SITE REPORT : ELECTRICITÉ DE FRANCE

Slurm User Group

Cécile Yoshikawa
September 27, 2016

HPC AT EDF

- About EDF
- How do we do HPC at EDF?
- Our HPC infrastructures
- Our in-house OS dedicated to scientific IT needs

ABOUT EDF

- **World's biggest electric utility**
  - 75B € in annual revenue, 37.6M clients worldwide
  - 160,000 employees worldwide

- **Main activities**
  - Electricity Generation & Engineering
  - Electricity Transmission & Distribution
  - Research & Development
  - Optimization & Trading
  - Products & Services

- **Importance of R&D and engineering divisions**
  - 650M € Net R&D budget in 2015
  - 541 patented & protected innovations

HPC AND SCIENTIFIC IT NEEDS

- **Modeling**
  - Approximate reality with a model
- **Simulations on a wide range of fields:**
  - Execution of numerical codes
  - Structural and fluid mechanics, neutronics for nuclear plant maintenance
  - Materials for renewable energies
- **In-house developed codes (Often Open Source):**
  - Structures and Thermomechanics Analysis: Code_Aster (www.code-aster.org)
  - CFD: Code_Saturne (www.code-saturne.org)
  - Pre and post-processing with SALOME (http://www.salome-platform.org)
- **High Performance Visualization**
  - Parallel rendering

