A fully configurable HPC web portal for managing Slurm jobs

Patrice Calegari

Slurm User Group SLUG'19
Salt Lake City, USA - September 18, 2019

Atos

We will talk about...

Context of the projects

XCS - eXtreme factory Computing Studio

BEM - Bull Efficiency Manager

Conclusion and future work

Context of the projects

Bull/Atos HPC & AI Software R&D

- Our division, Atos BDS (Big Data & Security) is in charge of developing supercomputing hardware and middleware.
- Our domains of interests: HPC, AI and Quantum simulations.
- User experience (UX) is extremely important
- Security is critical in all our activities (and those of our clients)
- We contribute to Slurm community and integrate Slurm in our HPC stack for more than 10 years

XCS
eXtreme factory Computing Studio

Extreme factory Computing Studio v3 (XCS3)

Introduction

- **Modular HPC, AI & Quantum portal**
  - as-a-Service cornerstone application,
  - supports **Slurm** (and other schedulers)
  - Role Based Access Control (RBAC)
  - supports AD, LDAP (with Kerberos)
  - XCS = REST API service + GUI
- **Fully customizable user interface**
  - Responsive Web Design (RWD) GUI
  - Single Page Application (SPA) with configurable dashboards: layout, components, languages, themes

Latest release: XCS 3.8.0 (April 5, 2019)

