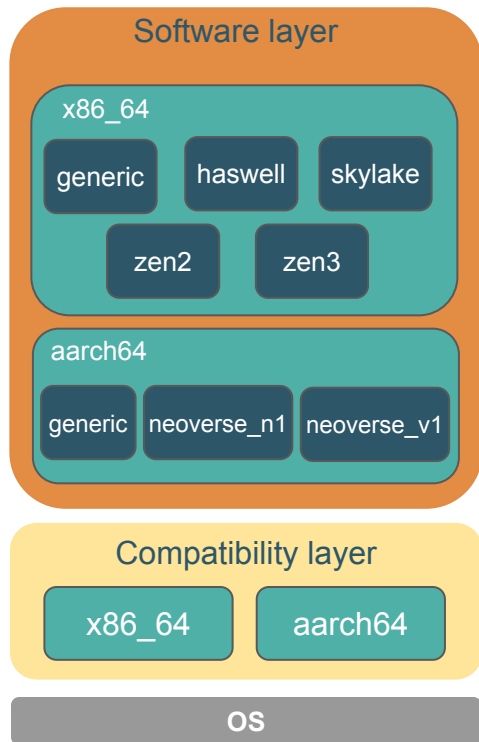


We need to collaborate more!

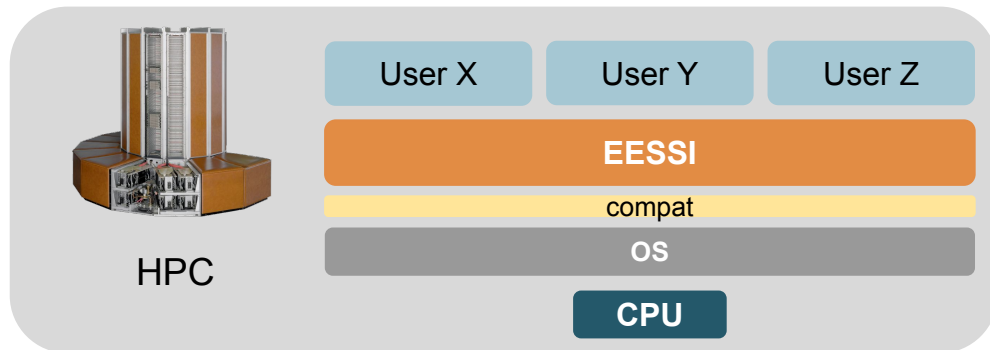
- **Too much software** for a single support team to handle
- Different systems (CPU, OS, ...) => different problems
- **EasyBuild is not sufficient anymore...**
- **Duplicate work** across HPC sites and scientists
- **Diverse software stacks** across different platforms



EESSI to the rescue



E E S S I
EUROPEAN ENVIRONMENT FOR
SCIENTIFIC SOFTWARE INSTALLATIONS



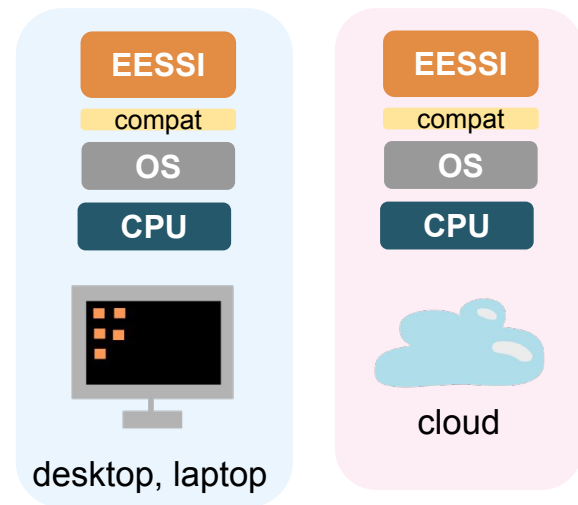
**Shared repository of
(optimized) scientific
software installations**

