



**doitnow**

HPC Services

# OOD Meets EESSI:

**Accessing and Distributing  
Scientific Software with Ease**

Christian Bustelo  
R&D - Do IT Now

Arturo Gimeno  
Support - Do IT Now



# A global HPC Services Company

---

Servicing HPC customers world-wide

- 100% independent capital (Self-funded company)
- 140+ HPC experts** at your service
- 30+ years of experience** in the EMEA HPC Market
- 200+ clients through all industry verticals** (Aereospace, Automotive, Chemical, Energy, FSI, Life Sciences, Manufacturing, Oil & Gas)
- Services managed
  - 12k+ users/year**
  - 150+ clusters/year**
  - 200+ Training sessions/year**
- Installation services
  - 5 clusters in the Top 500**

## Do IT Now Team

---

### EMEA:

France (Montpellier) – (22 people)

Germany (Munich) – (2 people)

Italy (Torino, Maranello) – (60 people)

Spain (Barcellona) – (41 people)

### APAC:

New Zealand (Auckland) – (11 people)

**USA:** 3 people

**BRASIL:** 4 people

**Total 140+ people**

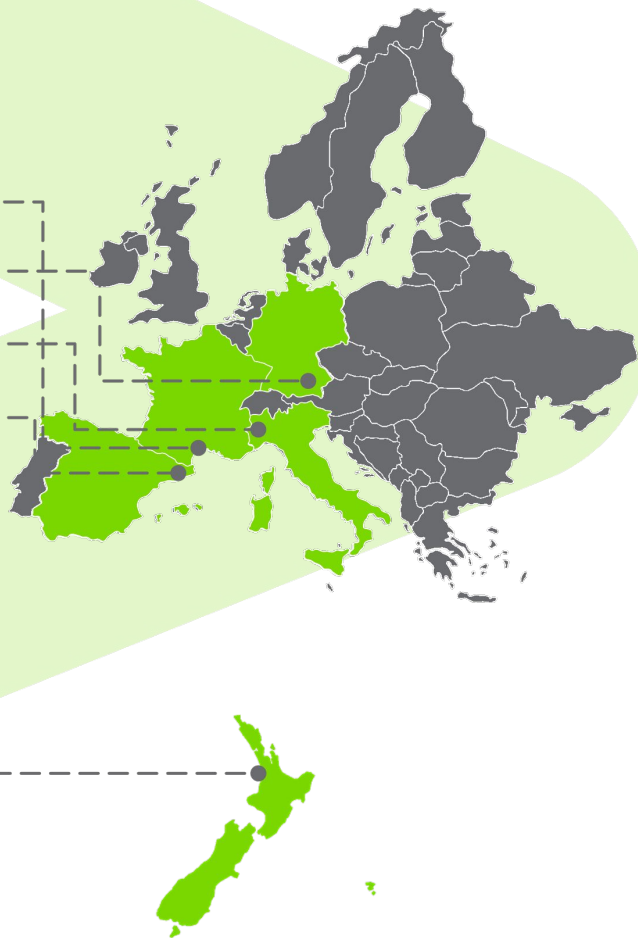
France - Montpellier

Germany - Munich

Italy - Turin


Spain - Barcelona

New Zealand - Auckland




## What we offer?


---



CONSULTING



INSTALLATION



SUPPORT



R&D-as-a-Service

## Open Ondemand & EESSI

---

**OPEN**

### 1. nDemand:

- Make HPC resources more accessible for everyone + Point of gathering of multiple types of resources.



### 2. **E E S S I:**

- EESSI is a shared repository of optimized scientific software installations



# MultiXscale Centre-of-Excellence in a nutshell

- 4-year project (started in Jan 2023), ~€6M budget
- Collaboration between EESSI and CECAM (total of 16 partners)
  - EESSI primarily addresses technical aspects
  - CECAM network provides scientific expertise
- 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