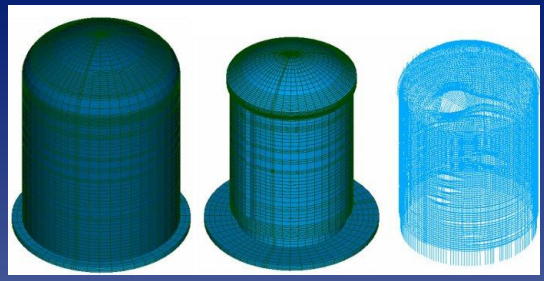

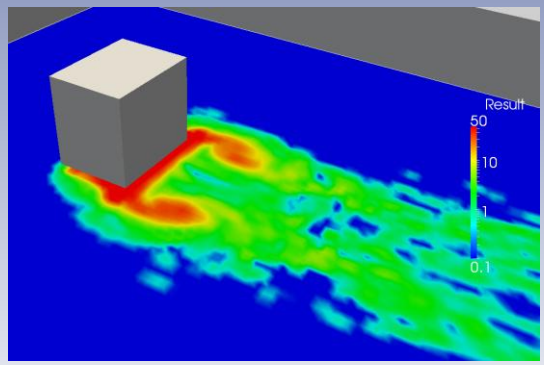

6 EXISTING CLUSTERS FOR 2 DIVISIONS

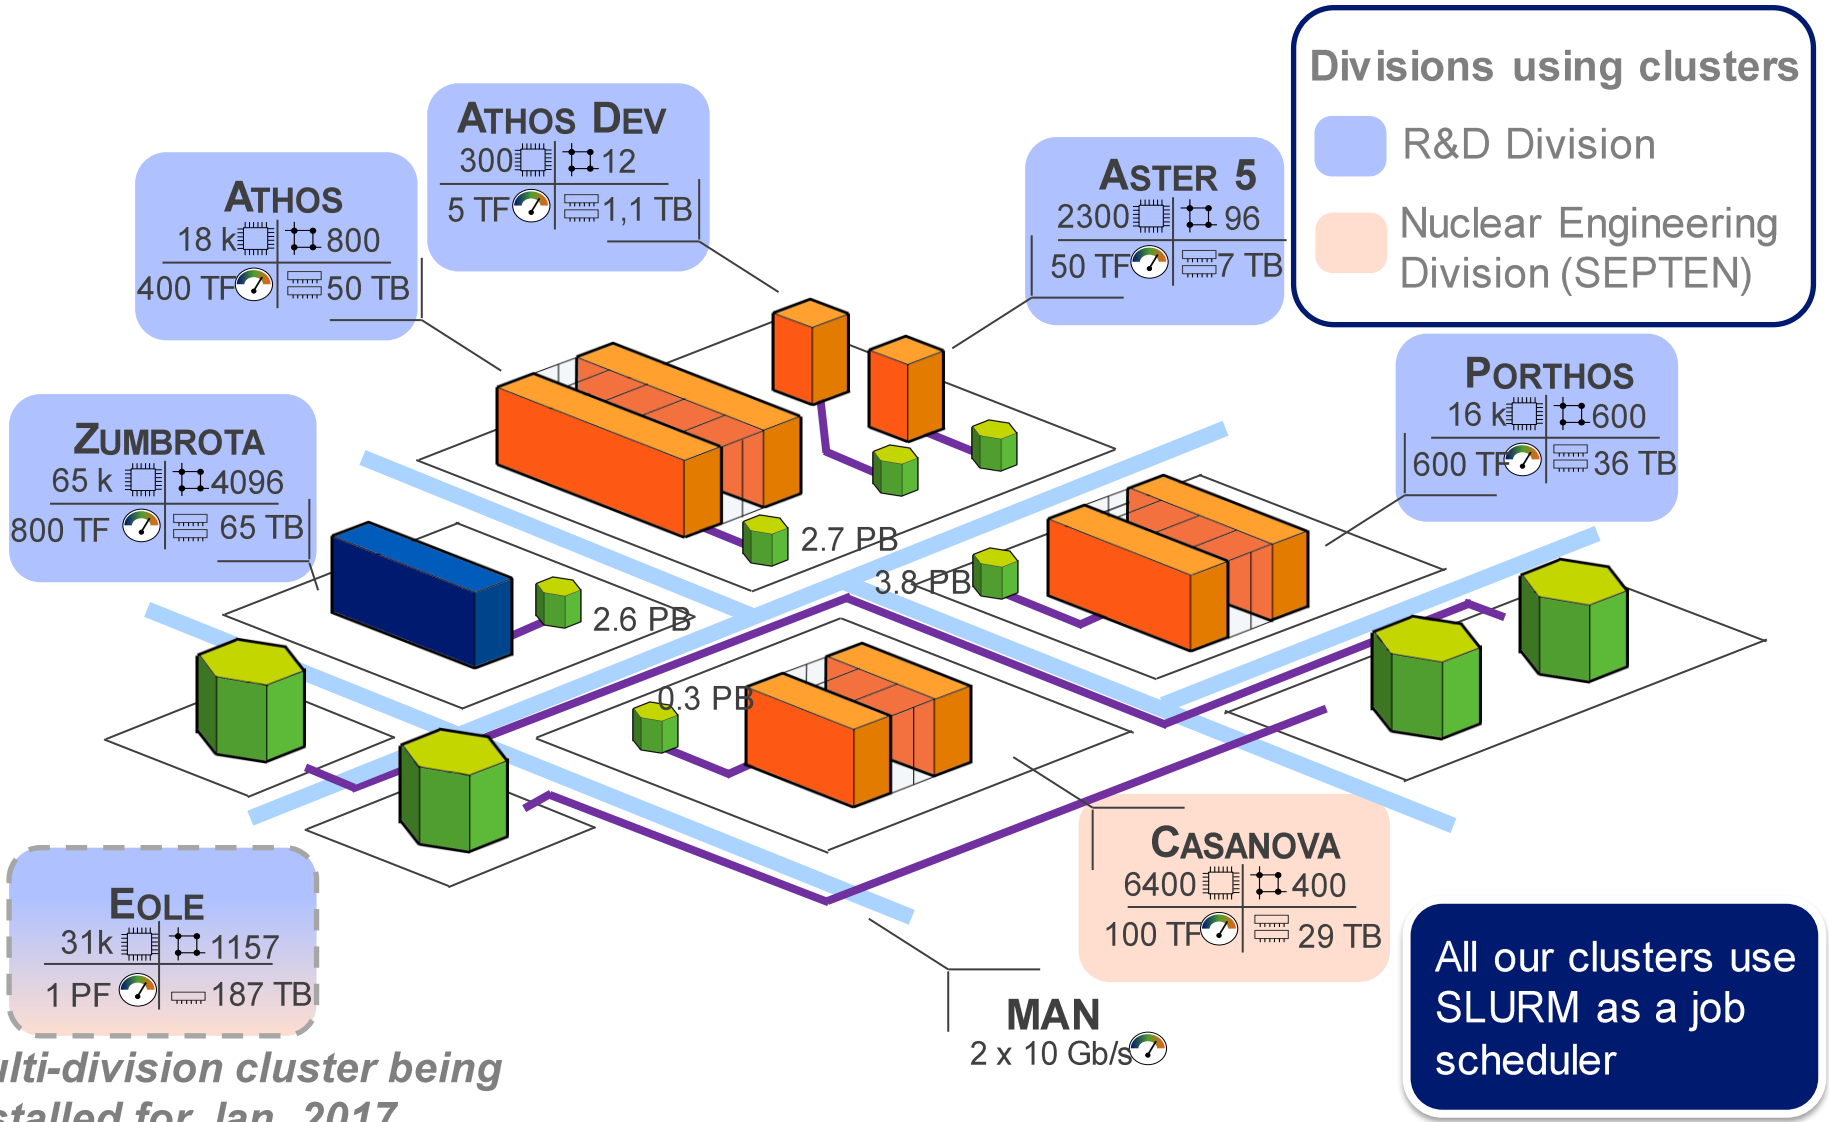

Multi-division cluster being installed for Jan. 2017

SCIBIAN

- **A Debian-based distribution: www.scibian.org**
  - Customizations to meet scientific IT needs
  - Initially an EDF custom distribution (Calibre)
    - Same distrib for workstations, servers & clusters
  - Being turned into an Open Source community project
    - Kick-off event on Sept. 30th at La Défense, Paris

- **Longer support for each major release:** beyond Oldstable

- **HPC with Scibian:**
  - Debian packaging of HPC dedicated SW:
    - GPFS, OFED, Mellanox IB stack, OPA to come
  - Custom Deployment System for diskless nodes
  - Tools on top of SLURM:
    - SLURM Dashboard, JobMetrics, NEOS

OUR SLURM USAGE AT EDF

- The functionalities we use
- A new challenge with our new cluster

DETAILS ON SLURM USAGES (1/2)

- **SLURM 15.08.8 on all clusters**

- **Separate partitions depending on node types**
  - Standard nodes
  - Large memory nodes
  - Graphical nodes with a GPU card

- **Several QOS on each cluster**
  - Selected partition
  - Required number of cores
  - Walltime

- **LUA Plugin for job submission**
  - Automatically route jobs into the proper QOS
  - https://github.com/edf-hpc/slurm-llnl-misc-plugins/tree/master/job_submit

DETAILS ON SLURM USAGES (2/2)

- **Accounting used on each cluster**
  - One dedicated database per cluster
    - Easy to maintain & to decommission
    - MariaDB in mode multi-master
  - One additional global PostgreSQL database collecting data from the per cluster databases, log files, LDAP information

- **Scheduling Policy**
  - Multi-factor Job Priority
  - Classic Fairshare Algorithm for existing clusters
  - Fair Tree Fairshare Algorithm for our new cluster

- **CPU and Memory as consumable resources**

- **Task Plugin: cgroup on the most recent clusters**
  - Memory controller (ConstrainRAMSpace=yes) to be used in our new cluster

A NEW CLUSTER SHARED BY 2 DIVISIONS (1/2)

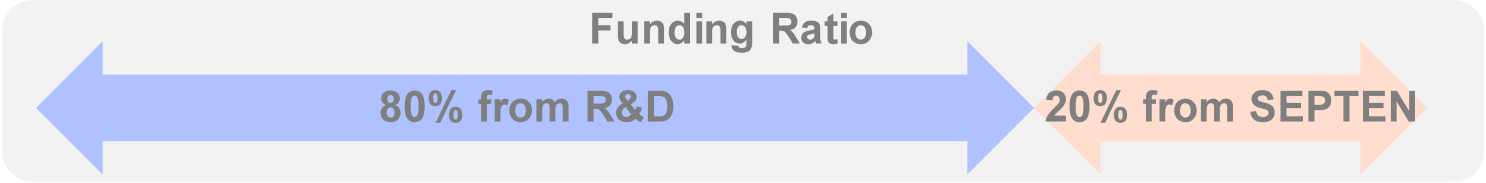

- **New constraints to share the resources**
  - Ideal resource ratio: 80% for R&D, 20% for SEPTEN
    - But no static sharing
  - If some resources of a division are unused, the other division should be able to use them
  - A division should be able to use all the resources it is entitled to within 8h