**Same software stack
everywhere!**

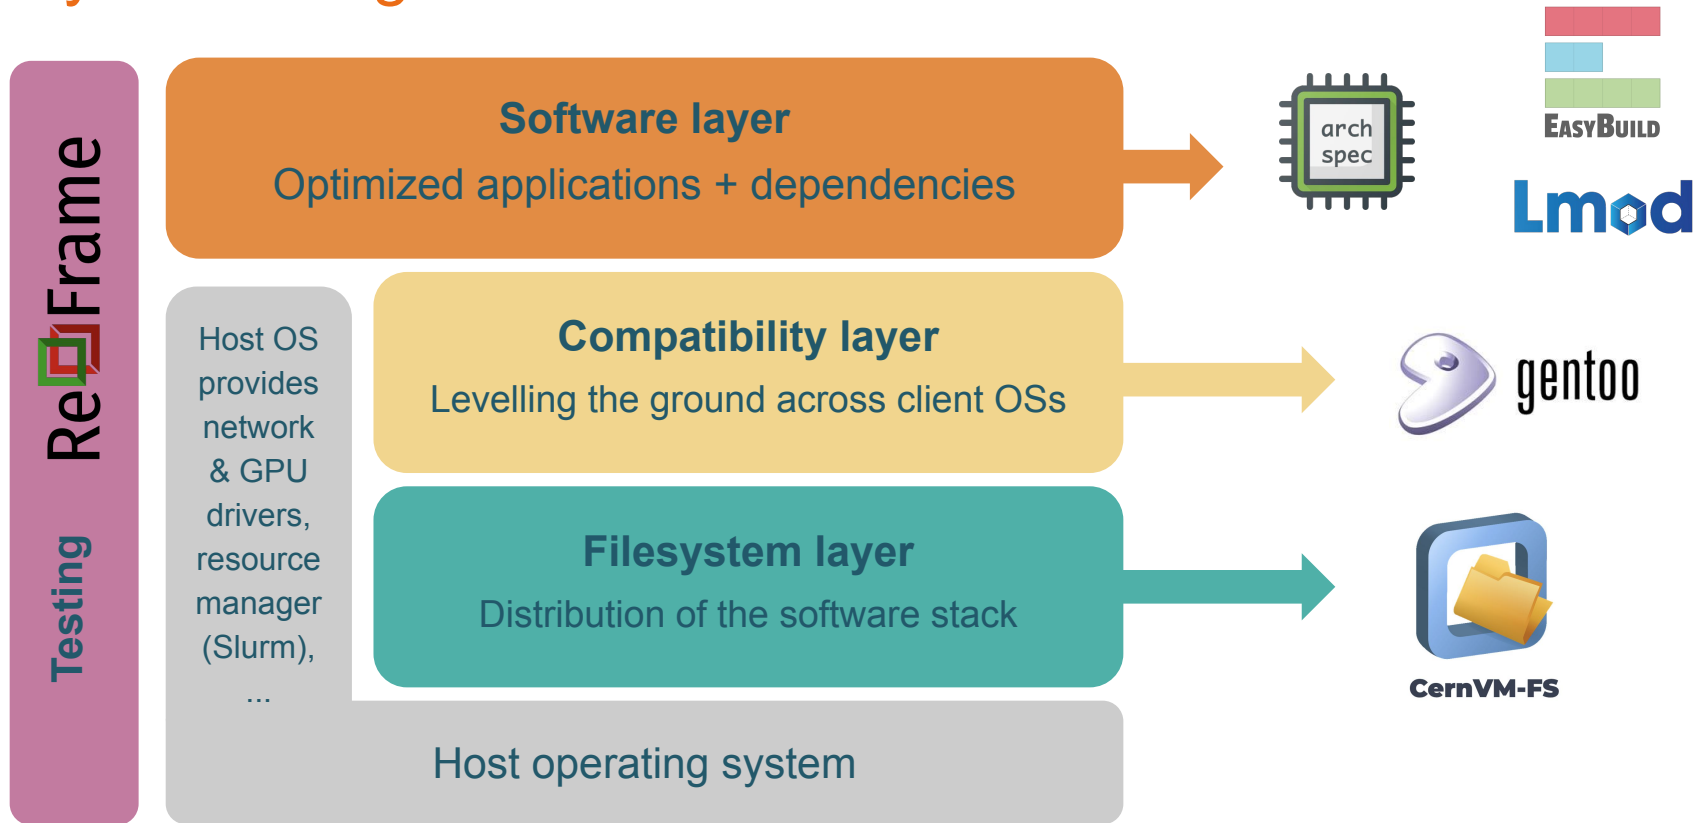
eessi.io

eessi.io/docs (try out the pilot setup!)