**Multi**  **scale**

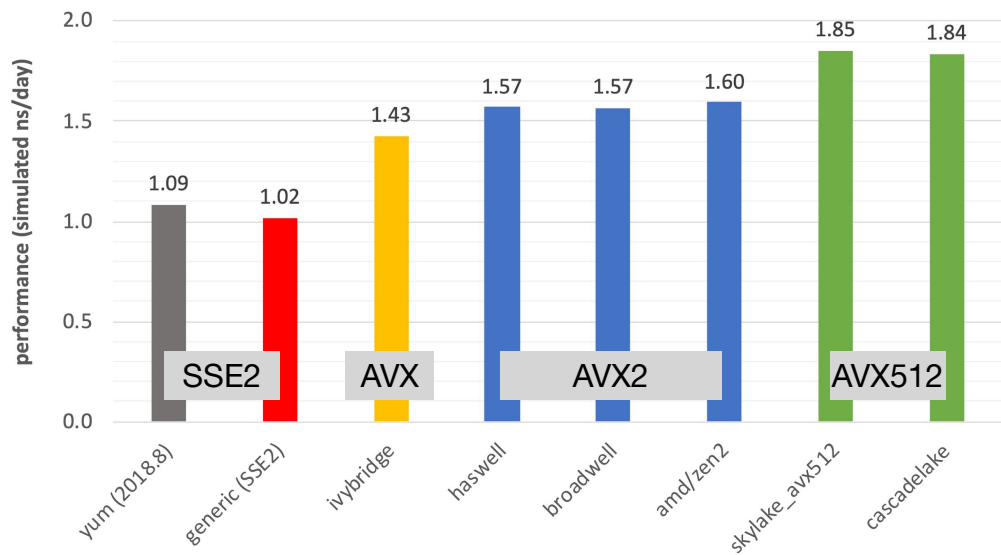
<https://multixscale.eu>

# The changing landscape of scientific computing

- Explosion of available scientific software applications (bioinformatics, AI boom, ...)
- Increasing interest in **cloud** for scientific computing (flexibility!)
- **Increasing variety in processor (micro)architectures** beyond Intel & AMD:
  - Arm is already here (see Fugaku, JUPITER, ...), RISC-V is coming (soon?)
- In strong contrast: available (wo)manpower in **HPC support teams is (still) limited...**

## Optimized scientific software installations

- Software should be optimized for the system it will run on (keep the P in HPC!)
- Impact on performance is often significant for scientific software!
- Example: GROMACS 2020.1 (PRACE benchmark, Test Case B)
- Metric: (simulated) ns/day, higher is better
- Test system: dual-socket Intel Xeon Gold 6420 (Cascade Lake, 2x18 cores)
- Performance of different GROMACS binaries, on exact same hardware/OS





What if you no longer have to install  
a **broad range of scientific software**  
from scratch on every laptop, HPC cluster,  
or cloud instance you use or maintain,  
**without compromising on performance?**



## EESSI in a nutshell

---

- European Environment for Scientific Software Installations (**E E S S I**)
- **Shared repository of (optimized!) scientific software installations**
- Uniform way of providing software to users, regardless of the system they use!
- Should work on any Linux OS (+ WSL, macOS via Lima) and system architecture
- From laptops and personal workstations to HPC clusters and cloud
- Support for different CPU (micro)architectures, interconnects, GPUs, etc.
- **Focus on performance, automation, testing, collaboration**


## Major goals of EESSI

---

- Avoid duplicate work (for researchers, HPC support teams, sysadmins, ...)
  - Tools that automate software installation process (EasyBuild, Spack) are not sufficient anymore
  - Go beyond sharing build recipes, work towards a shared software stack
- Providing a truly uniform software stack
  - Use the (exact) same software environment everywhere
  - Without sacrificing performance for “mobility of compute” (like is typically done with containers/conda)
- Facilitate HPC training, development of (scientific) software, ...


Testing  
**ReFrame**

**Software layer**  
Optimized applications + dependencies




Host OS provides network & GPU drivers, resource manager (Slurm), ...