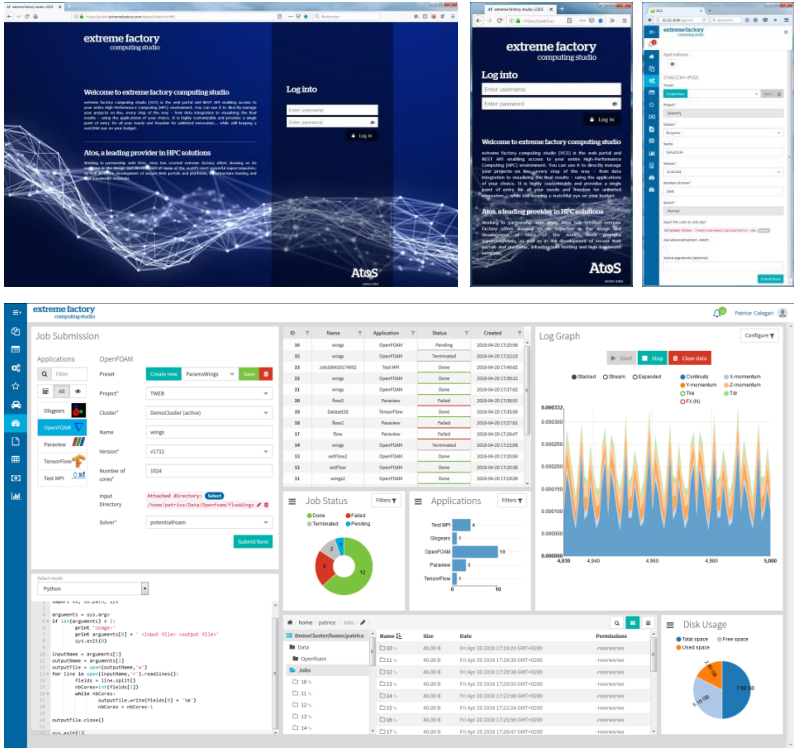

XCS REST API

https://public.extremefactory.com/demo/api/doc/api-full.html

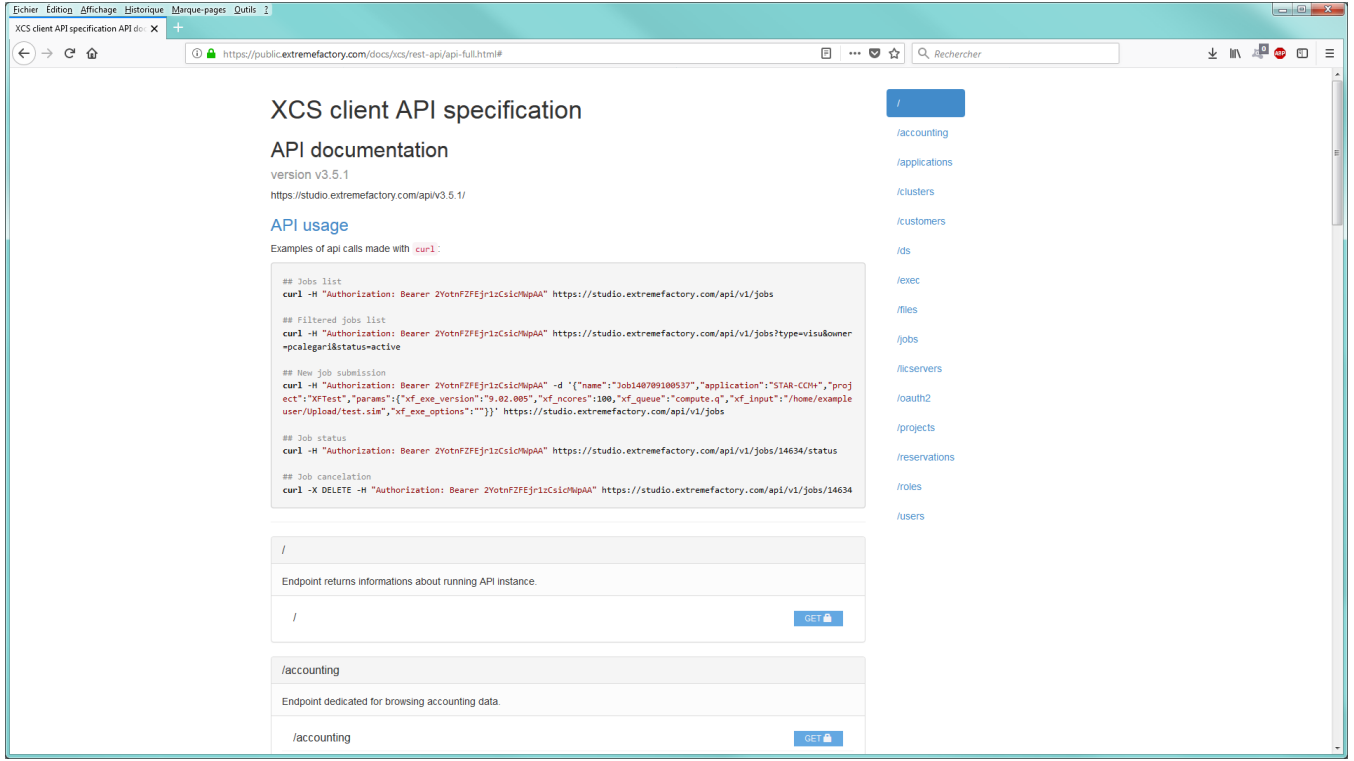

XCS REST API

https://public.extremefactory.com/demo/app/api/doc/api-full.html

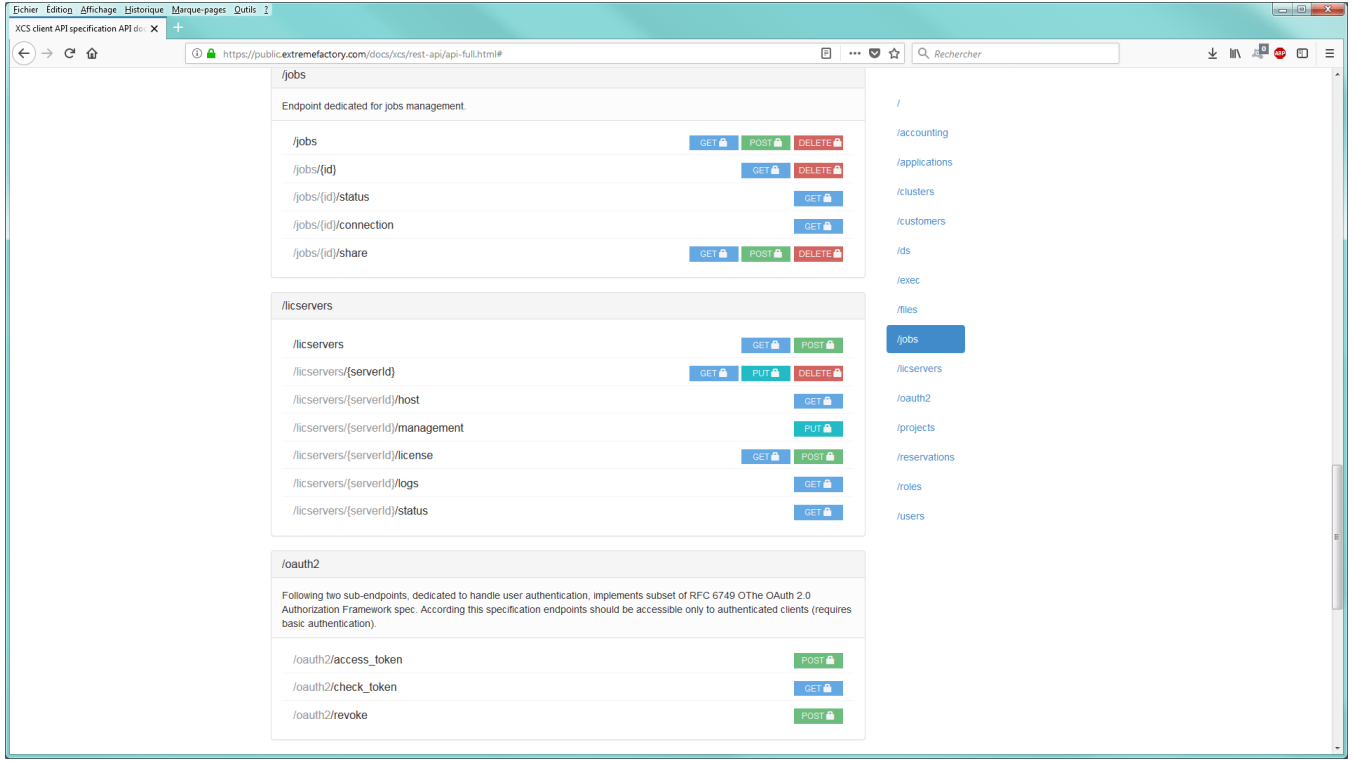

XCS REST API

https://public.extremefactory.com/demo/app/api/doc/api-full.html

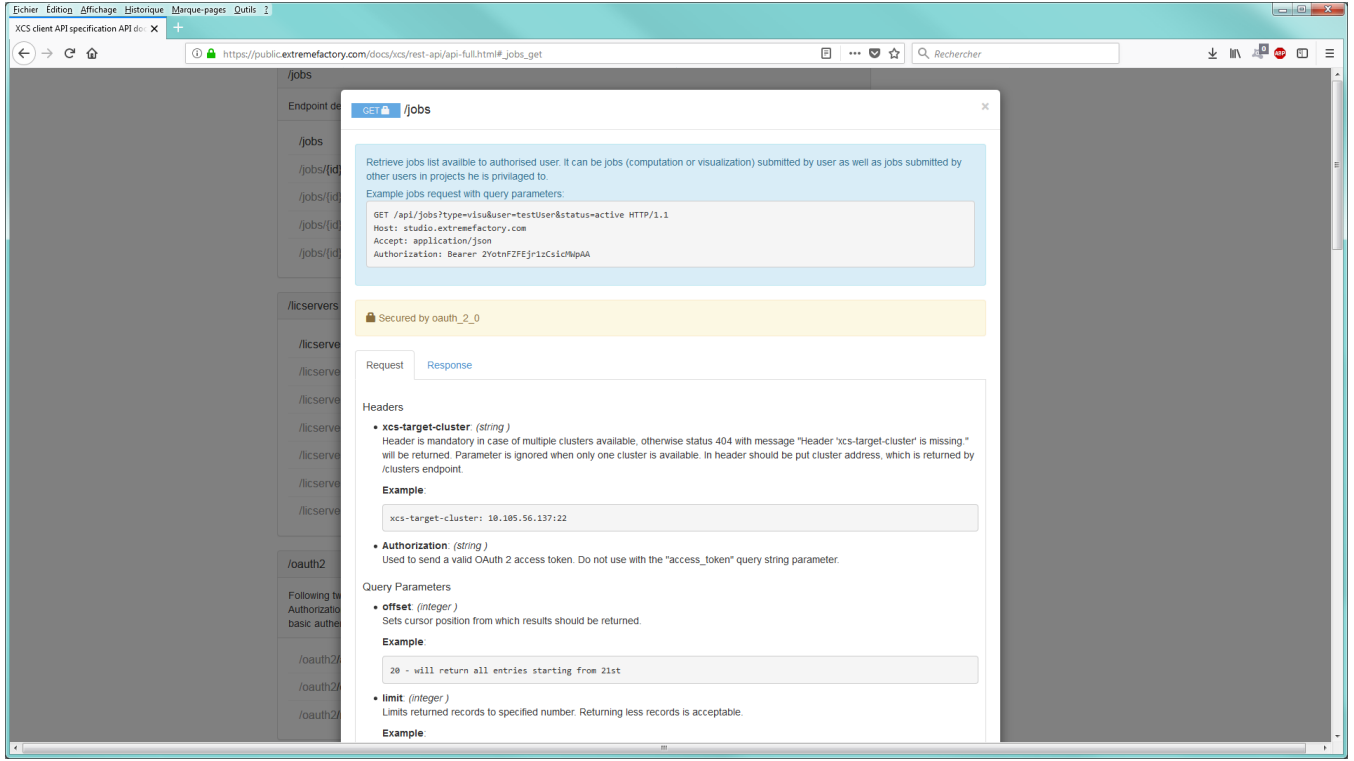

XCS user dashboard

Example 1: 8 components