A NEW CLUSTER SHARED BY 2 DIVISIONS (2/2)

- **Solution to be implemented**
  - Fair Tree Fairshare Algorithm
    - PriorityFlag = FAIR_TREE
  - An account per division with fairshare factors according to the sharing ratio

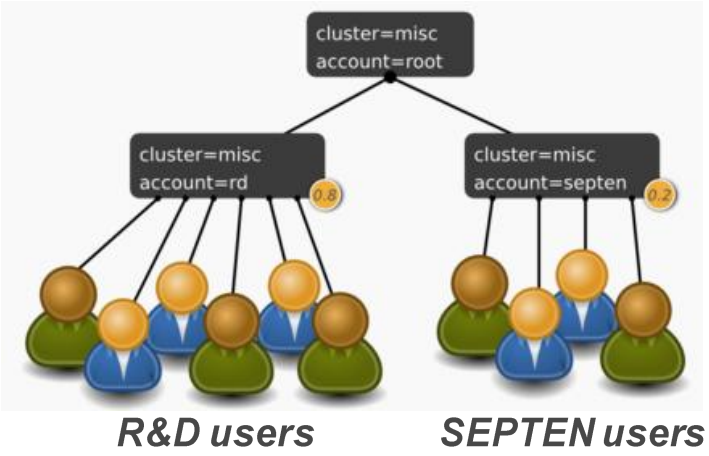

  - Jobs are classified in 2 types :
    - Short jobs < 8h
    - Long jobs between 8h and 7 days
  - 1 QoS for short jobs shared between the 2 divisions
  - 2 QoS for long jobs, 1 for each division with a GrpTRES limit on the number of nodes
  - Higher priority for the short job QoS

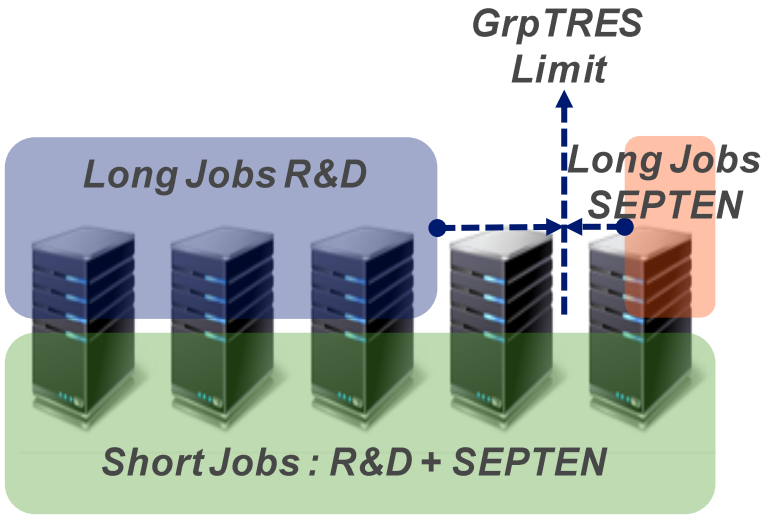

OUR IN-HOUSE DEVELOPED TOOLS TO WORK WITH SLURM

- SlurmWeb
- JobMetrics

SLURMWEB (1/4)

- A SLURM Dashboard for real time monitoring
- Sources: https://github.com/edf-hpc/slurm-web
- Documentation: https://edf-hpc.github.io/slurm-web