**Compatibility layer**  
Levelling the ground across client OSs



gentoo

**Filesystem layer**  
Distribution of the software stack

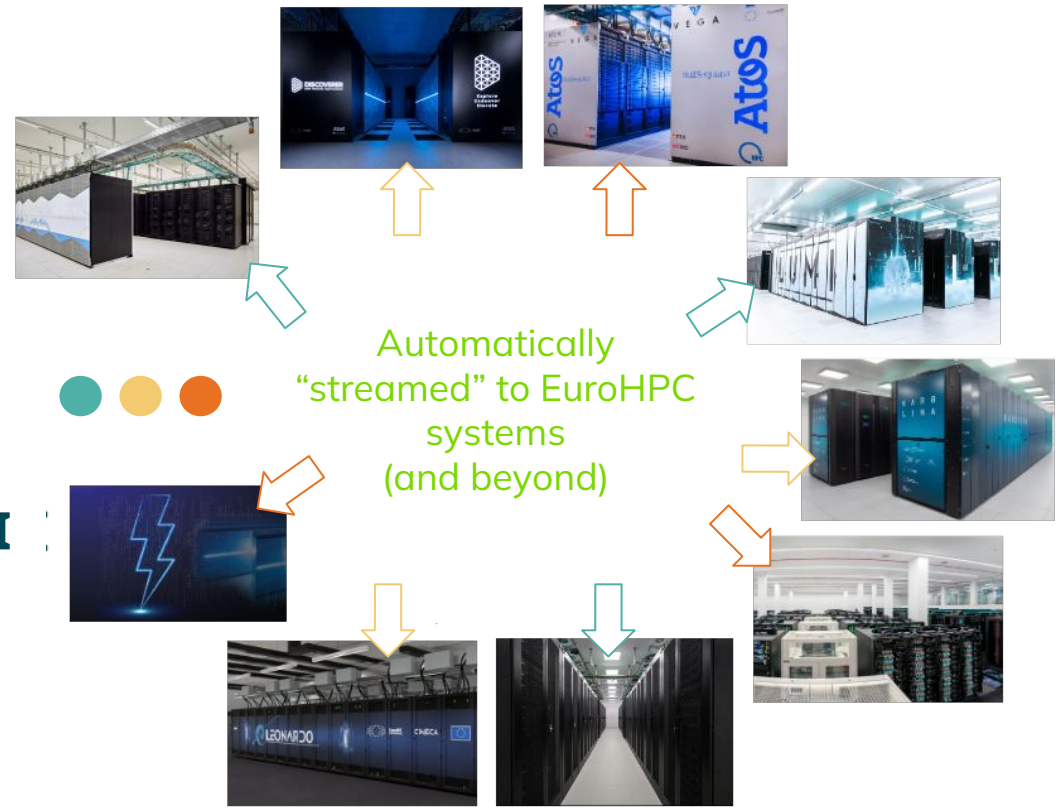


CernVM-FS

Host operating system



**E E S S I**  
EUROPEAN ENVIRONMENT FOR  
SCIENTIFIC SOFTWARE INSTALLATIONS

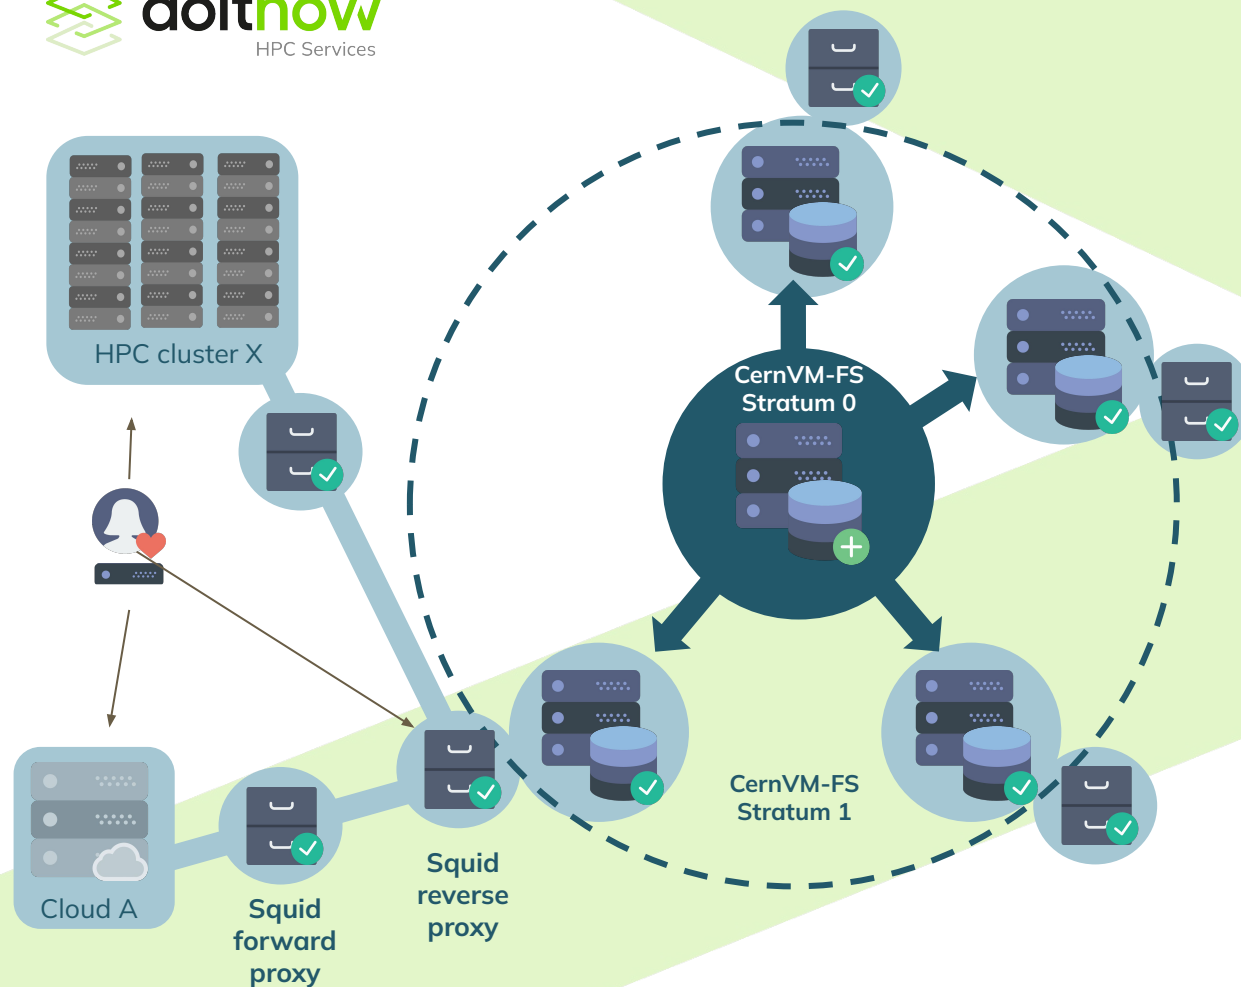





**CernVM-FS**

[cvmfs.readthedocs.io](https://cvmfs.readthedocs.io)

[github.com/EESSI/filesystem-layer](https://github.com/EESSI/filesystem-layer)



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

## EESSI user experience

```
$ source /cvmfs/software.eessi.io/versions/2023.06/init/bash
archdetect says x86_64/intel/haswell
archdetect could not detect any accelerators
Using x86_64/intel/haswell as software subdirectory.
{EESSI 2023.06} $ module load GROMACS/2024.4-foss-2023b
{EESSI 2023.06} $ gmx --version
:-) GROMACS - gmx, 2024.4-EasyBuild_4.9.4 (-:
```



