



A fully configurable HPC web portal for managing Slurm jobs

Patrice Calegari

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

We will talk about...

1

Context of the projects

2

XCS - eXtreme factory Computing Studio

3

BEM - Bull Efficiency Manager

4

Conclusion and future work

1

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

2

XCS

eXtreme factory

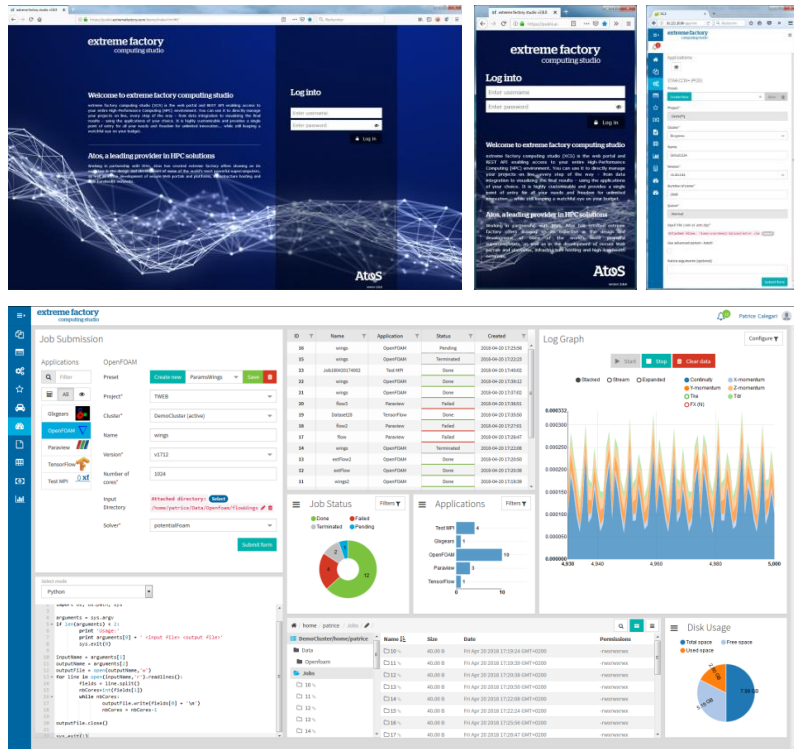
Computing Studio

Extreme factory Computing Studio v3 (XCS3)

Introduction

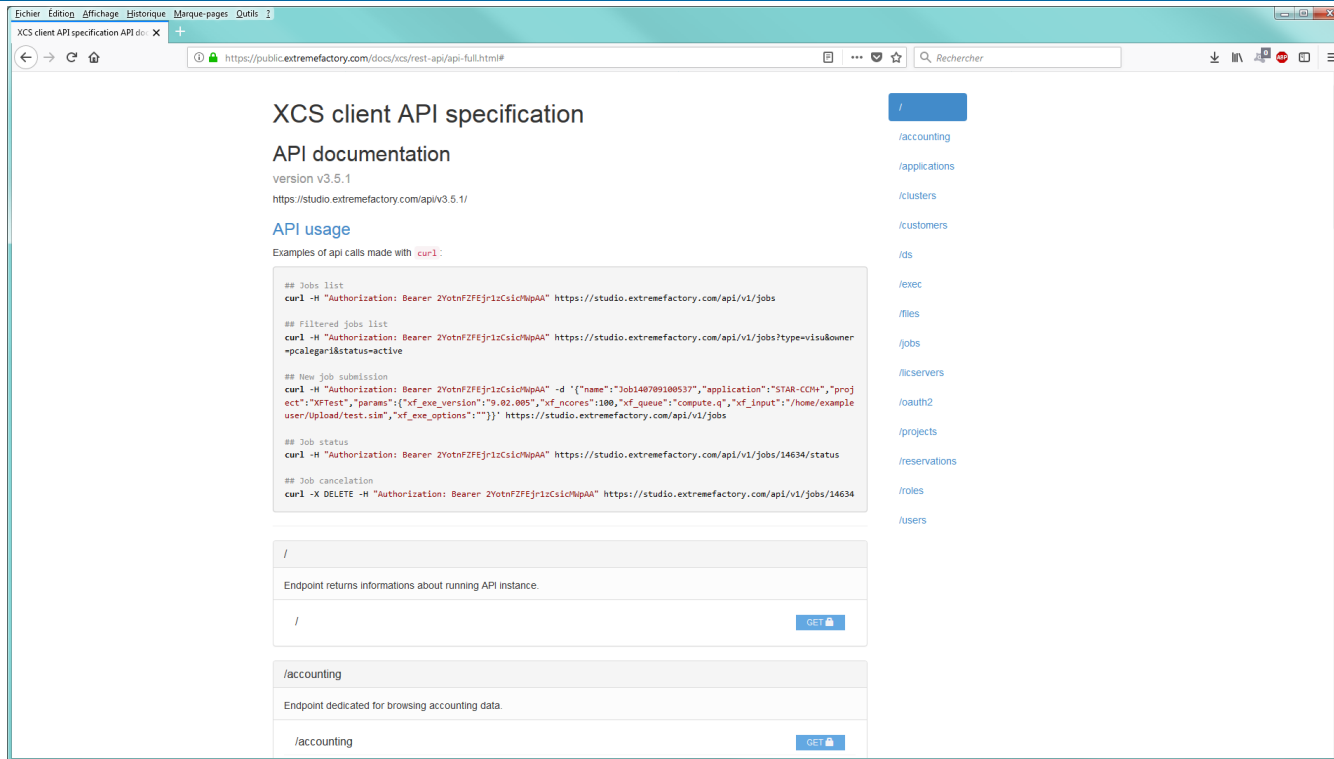
- ▶ **Modular HPC, AI & Quantum portal**
 - as-a-Service cornerstone application,
 - supports **Slurm** (and other schedulers)
 - Role Based Access Control (RBAC)
 - supports AD, LDAP (with Kerberos)
 - XCS = REST API service + GUI
- ▶ **Fully customizable user interface**
 - Responsive Web Design (RWD) 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>