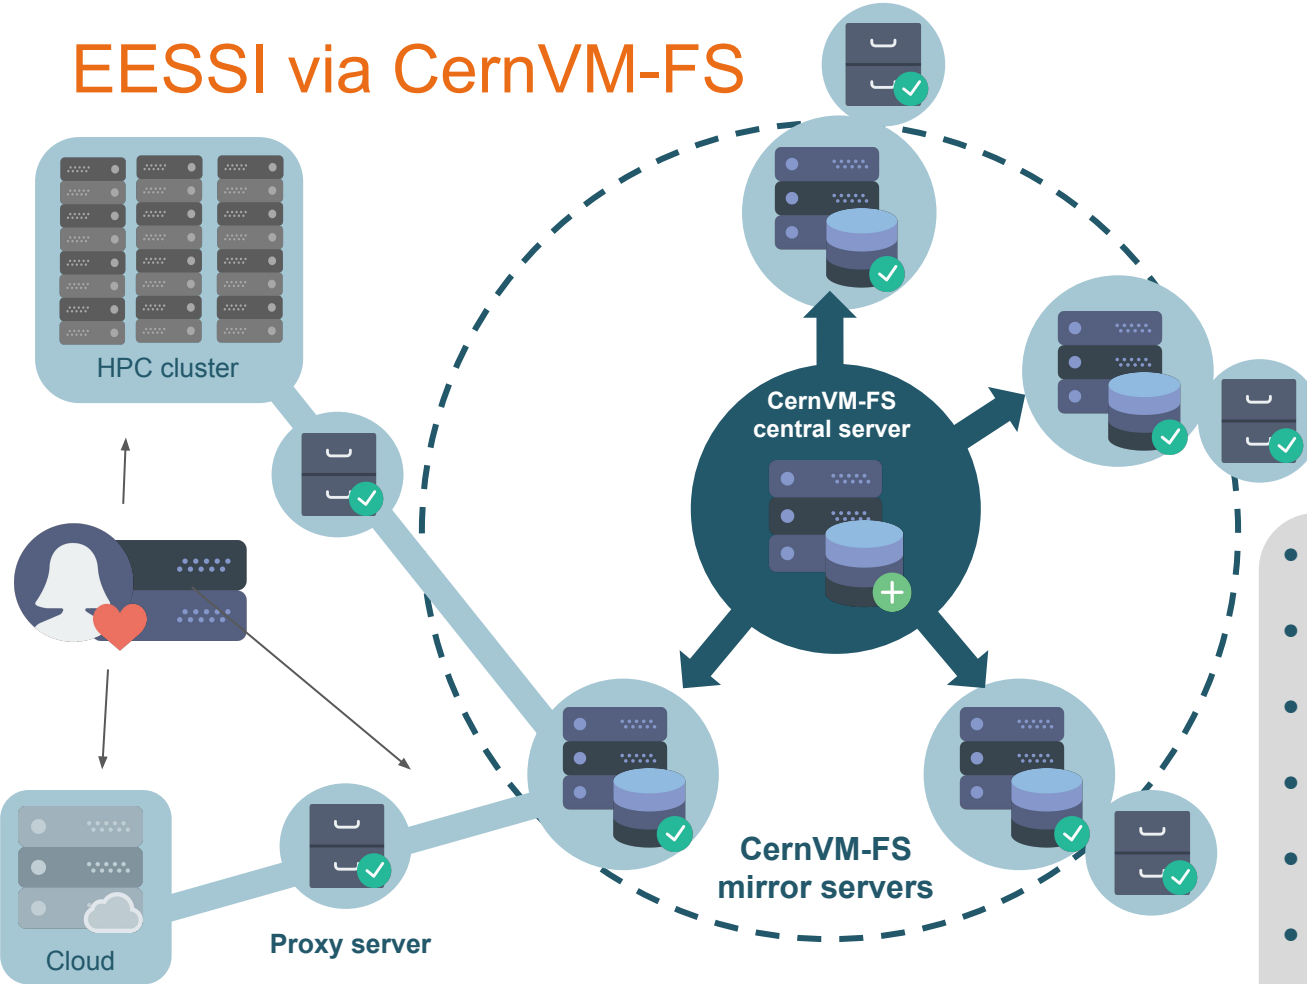
eessi.io/docs/support



Layered design of EESSI



EESSI via CernVM-FS



CernVM-FS

cvmfs.readthedocs.io

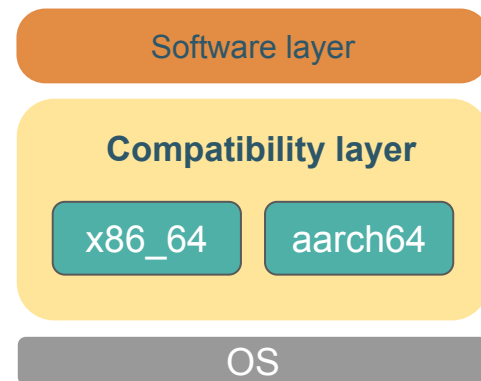
- Global distribution of software installations
- Centrally managed software stack
- Redundant network of “mirrors”
- Multiple levels of caching
- Clients automatically get software updates
- **Same software stack everywhere:**
laptops, HPC clusters, cloud VMs, ...

EESSI compatibility layer



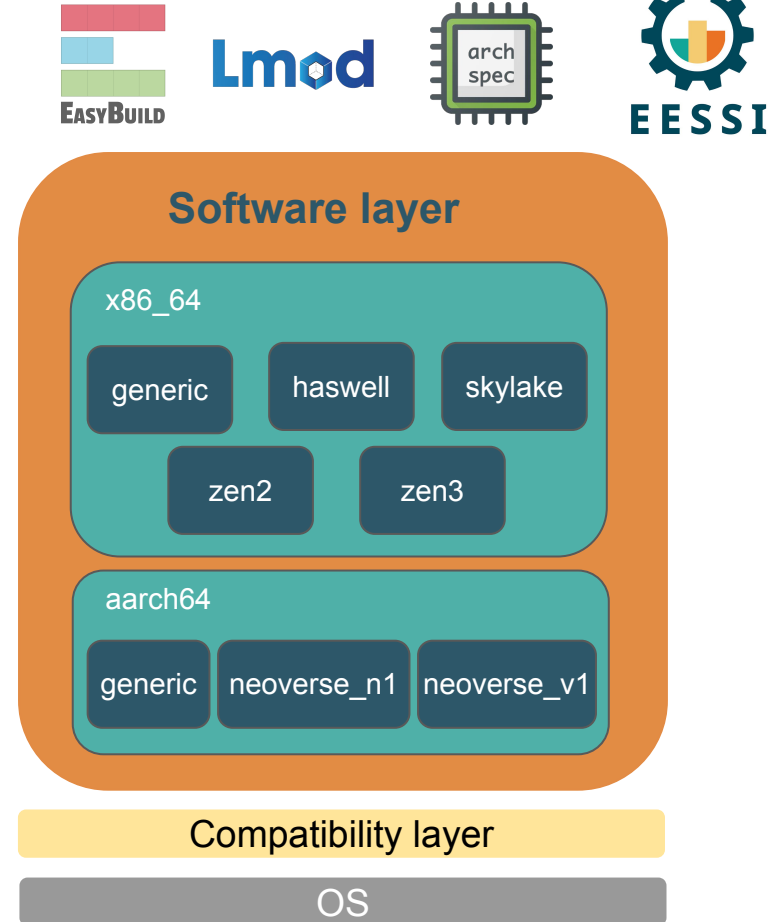
- “*Containers without the containing*”
- **Minimal collection of tools and libraries** (incl. glibc, bash, Python, Lmod, ...)
- **Built from source per CPU family** (x86_64, aarch64, ...) with Gentoo Prefix
- Installations included in software layer **only link to compat layer** (RPATH)
- Ensures **compatibility** with any client system running Linux

```
$ ls /cvmfs/software.eessi.io/versions/2023.06/compat/linux/aarch64/  
bin etc lib lib64 opt reprod run sbin stage1.log stage2.log  
stage3.log startprefix tmp usr var
```