Local client cache

Mirror server

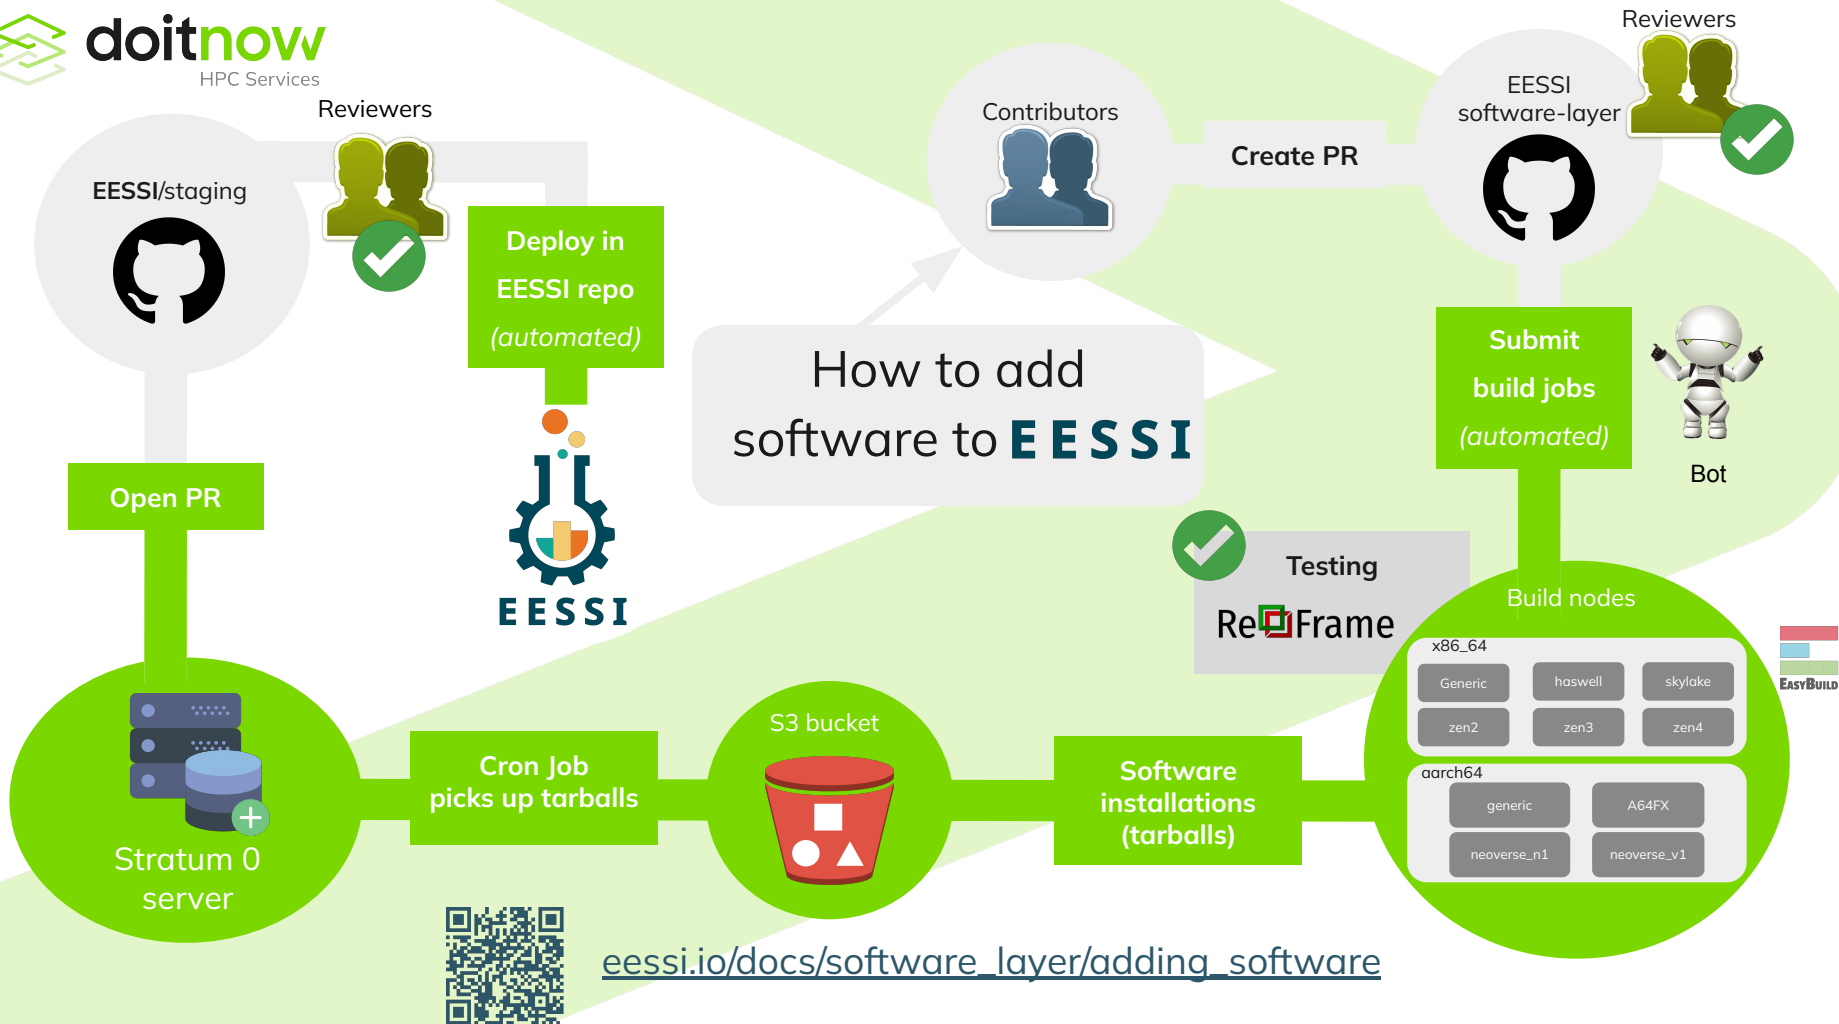
Central server

## Supported system architectures

---

- Different generations of x86\_64 (Intel, AMD) and Arm 64-bit CPUs; RISC-V is WIP
  - Including A64FX (Deucalion, WIP) & NVIDIA Grace (JUPITER, coming soon)
  - Also works on laptops, in virtual machines in the cloud, on Raspberry Pi boards, etc.
- Different accelerators: NVIDIA GPUs (today) + AMD GPUs (soon)
  - For now, only software installations for AMD Rome (Zen2) + NVIDIA A100 are available
- Various interconnects like Infiniband, via “fat” MPI libraries
  - Support for injecting a vendor-provided MPI library is available
- Goal is to support system architecture of all (current & future) EuroHPC systems





[eessi.io/docs/software\\_layer/adding\\_software](https://eessi.io/docs/software_layer/adding_software)

## How about integration?

**OPEN**  **nDemand**

## Potential use cases

---

### 1. Training

- Offer continuous formation without being limited by environments



### 2. Deploying

- Avoiding dependences of containers and or continuous installations



### 3. Sharing software

- Improve collaboration and scientific software development and sharing between centers



# Training

---

**OPEN**

- **OnDemand** project → Training?
  - User training for learn how to work with OOD
  - Sysadmin training for learn how to manage and solve problems
  
- Currently using the cluster of the client
  - Problems in case of full installation
  - Difficult to have full control of the environment
  - Sometimes not enough resources for big trainings



# Training



## 1. Setup Environment

Prepare a cluster structure

We can prepare a basic cluster environment with cloud providers (Aws, Azure...)