- Information about jobs

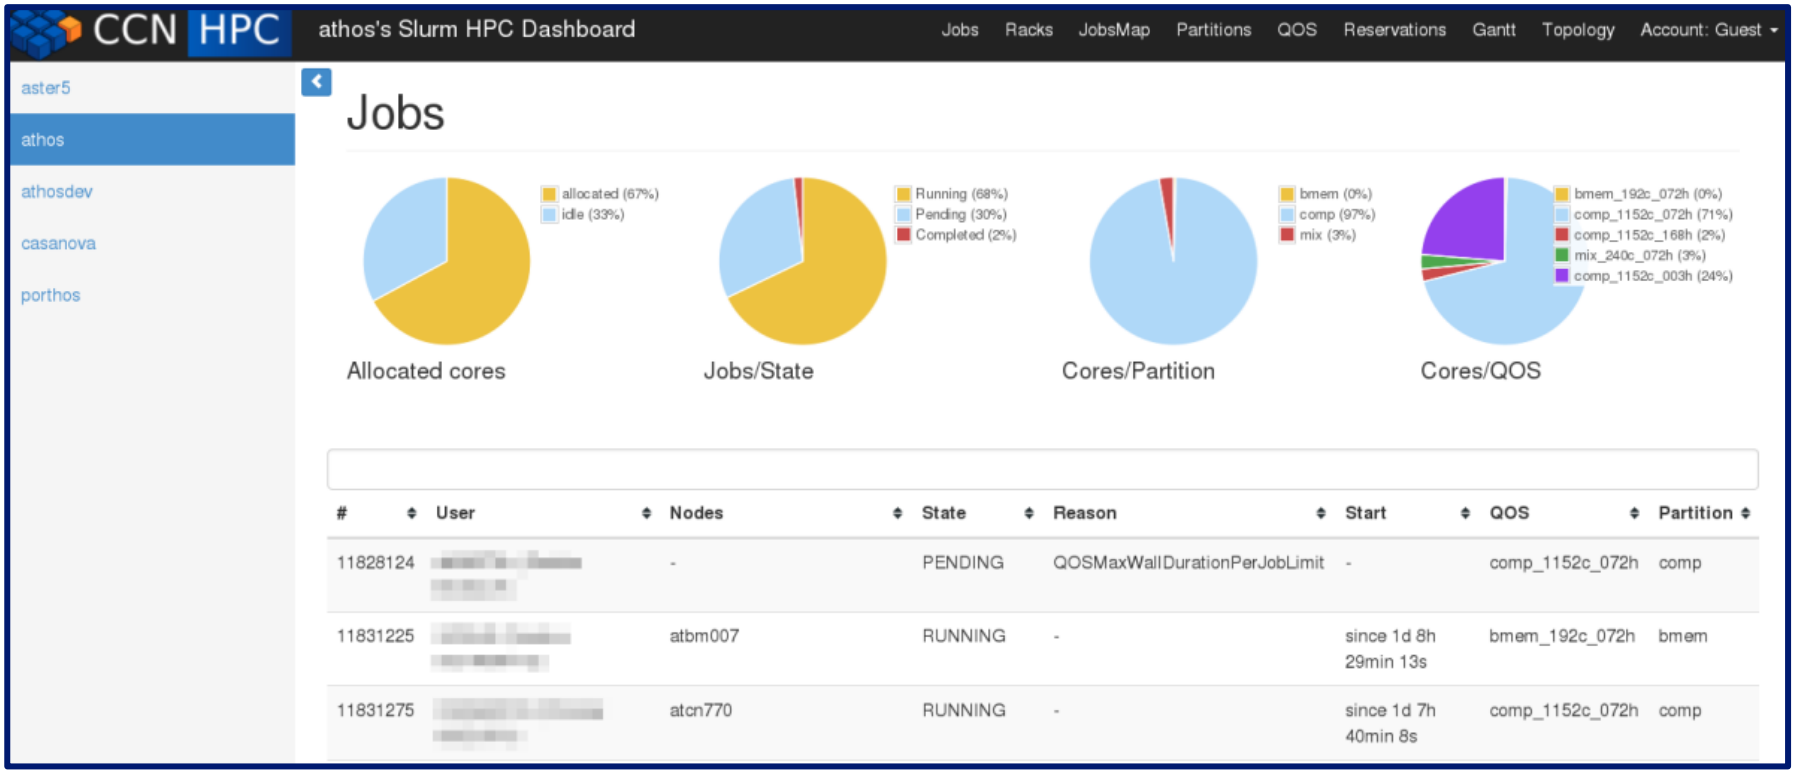

SLURMWEB (2/4)