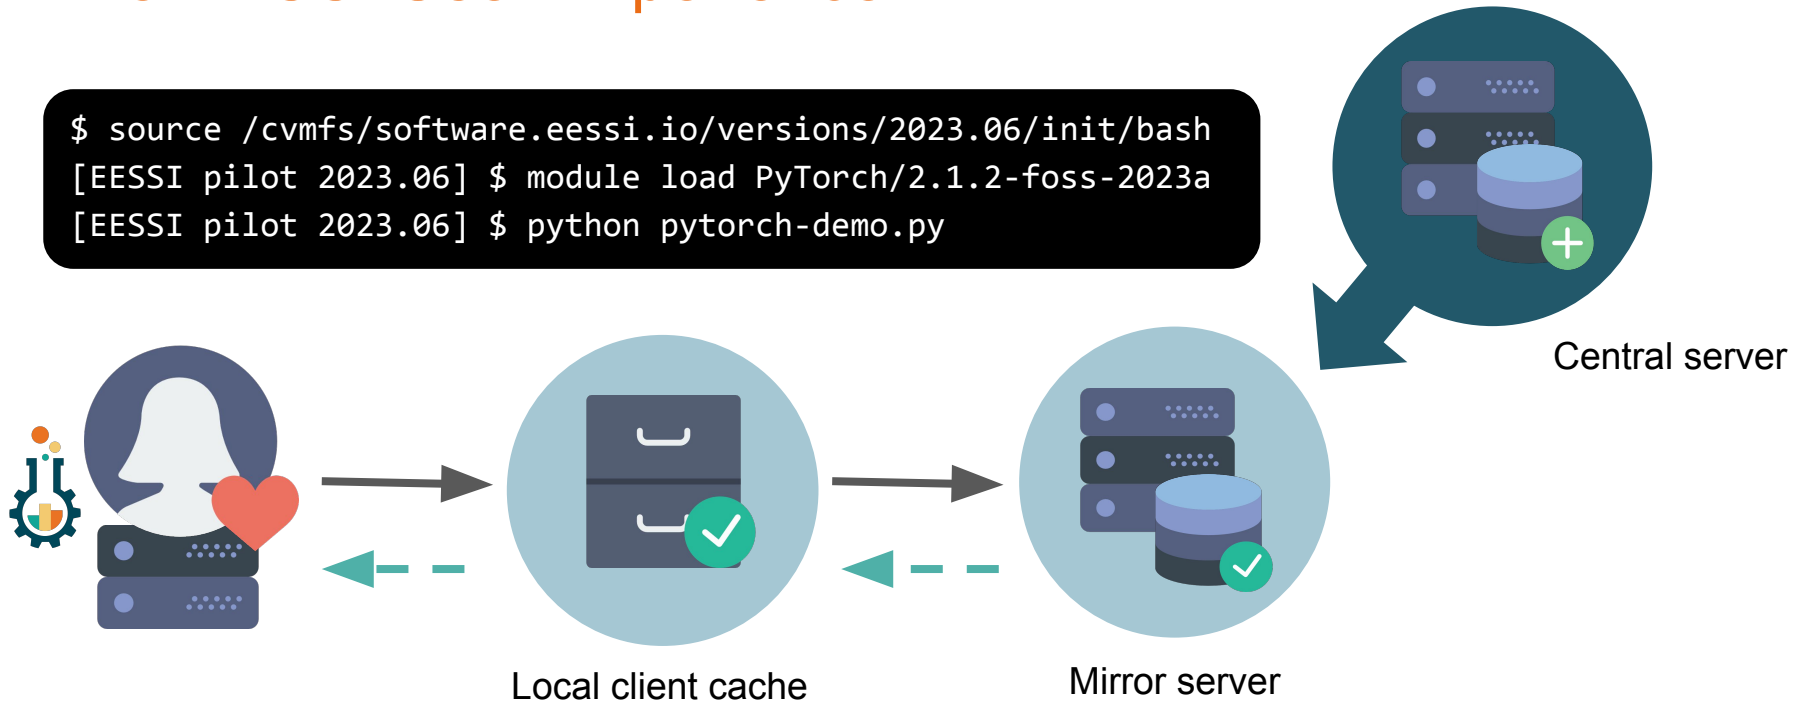
EESSI software layer

- Installations of scientific software applications
- **Optimized for specific CPU targets**
- Works on any client system running Linux, since we only link to libraries in compat layer
- Built using EasyBuild
- Environment modules as user interface (via Lmod)
- Detection of host CPU via archspec
- **Best subset of software installations for host CPU is automatically selected**






The EESSI User Experience

```
$ source /cvmfs/software.eessi.io/versions/2023.06/init/bash  
[EESSI pilot 2023.06] $ module load PyTorch/2.1.2-foss-2023a  
[EESSI pilot 2023.06] $ python pytorch-demo.py
```



EESSI provides **on-demand streaming**
of (scientific) software (like music, TV-series, ...)