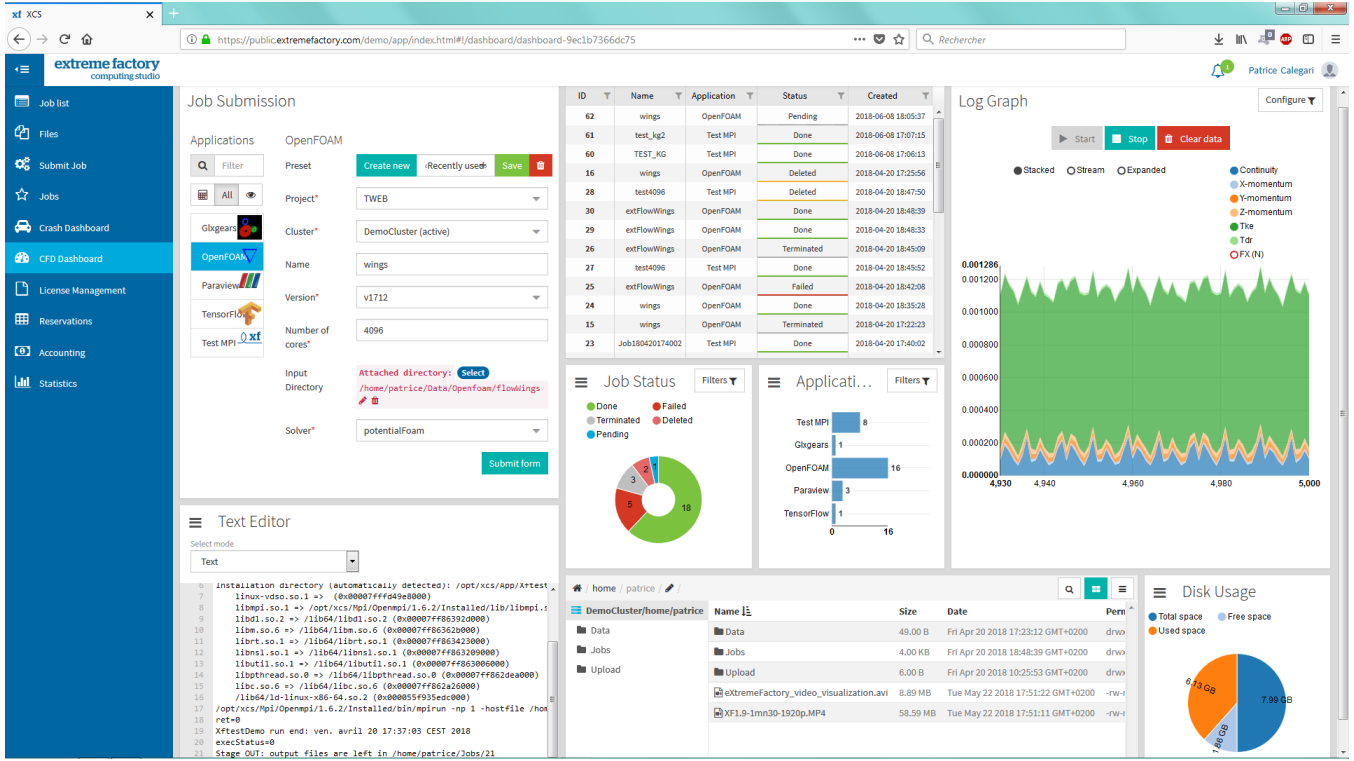

XCS user dashboard

Example 2: 1 component

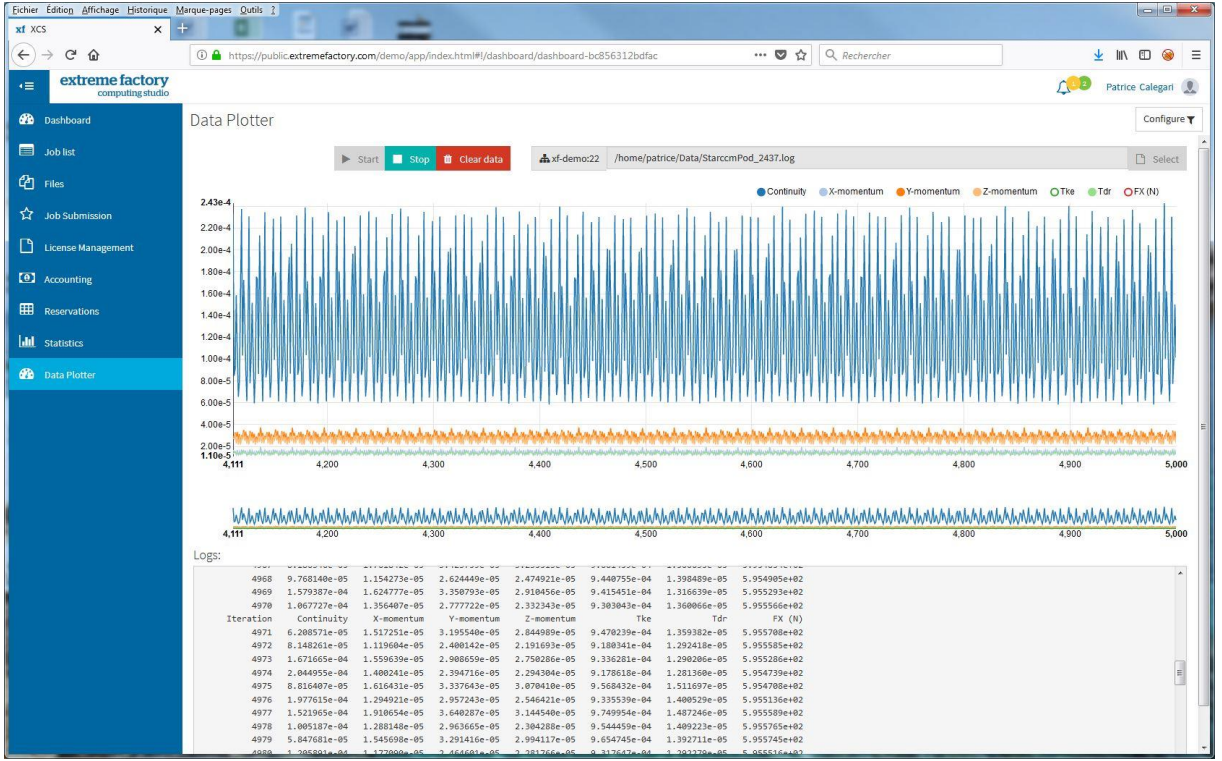

XCS user dashboard

Example 3: 6 components with edited theme

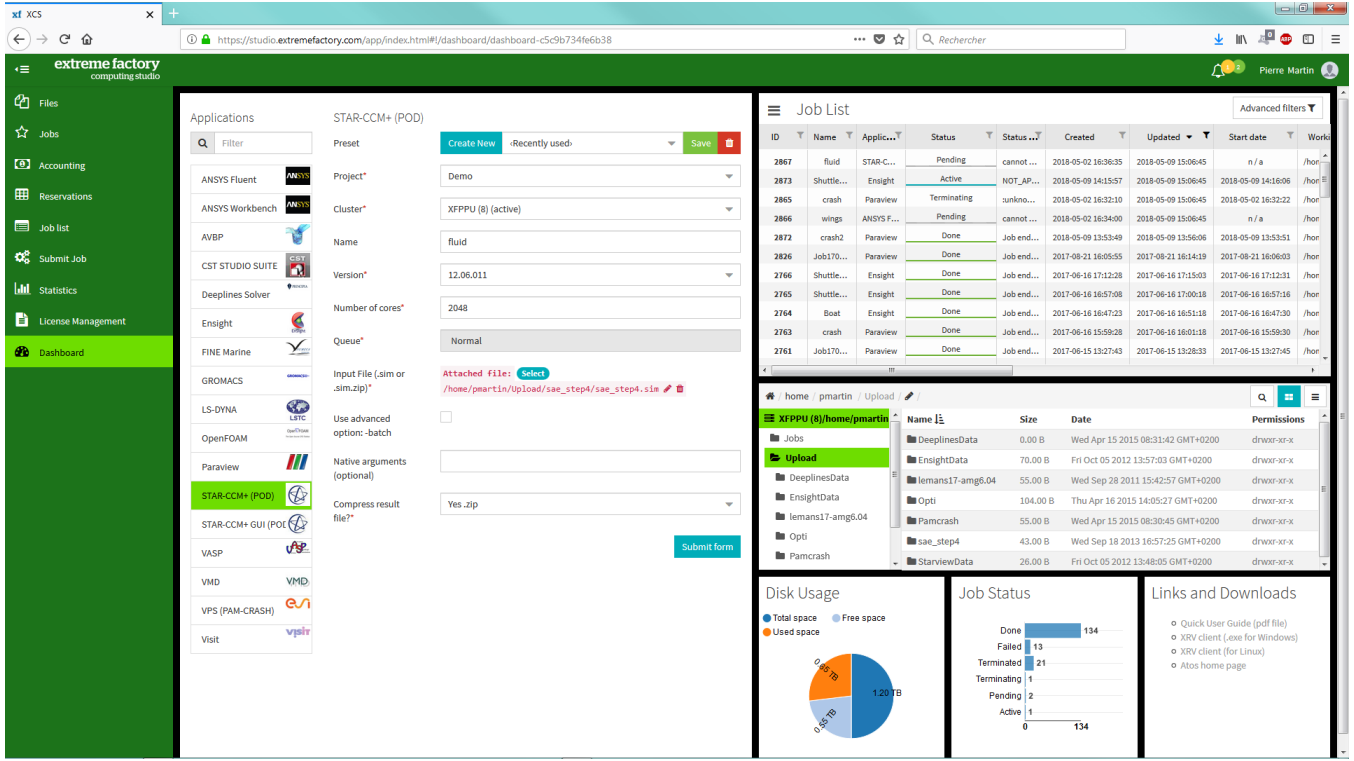

XCS dashboard main menu

import/export dashboards

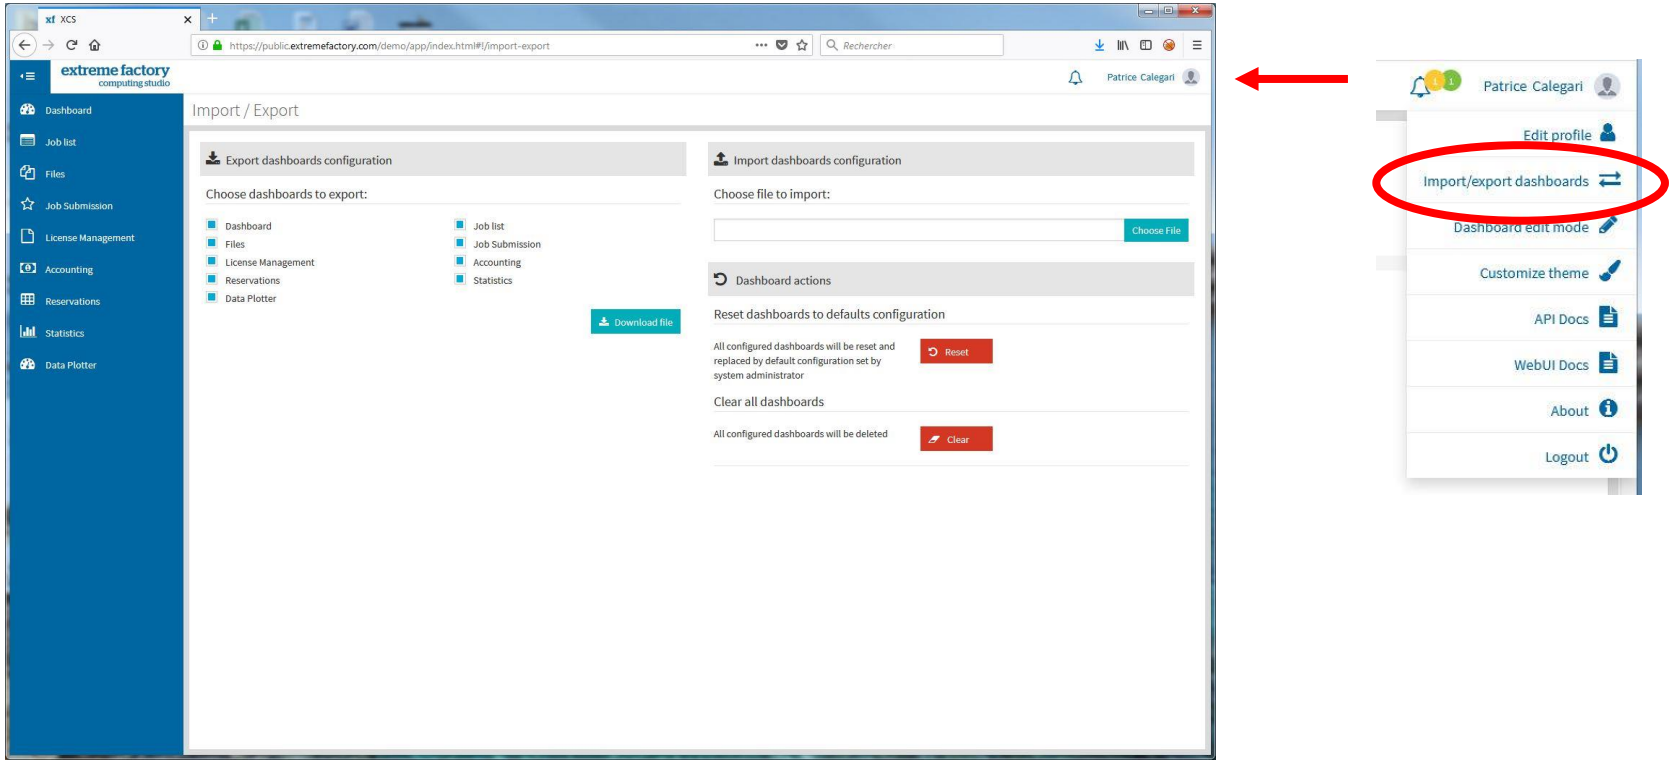

XCS dashboard main menu

REST API documentation

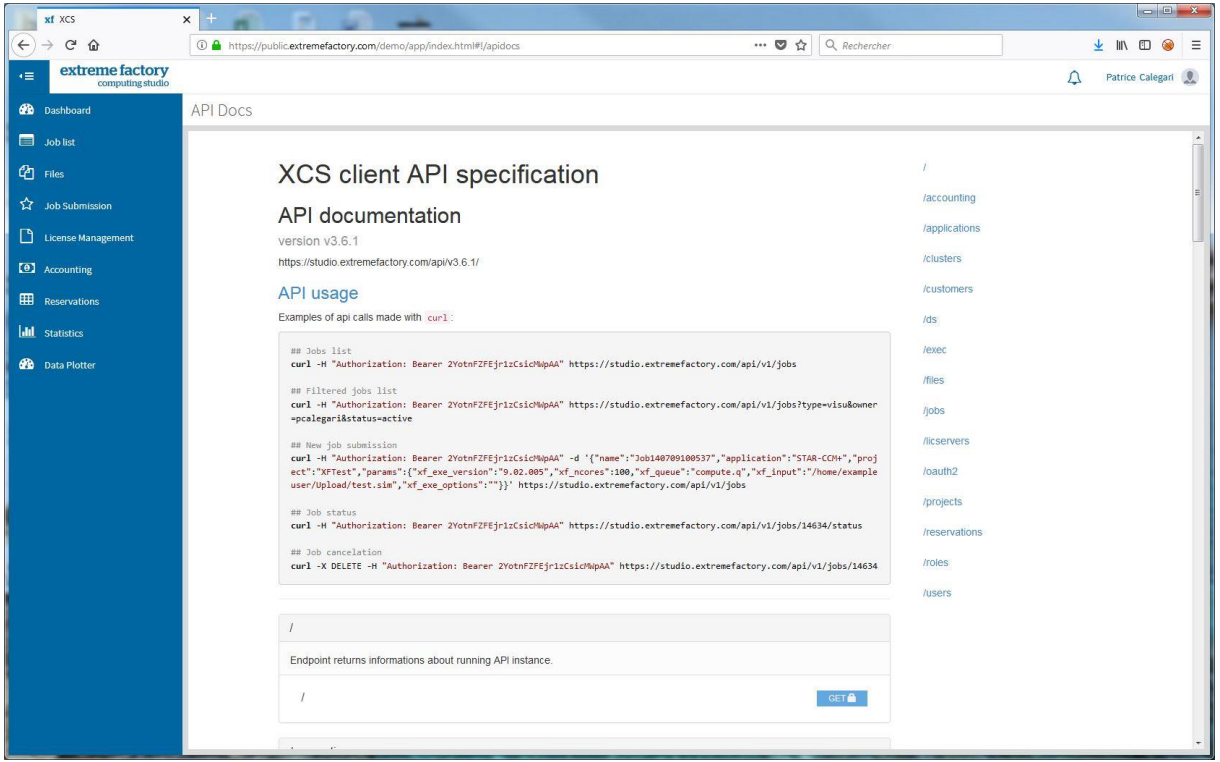

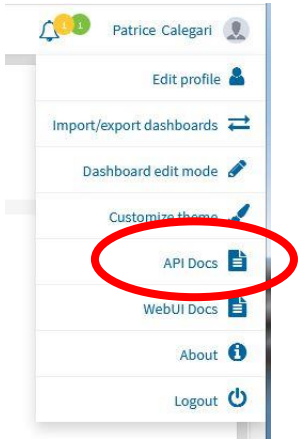

XCS Fundamental concepts

Key software product for HPCaaS solutions

Give users and admins access to resources through web services

- Use of a GUI in a web browser that relies on a REST API

Be compatible with « all possible » environments

- Software, frameworks, middleware

Never be intrusive

- The solution should be used in existing environments without modifying them

Keep all the intelligence in the REST API server

- The goal of the GUI is only to be the HMI (Human Machine Interface)