- **Information about racks and nodes**

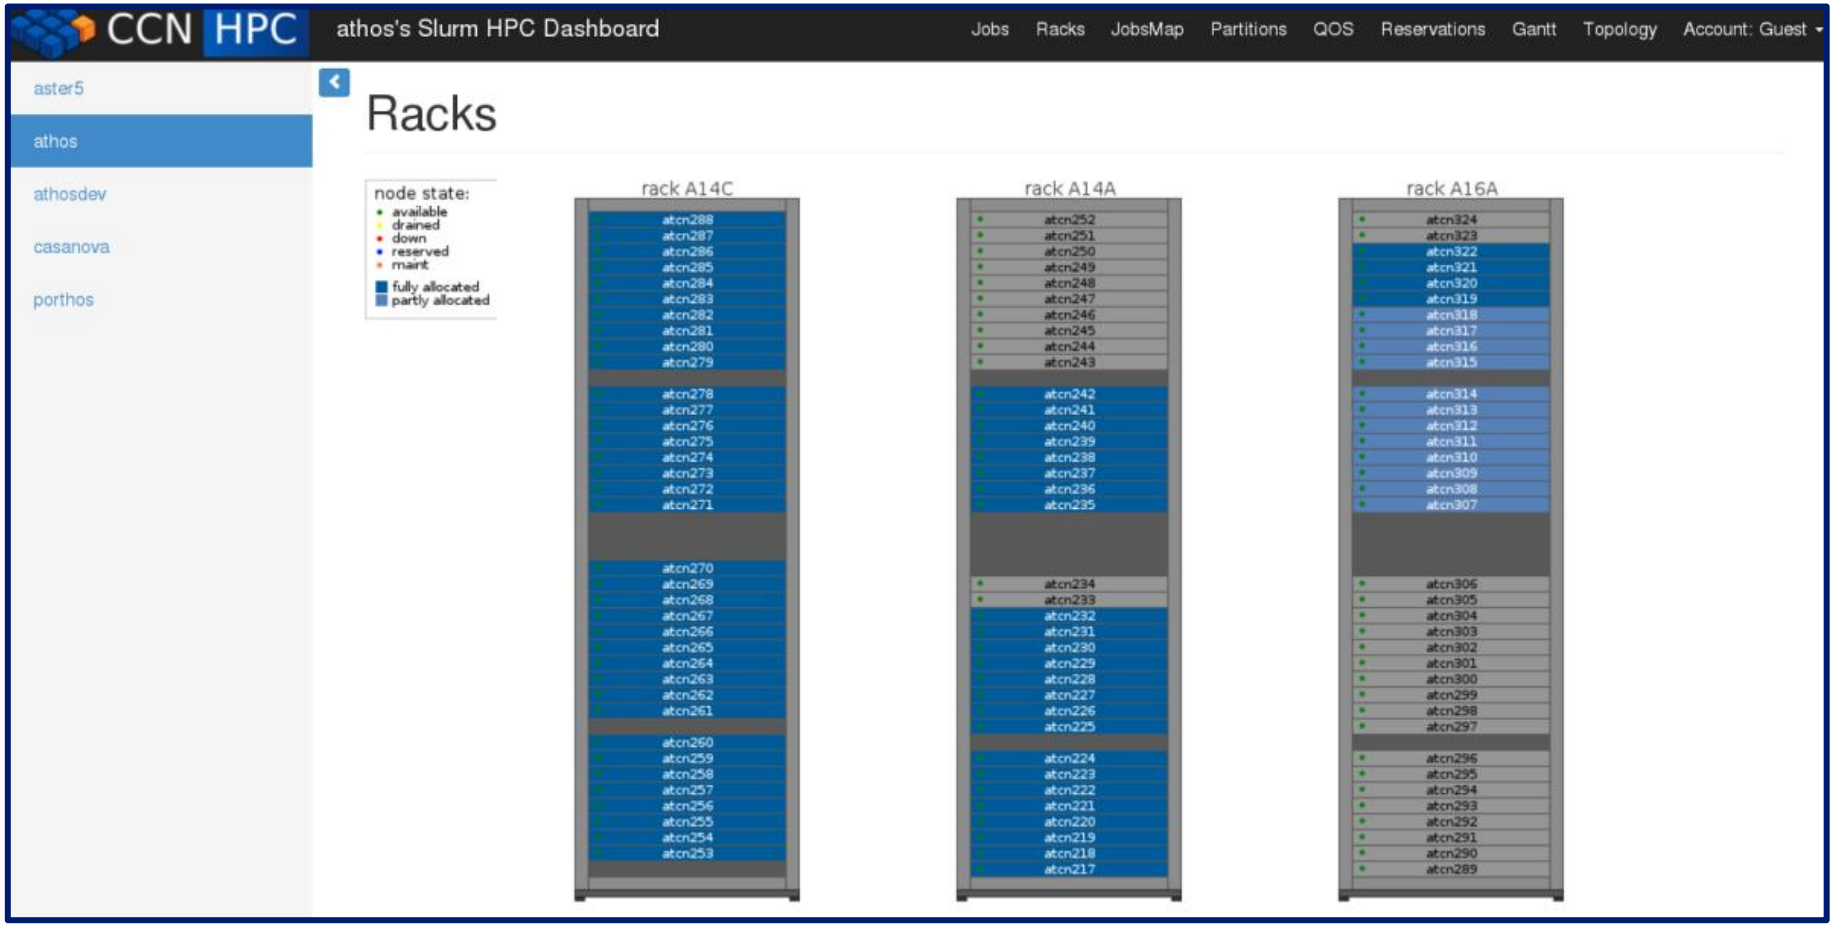

SLURMWEB (3/4)

- **Mapping between nodes and jobs**

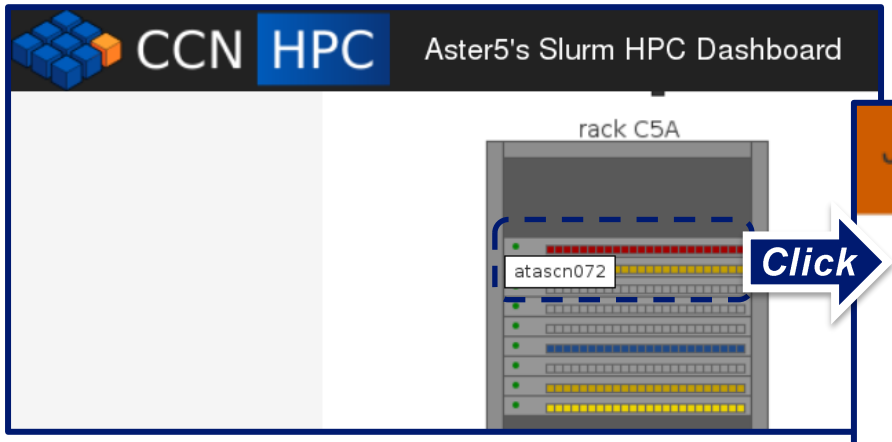

Rack View with nodes running a job

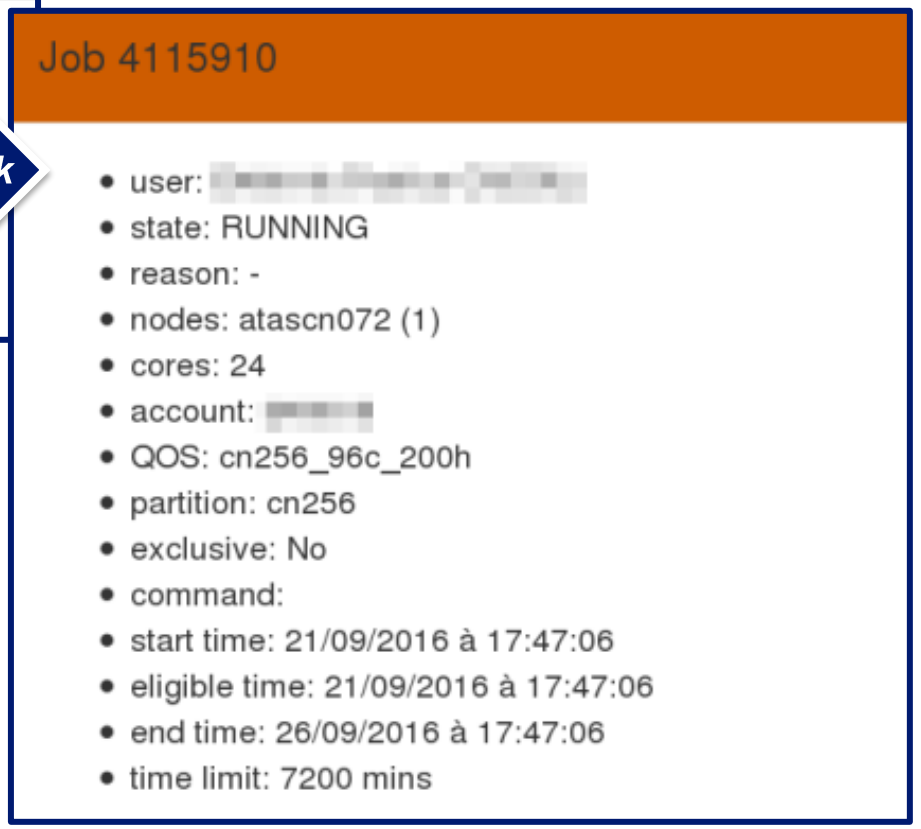

Information about the job running on the selected node

SLURMWEB (4/4)