## 2. Prepare OOD

Automatized installation

With automatized installers for Open OnDemand we can deploy and launch over base cluster



## 3. Software Access

Enable use of Software

Enabling EESSI access for the compute nodes allows us to have access to the scientific software from the OOD apps

## Training

---



- Avoid having local cluster and give access to everyone every different training .
- Avoid having configured clusters on the cloud unused.
- No dependencies on the current status of the client cluster



# Deploying

---

- New applications on the portal means:
  - Prepare Application in OOD Server
  - Installation of Software on the compute nodes
  
- Compute nodes:
  - Classical installation
  - EasyBuild Modules
  - Containers and Images



**EASYBUILD**.io  
building software with ease



# Deploying

Home / My Interactive Sessions / Jupyter Notebook EESSI

### Interactive Apps

- Desktops
  - Classiq
  - Desktop XFCE-GNOME
  - Desktop w.Container
- GUIs
  - Avogadro2
  - DIA
  - MATLAB
  - ParaView

## Jupyter Notebook EESSI

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

Account:

Number of hours:

I would like to receive an email when the session starts

\* The Jupyter Notebook EESSI session data for this session can be accessed under the data root directory.

```
GNU nano 2.9.8                                     template/script.sh.erb

#!/usr/bin/env bash

# Benchmark Info
echo "TIMING - Starting main script at: $(date)"

# Set working directory to home directory
cd "${HOME}"

#
# Start Jupyter Notebook Server
#

# Purge the module environment to avoid conflicts
module purge

# Prepare EESSI
source /cvmfs/software.eessi.io/versions/2023.06/init/bash

# Load the module
module load JupyterNotebook

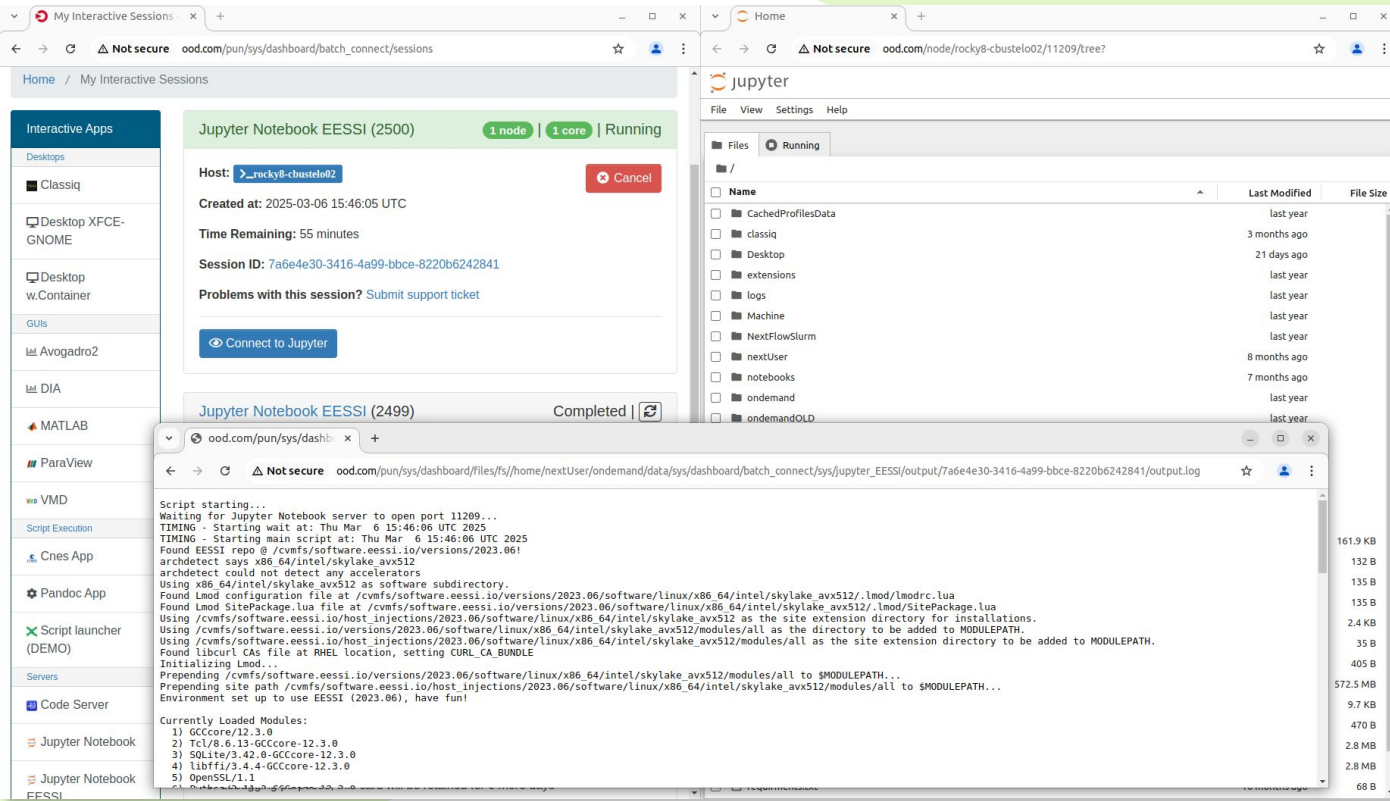
# List loaded modules
module list
[]

# Benchmark Info
echo "TIMING - Starting jupyter at: $(date)"

# Launch the Jupyter Notebook Server
#set -x
jupyter notebook --config="${CONFIG_FILE}" <%= context.extra_jupyter_args %>
```