XCS architecture

current v3

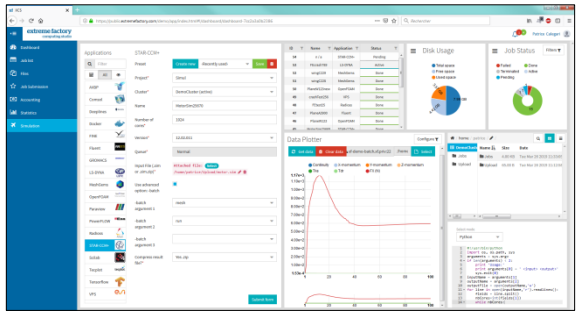

XCS web User Interface

HTTPS

DC API

Job submission DC

XCS DCs

Data mngmt DC

XCS GUI web server
- Dashboards
- Web Design

HTTPS

XCS REST API web server

Security service

SSH

HPC cluster integration layer
- Slurm
- HPC applications

XCS Data base

Directory service

DC = Dashboard Component

Slurm job submission workflow with XCS

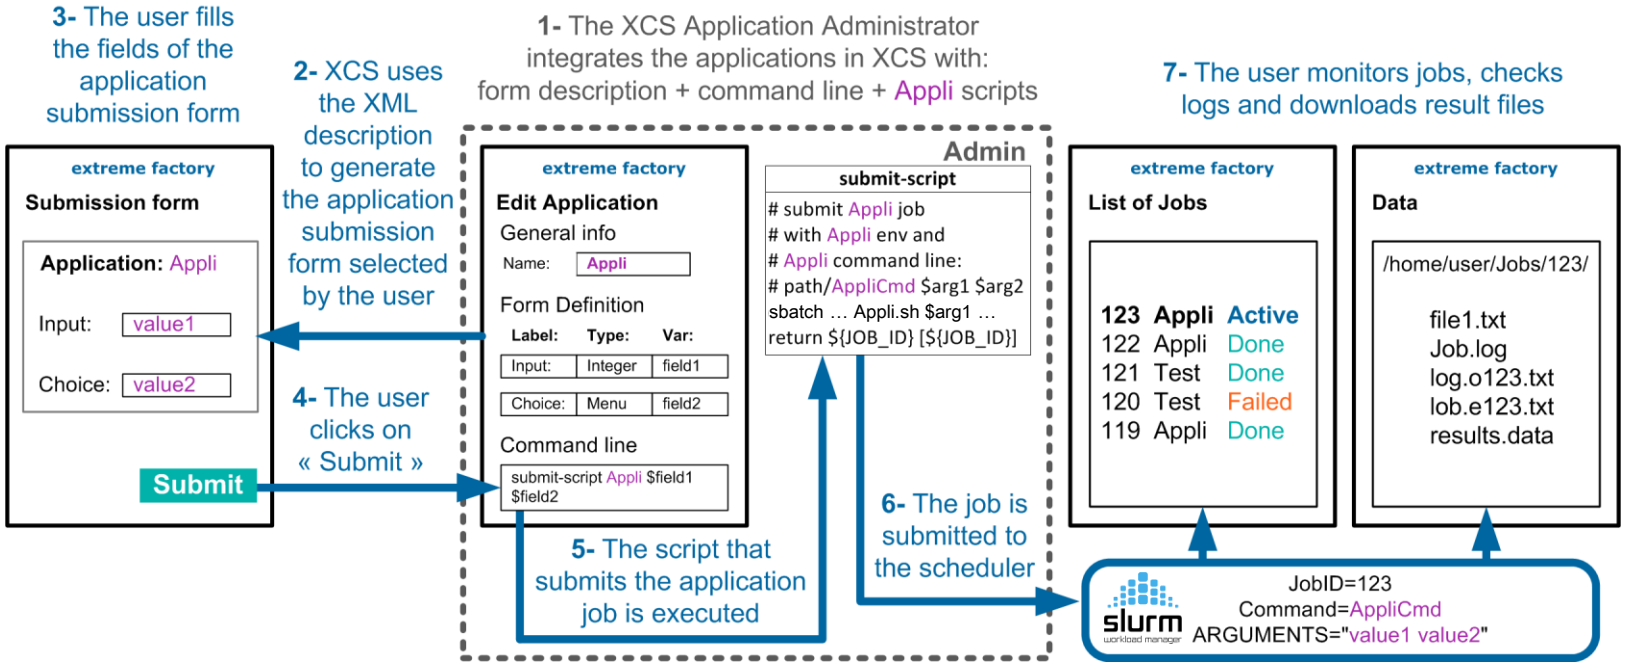

XCS application administrator dashboard

HPC application general information

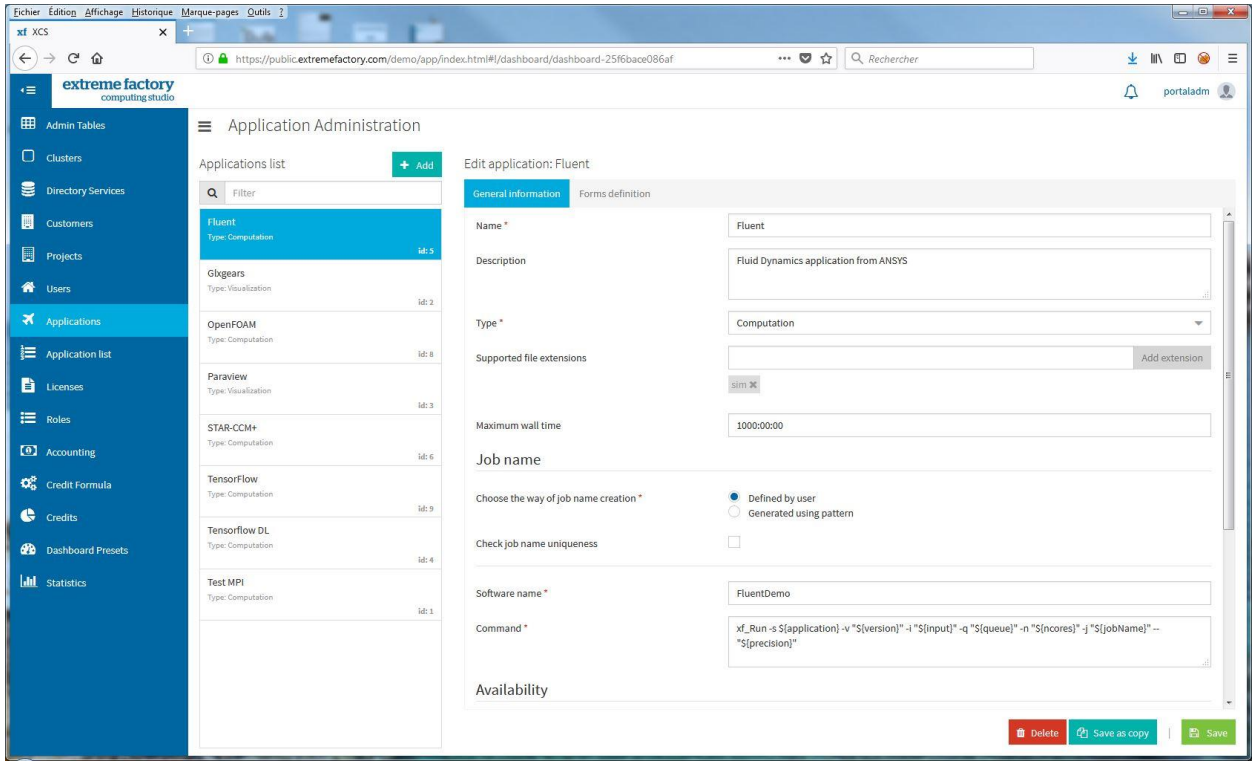

XCS application administrator dashboard

HPC application form definition

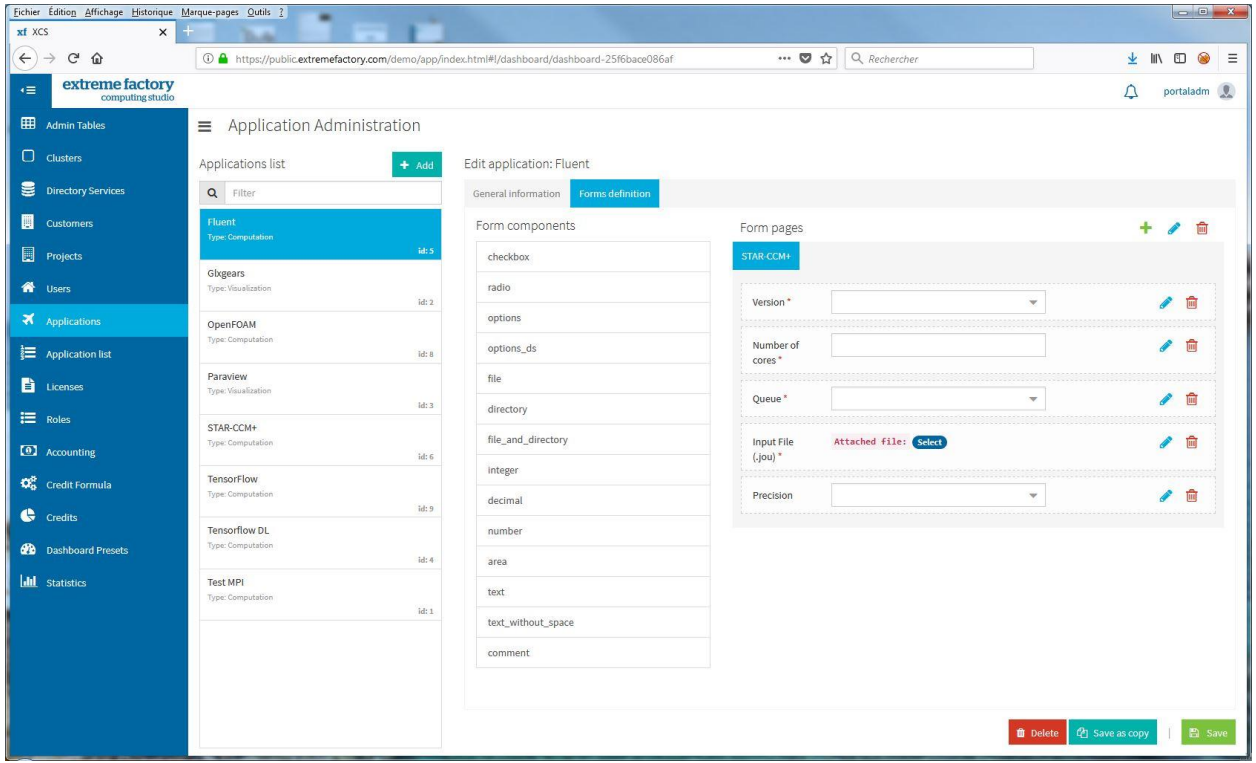

BEM
Bull Efficiency Manager

Bull Efficiency Manager (BEM)

Introduction

- **Slurm** has been enhanced by Bull/Atos to provide additional functionality including topology-aware resource allocation and advanced placement policies,
- **Bull Efficiency Manager (BEM)** is the web application running upon the **Slurm** workload manager to show cluster details interactively,
- **BEM** dashboards show information in graphs and tables for both current and previous archived data about cluster resources.