- **Software Architecture**

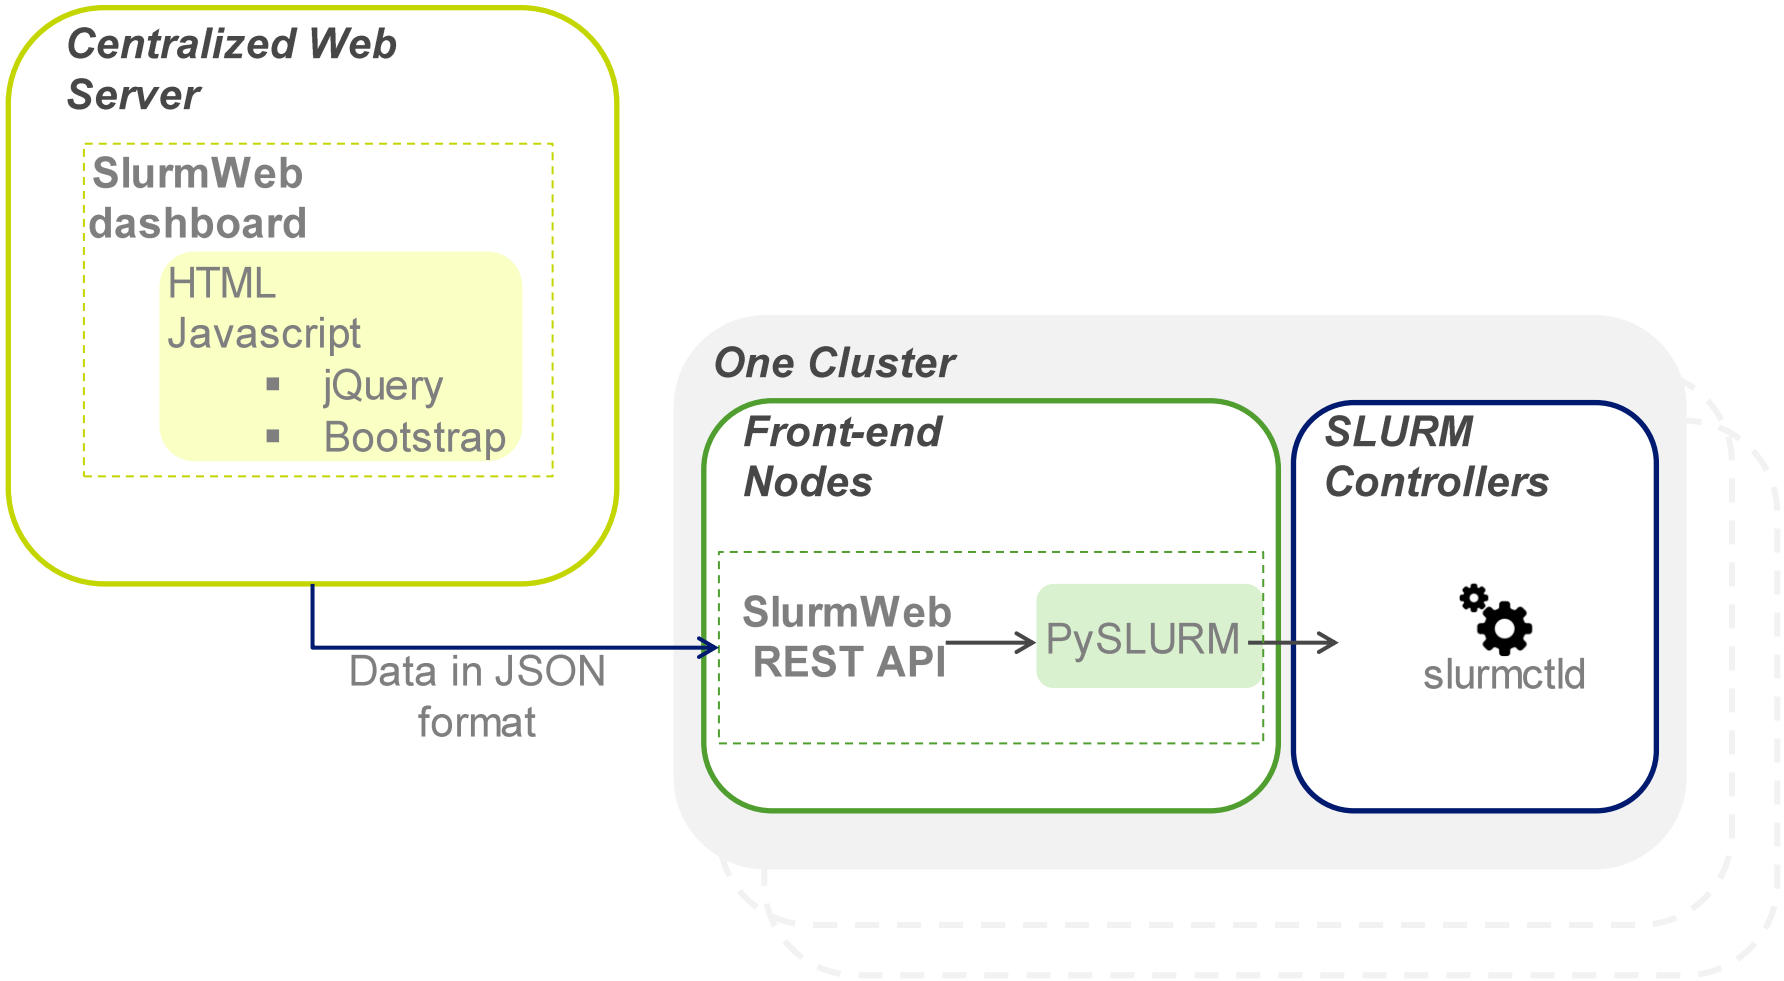

JOBMETRICS (1/2)

- user: [redacted]
- state: RUNNING
- reason: -
- nodes: pocn[250,280,290-303] (16)
- cores: 448
- account: rdusers
- QOS: cn_0448c_024h
- partition: cn