# Deploying



The screenshot displays a web dashboard for managing interactive sessions. On the left, a sidebar lists various applications including Desktops, Classiq, Desktop XFCE-GNOME, Desktop w.Container, GUIs, Avogadro2, DIA, MATLAB, ParaView, VMD, Script Execution, Cnes App, Pandoc App, Xcrt launcher (DEMO), Servers, and Code Server. The main content area shows two Jupyter Notebook sessions:

- Jupyter Notebook EESSI (2500)**: Status: **Running** (1 node, 1 core). Host: `>_rocky8-cbustelo02`. Created at: 2025-03-06 15:46:05 UTC. Time Remaining: 55 minutes. Session ID: `7a6e4e30-3416-4a99-bbce-8220b6242841`. A **Connect to Jupyter** button is visible.
- Jupyter Notebook EESSI (2499)**: Status: **Completed**.

The log for the completed session (2499) is shown in a separate window, detailing the startup process:

```
Script starting...
Waiting for Jupyter Notebook server to open port 11209...
TIMING - Starting wait at: Thu Mar 6 15:46:06 UTC 2025
TIMING - Starting main script at: Thu Mar 6 15:46:06 UTC 2025
Found EESSI repo @ /cvmfs/software.eessi.io/versions/2023.06!
archdetect says x86_64/intel/skylake_avx512
archdetect could not detect any accelerators
Using x86_64/intel/skylake_avx512 as software subdirectory.
Found Lmod configuration file at /cvmfs/software.eessi.io/versions/2023.06/software/linux/x86_64/intel/skylake_avx512/.lmod/lmodrc.lua
Found Lmod SitePackage.lua file at /cvmfs/software.eessi.io/versions/2023.06/software/linux/x86_64/intel/skylake_avx512/.lmod/SitePackage.lua
Using /cvmfs/software.eessi.io/host_injections/2023.06/software/linux/x86_64/intel/skylake_avx512 as the site extension directory for installations.
Using /cvmfs/software.eessi.io/versions/2023.06/software/linux/x86_64/intel/skylake_avx512/modules/all as the directory to be added to MODULEPATH.
Using /cvmfs/software.eessi.io/host_injections/2023.06/software/linux/x86_64/intel/skylake_avx512/modules/all as the site extension directory to be added to MODULEPATH.
Found libcurl CAS file at RHEL location, setting CURL_CA_BUNDLE
Initializing Lmod...
Prepending /cvmfs/software.eessi.io/versions/2023.06/software/linux/x86_64/intel/skylake_avx512/modules/all to $MODULEPATH...
Prepending site path /cvmfs/software.eessi.io/host_injections/2023.06/software/linux/x86_64/intel/skylake_avx512/modules/all to $MODULEPATH...
Environment set up to use EESSI (2023.06), have fun!

Currently Loaded Modules:
 1) GCCcore/12.3.0
 2) Tcl/8.6.13-GCCcore-12.3.0
 3) SQLite/3.42.0-GCCcore-12.3.0
 4) libffi/3.4.4-GCCcore-12.3.0
 5) OpenSSL/1.1
```

The log window also shows a file browser view of the output directory, listing files such as `cachedProfilesData`, `classiq`, `Desktop`, `extensions`, `logs`, `Machine`, `NextFlowSlurm`, `nextUser`, `notebooks`, `ondemand`, and `ondemandOLD` with their respective file sizes.