- EuroHPC Centre of Excellence: 4 year project (2023-2026), €6M budget (50% for EESSI)
- Collaboration between EESSI and CECAM: total of 16 partners (academic + industry)
- EESSI focuses on technical aspects: providing a shared stack of scientific software
- Scientific target: multiscale simulations with 3 key use cases
 - Helicopter design and certification for civil transport 
 - Battery applications to support the sustainable energy transition 
 - Ultrasound for non-invasive diagnostics and biomedical applications 

Demo: getting access to EESSI (in Tier-1 cloud VM)

- Getting access to EESSI is... easy
- Only need to install + configure CernVM-FS (eessi.io/docs/getting_access)
- Can also access EESSI via a Singularity/Apptainer container image
- Rich software stack, including *optimized* installations of TensorFlow, PyTorch, ...
- Initial GPU support available (eessi.io/docs/gpu)
- Check out EESSI demo script for TensorFlow (eessi.io/docs/using_eessi/eessi_demos)
- Attractive for VMs in Tier-1 cloud, similar user experience as on VSC clusters



Demo: using EESSI

eessi.io/docs/using_eessi/eessi_demos



```
/cvmfs/software.eessi.io/versions/2023.06/software
```

```
`-- linux
  |-- aarch64
  |   |-- generic
  |   |-- neoverse_n1
  |   `-- neoverse_v1
  `-- x86_64
      |-- amd
      |   |-- zen2
      |   `-- zen3
      |-- generic
      `-- intel
          |-- haswell
          `-- skylake_avx512
              |-- modules
              `-- software
```

```
$ source /cvmfs/software.eessi.io/versions/2023.06/init/bash
Found EESSI pilot repo @
/cvmfs/software.eessi.io/versions/2023.06!
```

```
archspeg says x86_64/amd/zen2
Using x86_64/amd/zen2 as software subdirectory
```

```
...
Environment set up to use EESSI pilot software stack, have fun!
```

```
$ module load R/4.3.2-gfbb-2023a
```

```
$ which R
/cvmfs/software.eessi.io/versions/2023.06/software/linux/x86_64/
amd/zen2/software/R/4.3.2-gfbb-2023a/bin/R
```

```
$ R --version
R version 4.3.2
```

Questions?



Demo: Running an LLM on Tier-1 cloud

Let's try and get a trained LLM “deployed” in a Tier-1 cloud VM!

- Using PyTorch installation provided by EESSI + FastAPI
- Goal is to provide a REST API to which we can submit queries, and get a response
- Caveat: by default only port 80 is forwarded for HTTP, need to be changed in `modules/vm_with_pf_rules_with_ssh_access_with_nginx/main.tf`

plainenglish.io/blog/your-local-llm-using-fastapi

Demo: Running an LLM on Tier-1 cloud



Step 1: install all necessary software in a Python virtual environment, on top of EESSI

```
source /cvmfs/software.eessi.io/versions/2023.06/init/bash

module load PyTorch/2.1.2-foss-2023a

python3 -m venv llm_demo

source llm_demo/bin/activate

pip install fastapi uvicorn transformers

export PYTHONPATH=$PWD/llm_demo/lib/python3.11/site-packages:$PYTHONPATH
```

plainenglish.io/blog/your-local-llm-using-fastapi

Demo: Running an LLM on Tier-1 cloud



Step 2: “deploy” the LLM model (quick and dirty)

```
# copy-paste Python code in file named app.py  
  
# run app via uvicorn tool (ideally in a screen session)  
  
IP=193.190.80.3 # change this to match your VM!  
  
uvicorn --host $IP app:app
```

plainenglish.io/blog/your-local-llm-using-fastapi

Demo: Running an LLM on Tier-1 cloud

Step 3: query the LLM via the REST API

```
IP_PORT=193.190.80.3:58282 # change this to match your VM!  
  
curl -X POST -H "Content-Type: application/json" -d '{"prompt": "A  
supercomputer is"}' http://${IP_PORT}/generate
```

plainenglish.io/blog/your-local-llm-using-fastapi