HPC metrics: cluster porthos job 389734

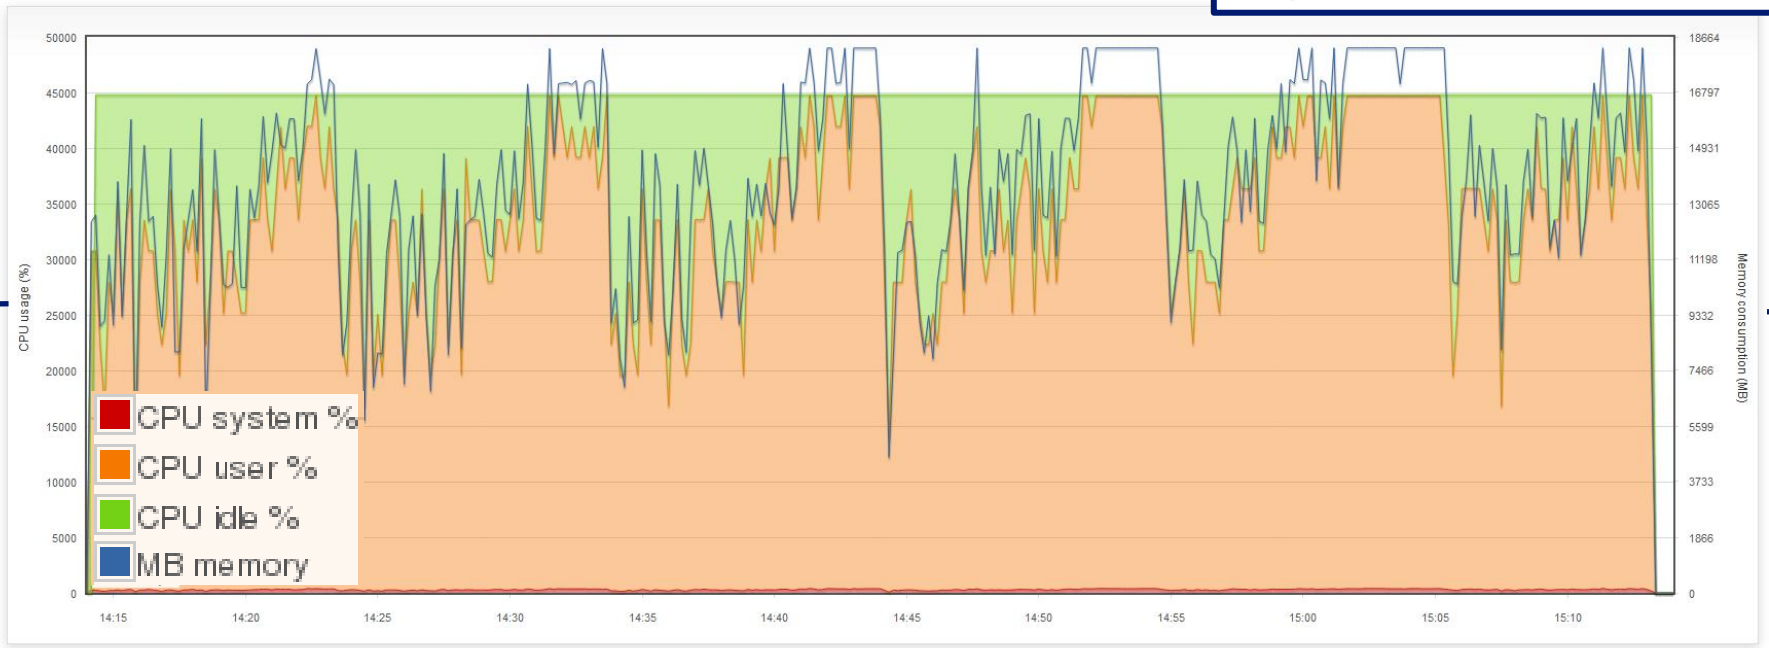

Real-time Total CPU Consumption of the job

100% : full consumption of one core

Real-time Total Memory Consumption of the job

Blue Graph

JOBMETRICS (2/2)

- **Web application to supply and display HPC job metrics such as**
  - Real-time CPU consumption for a job during its execution
  - Real-time memory consumption for a job during its execution
- **Prerequisites**
  - Use of *cgroups* task plugin to distinguish resource consumption in case of several jobs running simultaneously on one node
- **Implementation**
  - *Collectd* running on each computation node to collect metrics

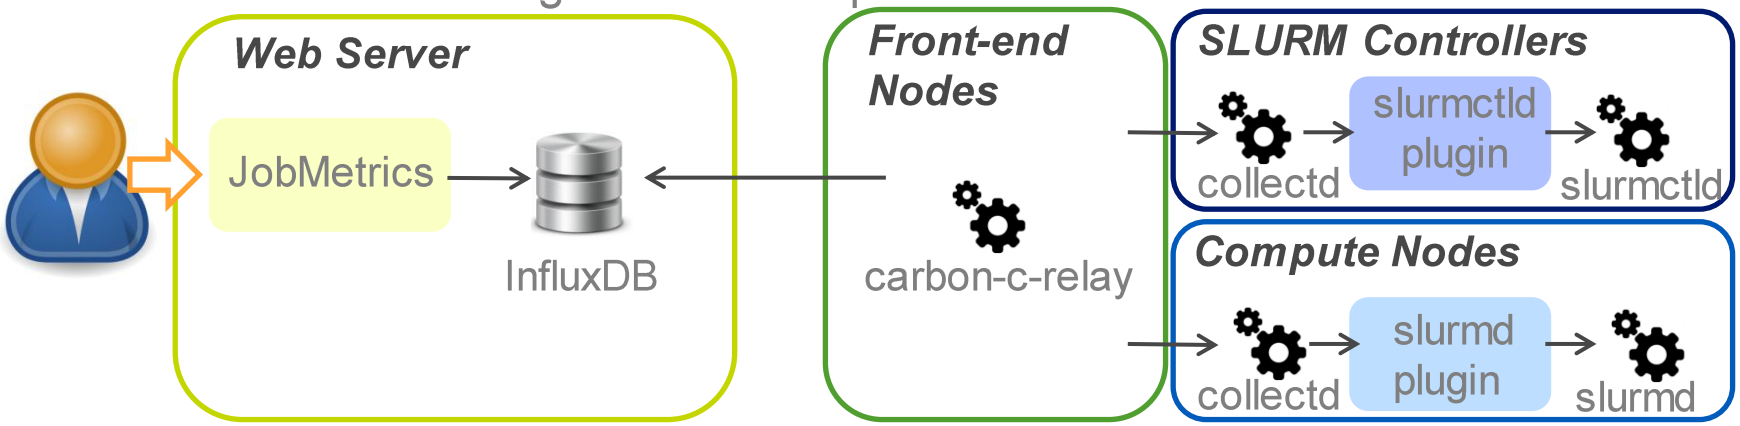

- **Sources : https://github.com/edf-hpc/jobmetrics**

WHAT IS NEXT?