## Sharing Software

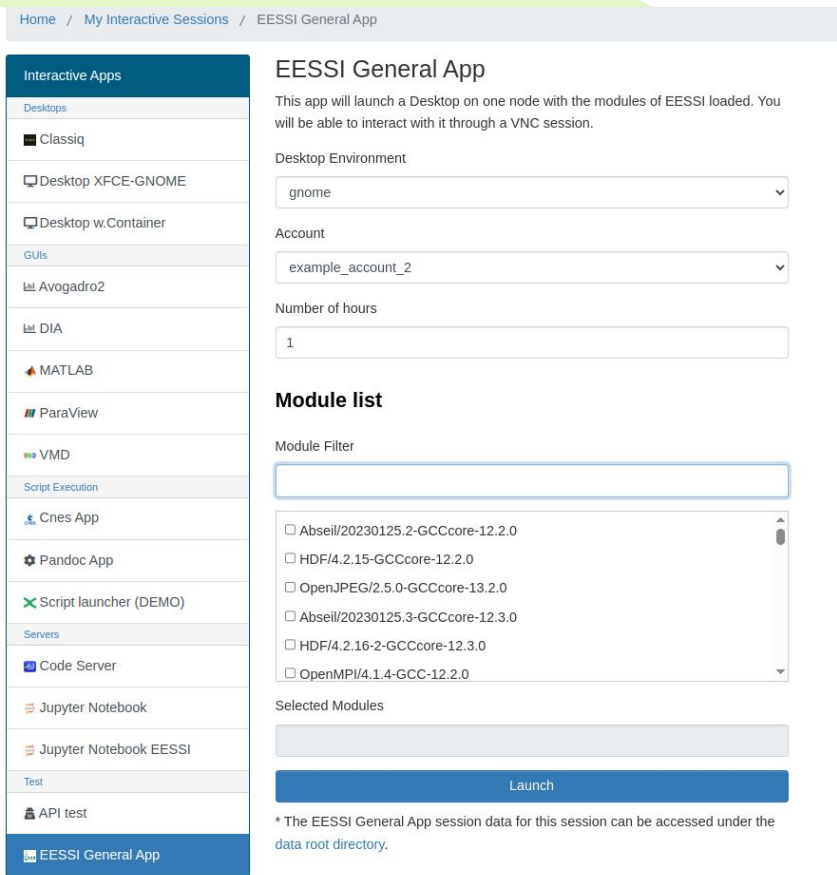
---

- How can we take advantage of **EESSI**'s ability to share software from **OPEN OnDemand** in the most user-friendly way?
- What can **OPEN OnDemand** do to help **EESSI** software sharing?



# Sharing Software

- VNC Desktop Based application
- Modules field to select software needed by the user
- Selected modules unwritable field to check the selected ones
- Search mechanism to find specific modules



Home / My Interactive Sessions / EESSI General App

### Interactive Apps

- Desktops
  - Classiq
  - Desktop XFCE-GNOME
  - Desktop w.Container
- GUIs
  - Avogadro2
  - DIA
  - MATLAB
  - ParaView
  - VMD
- Script Execution
  - Cnes App
  - Pandoc App
  - Script launcher (DEMO)
- Servers
  - Code Server
  - Jupyter Notebook
  - Jupyter Notebook EESSI
- Test
  - API test
- EESSI General App**

### EESSI General App

This app will launch a Desktop on one node with the modules of EESSI loaded. You will be able to interact with it through a VNC session.

Desktop Environment  
gnome

Account  
example\_account\_2

Number of hours  
1

### Module list


Module Filter

- Abseil/20230125.2-GCCcore-12.2.0
- HDF/4.2.15-GCCcore-12.2.0
- OpenJPEG/2.5.0-GCCcore-13.2.0
- Abseil/20230125.3-GCCcore-12.3.0
- HDF/4.2.16-2-GCCcore-12.3.0
- OpenMPI/4.1.4-GCC-12.2.0

Selected Modules

Launch

\* The EESSI General App session data for this session can be accessed under the data root directory.




## Sharing Software

- Adding software to the EESSI stack also grant access from all the other

**OPEN**  
**OnDemand** + **E E S S I**

integrations

- Accessible via this App simplifying the sharing between centers





**doitnow**

HPC Services



Christian Bustelo

HPC Engineer

[christian.bustelo@doitnowgroup.com](mailto:christian.bustelo@doitnowgroup.com)

Arturo

HPC Sysadmin

[arturo.gimeno@doitnowgroup.com](mailto:arturo.gimeno@doitnowgroup.com)