XCS architecture

current v3

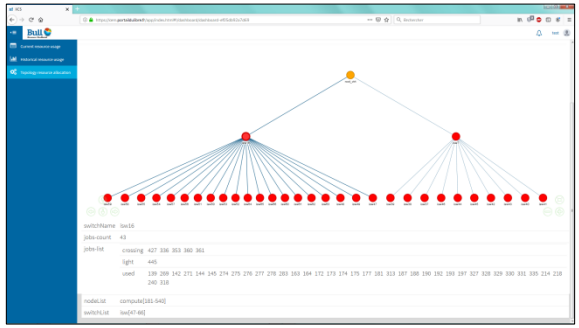

BEM web User Interface

HTTPS

DC API

Switch Topology DC

BEM DCs

Slurm usage history DC

XCS GUI web server

- Dashboards
- Web Design

HTTPS

BEM REST API web server

Security service

SSH

HPC cluster integration layer

- Slurm

BEM Data base

Directory service

DC = Dashboard Component

BEM

Login Page

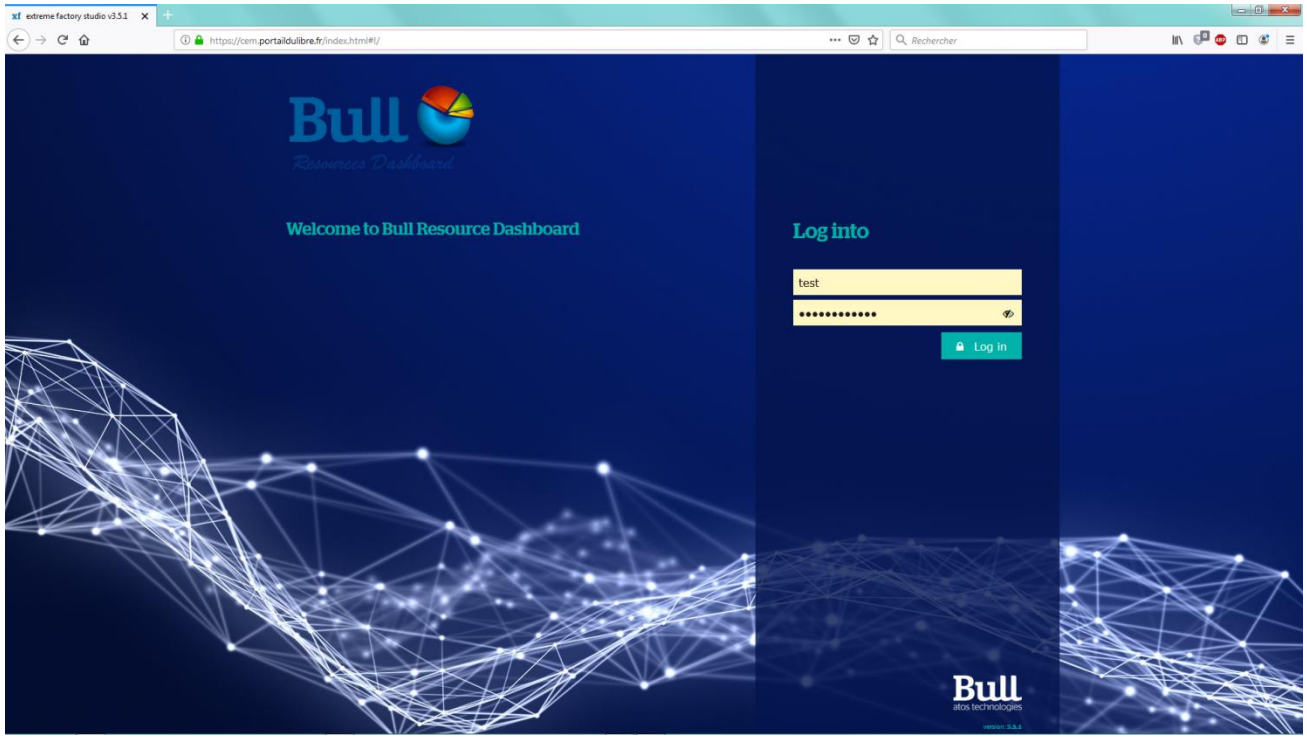

BEM

Current resource usage 1/3

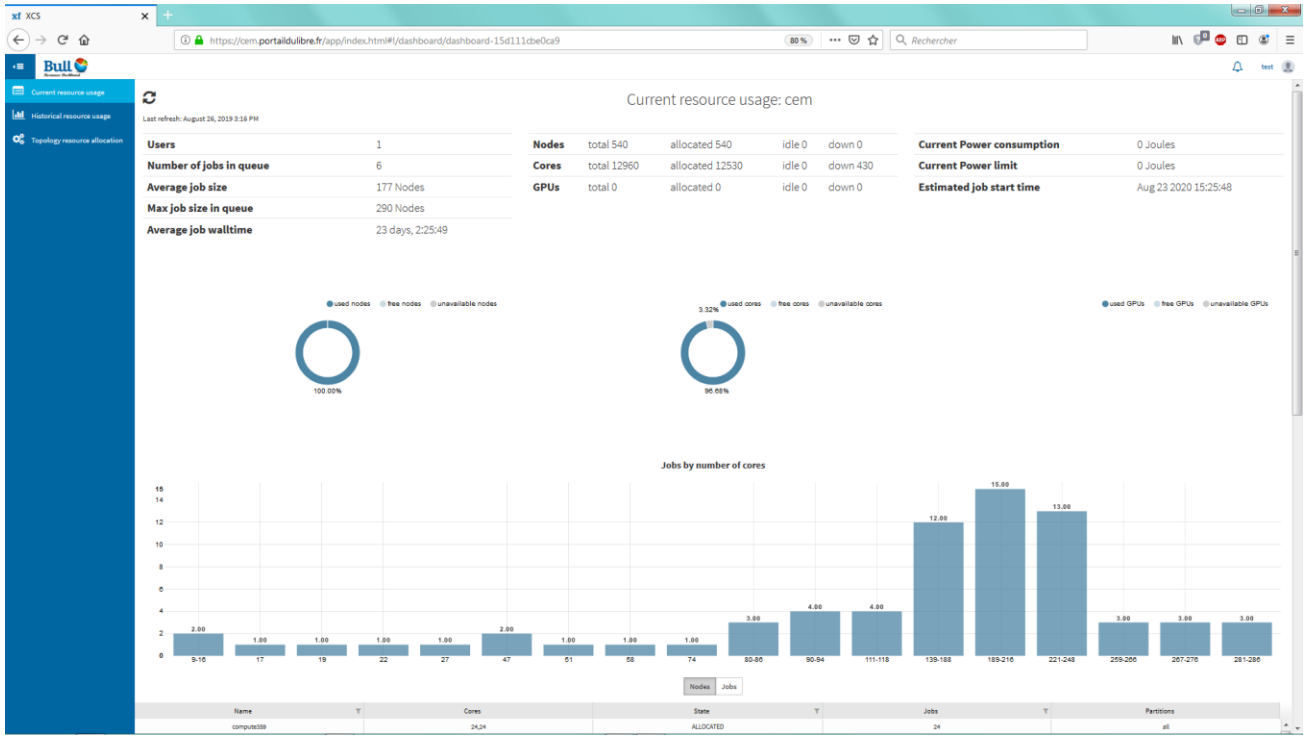

BEM

Current resource usage 2/3

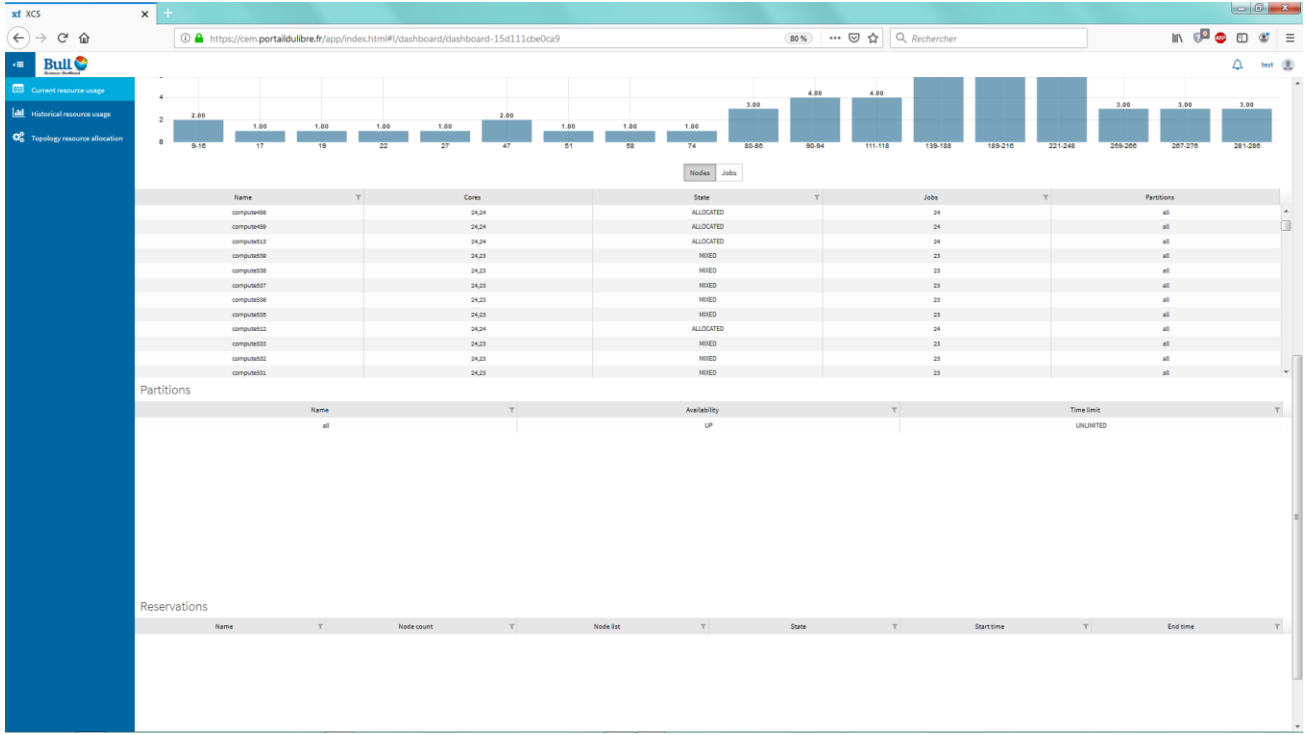

BEM

Current resource usage 3/3

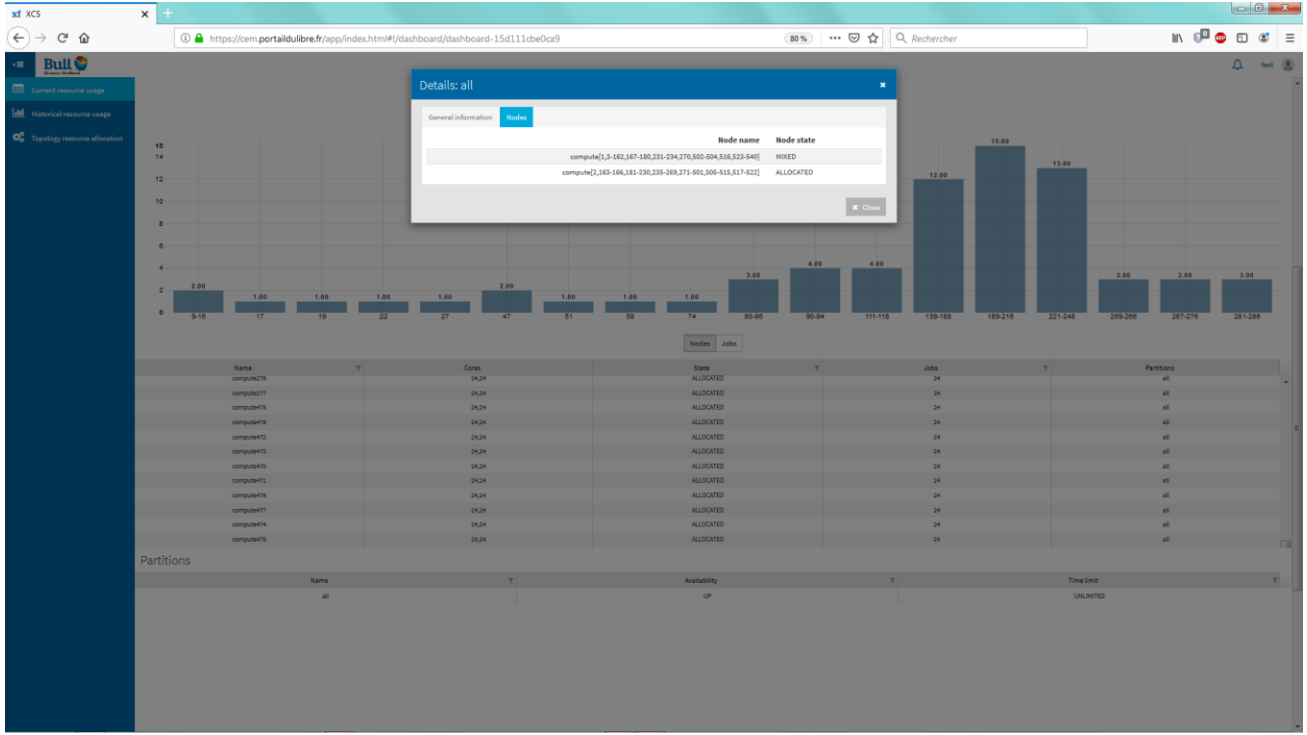

BEM

Historical resource usage

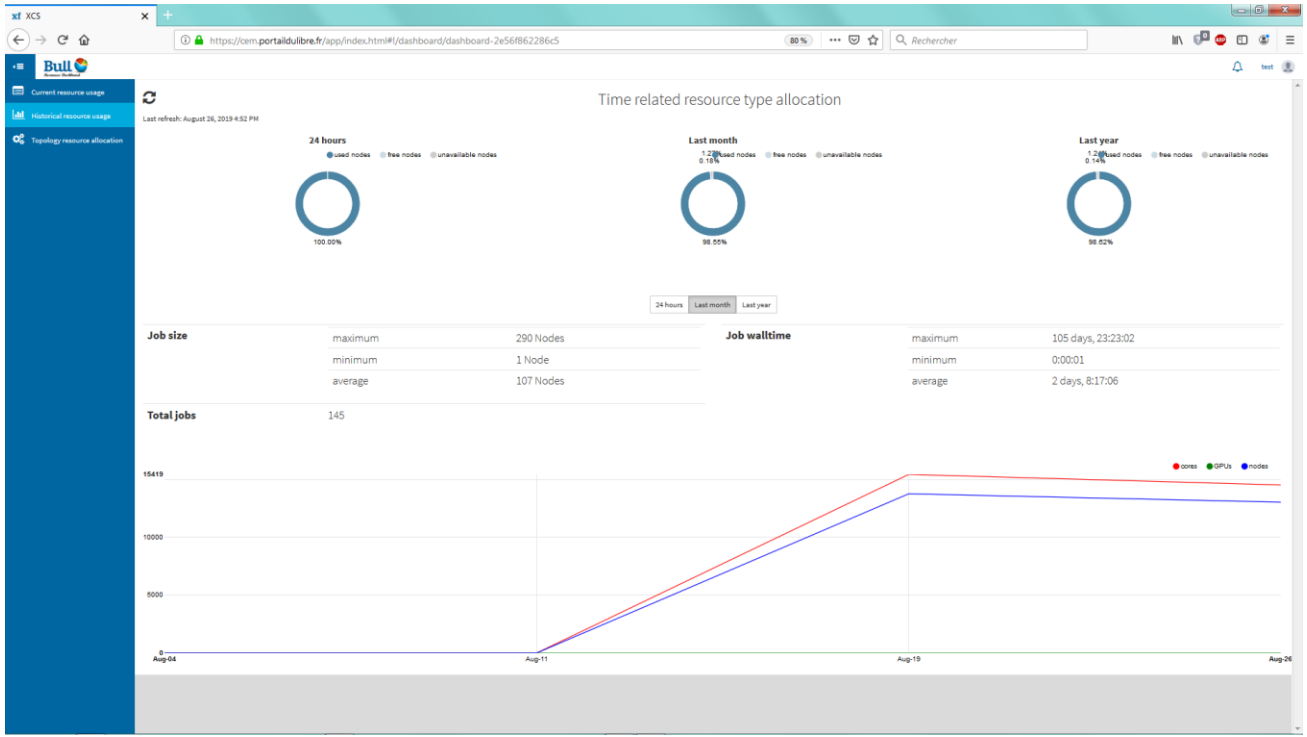

BEM

Topology resource allocation 1/3

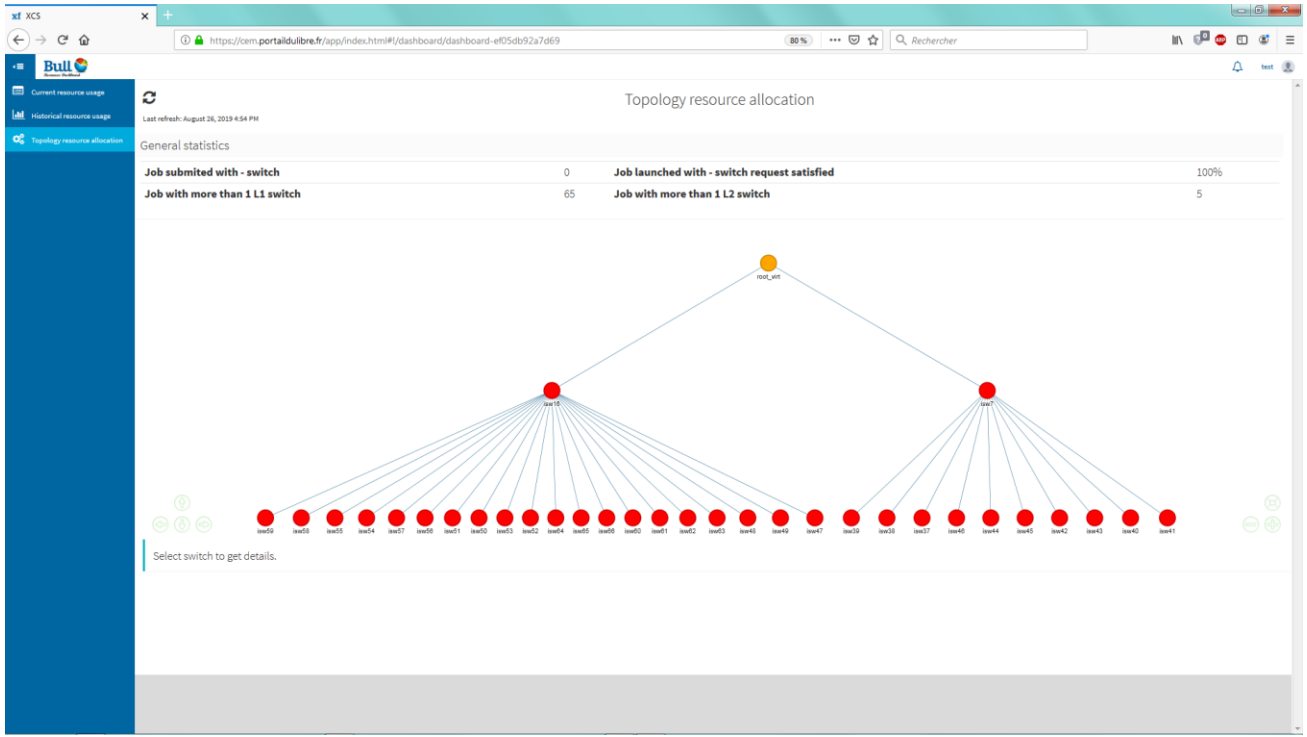

BEM

Current resource usage 2/3

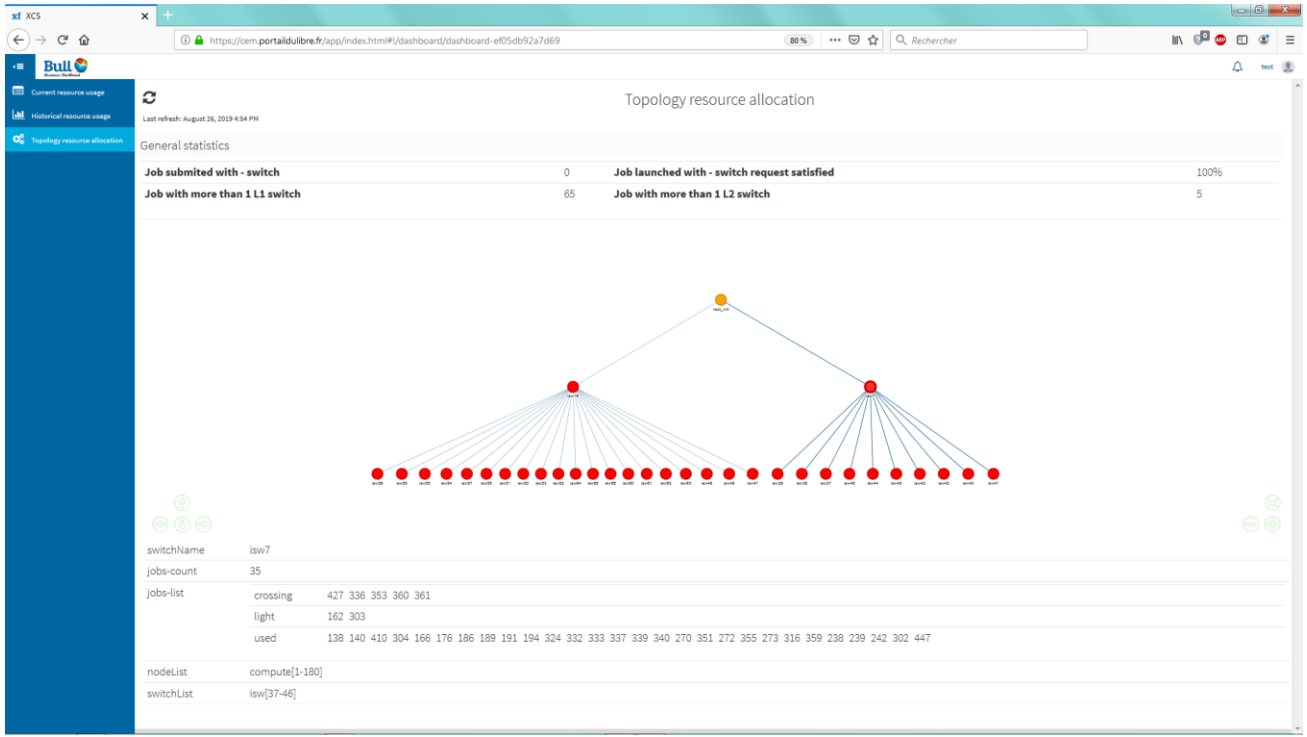

BEM

Current resource usage 3/3

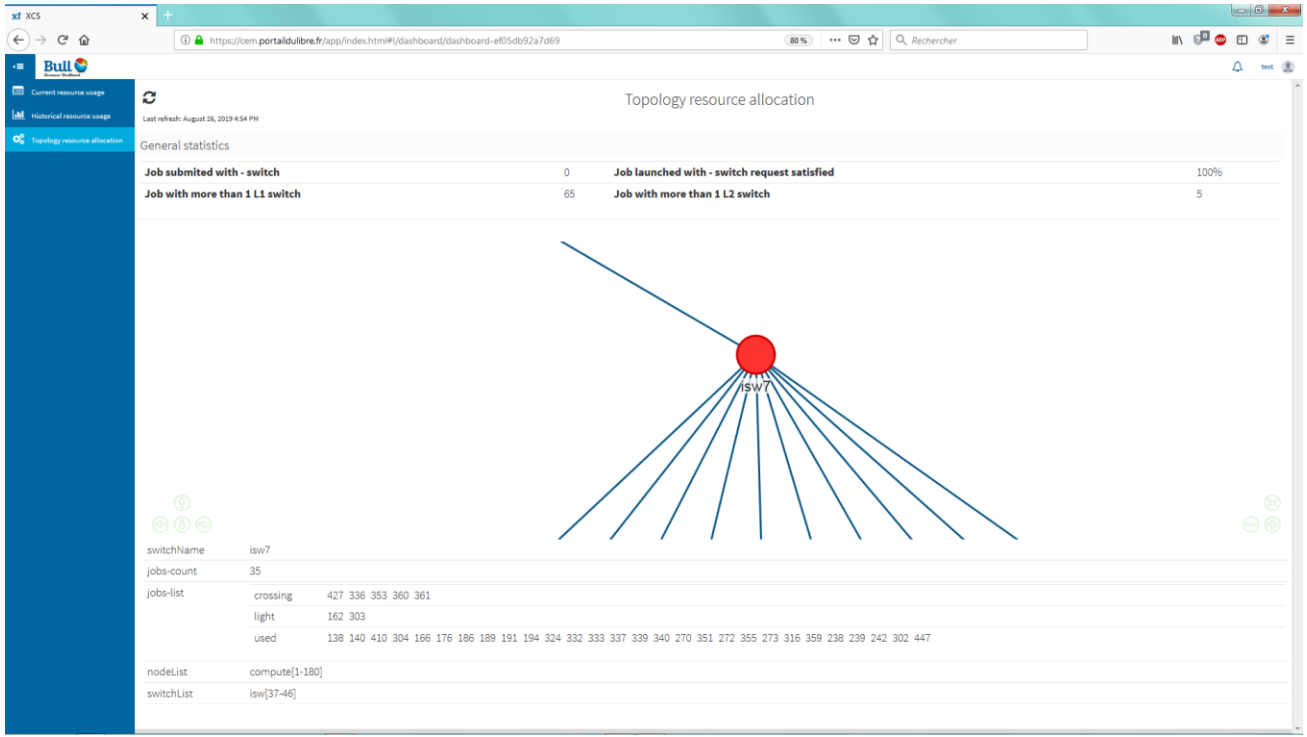

Conclusion & Future Work

Conclusions

- XCS is successfully used in production on many sites for several years and it evolves continuously
- BEM is still under development and the first Minimal Viable Product (MVP) is very promising