- SLURM jobs in containers

SLURM JOBS IN CONTAINERS

- **Initial Problem**
  - Natural OS life cycle
  - Some end-users want to use only qualified tools
    - Qualification sometimes takes a while
  - Developers want to test the newest tools available

- **Goal**
  - Allow more flexibility at the end of life of one OS version
    - Be able to run jobs on an old OS version
  - Allow early code porting
    - Be able to run jobs on the upcoming OS version

=> Run jobs on several Scibian versions dynamically

- **Constraints**
  - Easy selection of the OS version
  - Serial and MPI jobs OK
  - No loss of performance

SLURM JOBS IN CONTAINERS

- **Technical approach**
  - **Usage**
    - Choice of OS version with an environment variable set up with *lmod*
  - **Containers**
    - systemd-nspawn to be launched by slurmd (*ChosLoc* parameter)

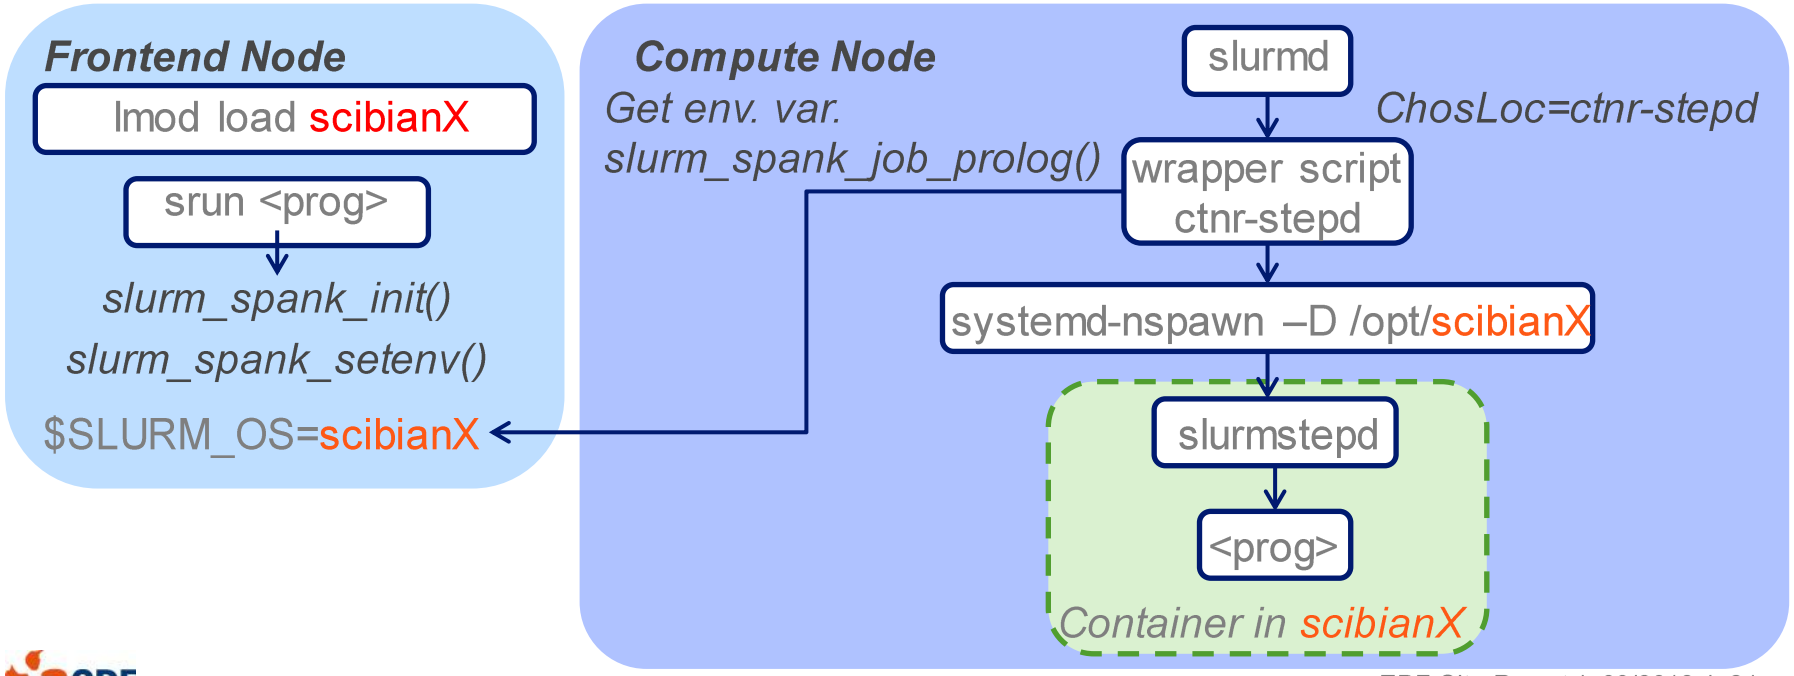

THANK YOU FOR LISTENING.
ANY QUESTIONS?

- All our tools are on Github:
  https://github.com/edf-hpc/
- Feel free to contact us:
  dsp-cspito-ccn-hpc@edf.fr
  cecile.yoshikawa@edf.fr