- Mobile devices are becoming a new standard way for doing “everything”, so such a web portal approach will soon be mandatory for new users (unexperienced users, young scientist of the new generation, non-technical managers, etc.)

On going and future work

- Unify both interfaces (XCS & BEM) and share a unique security service
- Add new features to administrate Slurm
- We develop a new web portal framework to federate all our HPC, AI & Quantum tools/microservices. It is an evolution of our current XCS solution with:
  - a generic web GUI framework
  - a security service (with flexible identity, authentication with SSO and authorization management).
  - global services (reverse proxy, gateway, discovery service, etc.)

XCS and BEM architecture

Complete solution to be developed in 2020

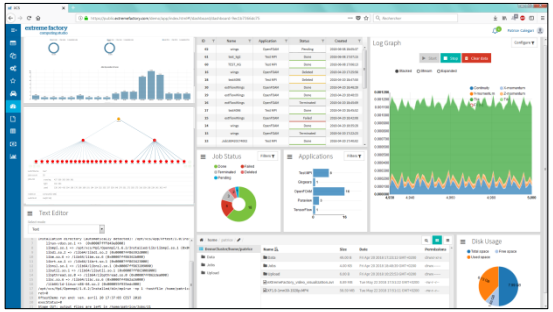

NEW unified web User Interface

HTTPS

DC API

Job submission DC

XCS DCs

Data mngmt DC

Unified GUI web server
- Dashboards
- Web Design

DC API

Switch Topology DC

BEM DCs

Slurm usage history DC

HTTPS

HTTPS

HTTPS

XCS REST API web server

Security service

BEM REST API web server

SSH

SSH

HPC cluster integration layer
- Slurm
- HPC applications

XCS Data base

Directory service

BEM Data base

BEM integration layer
- Slurm

DC = Dashboard Component

XCS and Slurm native REST service architecture

Possible evolution…

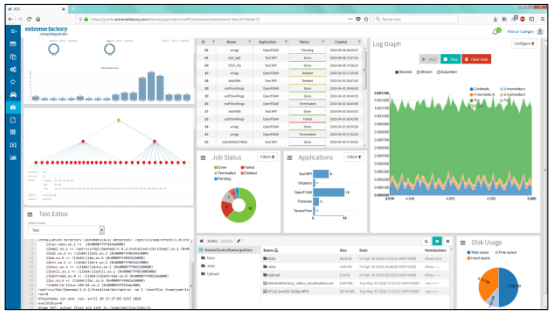

NEW unified web User Interface

HTTPS

DC API

Job submission DC

XCS DCs

Data mngmt DC

Unified GUI web server

- Dashboards
- Web Design

DC API

Slurm job specific DC

Slurm DCs

Slurm admin specific DC

HTTPS

HTTPS

HTTPS

HTTPS

XCS REST API web server

SSH

HPC cluster integration layer

- Slurm
- HPC applications

XCS Data base

Security service

Directory service

Slurm REST API deamon

slurm.restd

Slurm Data base

Slurm server

DC = Dashboard Component

Thank you

For more information please contact:
Mathis Clayer for Slurm topics (mathis.clayer@atos.net)
Patrice Calegari for GUI topics (patrice.calegari@atos.net)

Atos

More on HPC web portals

- **Web Portals for High-performance Computing: A Survey**
  - 36 page journal paper published by ACM
  - https://dl.acm.org/citation.cfm?id=3197385
- **Democratization of HPC through the Use of Web Portals: Different Strategies**
  - Panel at SC'19 in Denver, November 20th, 3:30pm-5pm
  - https://sc19.supercomputing.org/presentation/?id=pan102&sess=sess223