The screenshot shows a web browser window displaying the 'XCS client API specification' page. The page title is 'XCS client API specification' and the subtitle is 'API documentation'. The version is 'version v3.5.1' and the URL is 'https://studio.extremefactory.com/api/v3.5.1/'. There is a link for 'API usage' and a section for 'Examples of api calls made with curl:'. The examples include commands for listing jobs, filtering jobs, submitting a new job, checking job status, and canceling a job. On the right side, there is a navigation menu with links to various API endpoints: /, /accounting, /applications, /clusters, /customers, /ids, /exec, /files, /jobs, /licservers, /oauth2, /projects, /reservations, /roles, and /users. The main content area shows the endpoint '/' with a description 'Endpoint returns informations about running API instance.' and a 'GET' button. Below it, the endpoint '/accounting' is shown with a description 'Endpoint dedicated for browsing accounting data.' and a 'GET' button.

XCS REST API

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

The screenshot displays a web browser window showing the XCS REST API documentation. The address bar indicates the URL: <https://public.extremefactory.com/docs/xcs/rest-api/api-full.html#>. The page content is organized into sections for different API endpoints:

- /jobs**: Endpoint dedicated for jobs management.
 - `/jobs`: GET, POST, DELETE
 - `/jobs/{id}`: GET, DELETE
 - `/jobs/{id}/status`: GET
 - `/jobs/{id}/connection`: GET
 - `/jobs/{id}/share`: GET, POST, DELETE
- /licservers**:
 - `/licservers`: GET, POST
 - `/licservers/{serverId}`: GET, PUT, DELETE
 - `/licservers/{serverId}/host`: GET
 - `/licservers/{serverId}/management`: PUT
 - `/licservers/{serverId}/license`: GET, POST
 - `/licservers/{serverId}/logs`: GET
 - `/licservers/{serverId}/status`: GET
- /oauth2**: Following two sub-endpoints, dedicated to handle user authentication, implements subset of RFC 6749 OAuth 2.0 Authorization Framework spec. According this specification endpoints should be accessible only to authenticated clients (requires basic authentication).
 - `/oauth2/access_token`: POST
 - `/oauth2/check_token`: GET
 - `/oauth2/revoke`: POST

A right-hand sidebar contains a navigation menu with the following items: /, /accounting, /applications, /clusters, /customers, /ids, /exec, /files, /jobs (highlighted), /licservers, /oauth2, /projects, /reservations, /roles, and /users.

XCS REST API

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

The screenshot shows a web browser window displaying the XCS REST API documentation. The browser's address bar shows the URL: `https://public.extremefactory.com/docs/xcs/rest-api/api-full.html#jobs_get`. The page content is a modal window for the `GET /jobs` endpoint. The modal contains the following information:

- Endpoint:** `GET /jobs`
- Description:** Retrieve jobs list available to authorised user. It can be jobs (computation or visualization) submitted by user as well as jobs submitted by other users in projects he is privileged to.
- Example jobs request with query parameters:**

```
GET /api/jobs?type=vis&user=testUser&status=active HTTP/1.1
Host: studio.extremefactory.com
Accept: application/json
Authorization: Bearer 2YotnFZFEj1zCicMqAA
```
- Security:** Secured by `oauth_2_0`
- Request:** Response
- Headers:**
 - xcs-target-cluster (string)**: Header is mandatory in case of multiple clusters available, otherwise status 404 with message "Header 'xcs-target-cluster' is missing" will be returned. Parameter is ignored when only one cluster is available. In header should be put cluster address, which is returned by `/clusters` endpoint.
Example:

```
xcs-target-cluster: 10.185.56.137:22
```
 - Authorization (string)**: Used to send a valid OAuth 2 access token. Do not use with the "access_token" query string parameter.
- Query Parameters:**
 - offset (integer)**: Sets cursor position from which results should be returned.
Example:

```
20 - will return all entries starting from 21st
```
 - limit (integer)**: Limits returned records to specified number. Returning less records is acceptable.
Example:

XCS user dashboard

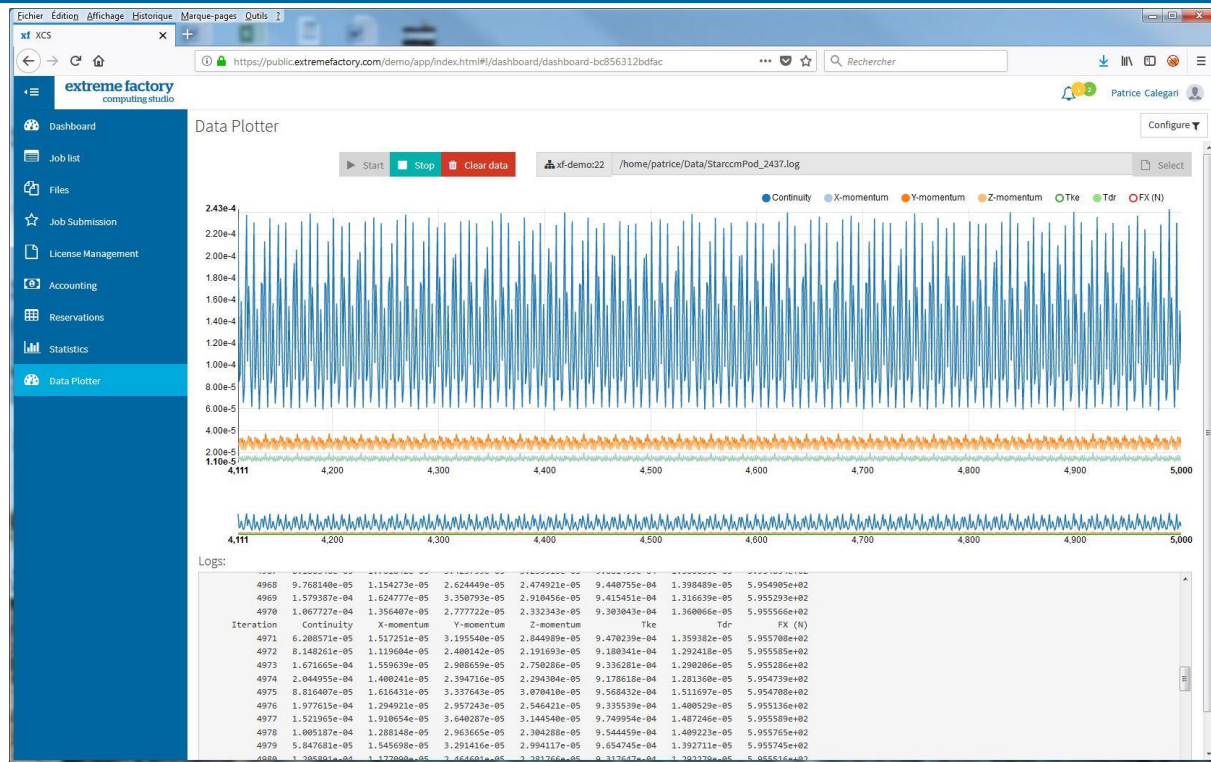
Example 1: 8 components

The screenshot displays the XCS user dashboard interface, which is divided into several functional areas:

- Job Submission:** A form for submitting new jobs. Fields include Project (TWEB), Cluster (DemoCluster), Name (wings), Version (v1712), Number of cores (4096), Input Directory, and Solver (potentialFoam). A "Submit form" button is at the bottom.
- Job Status:** A table listing jobs with columns for ID, Name, Application, Status, and Created. A pie chart below shows the distribution of job statuses: Done (green), Failed (red), Terminated (grey), Deleted (orange), and Pending (blue).
- Log Graph:** A line graph showing the progress of a job over time. The x-axis represents time (from 4,930 to 5,000) and the y-axis represents a numerical value (from 0.000000 to 0.001286). The graph shows a fluctuating green line, indicating the job's progress.
- Application Status:** A bar chart showing the number of jobs for different applications: Test MPI (8), Glgears (1), OpenFOAM (16), Paraview (3), and TensorFlow (1).
- Disk Usage:** A pie chart showing the total, free, and used space on the disk. The total space is 7.99 GB, with 6.13 GB used and 1.86 GB free.
- Text Editor:** A code editor showing the installation directory and the contents of a file named "xrcs".

XCS user dashboard

Example 2: 1 component



XCS user dashboard

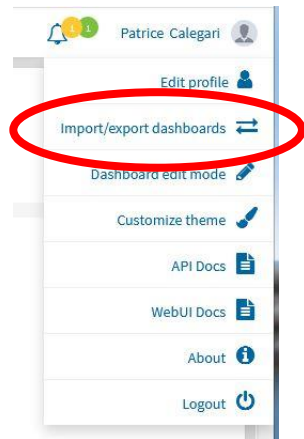
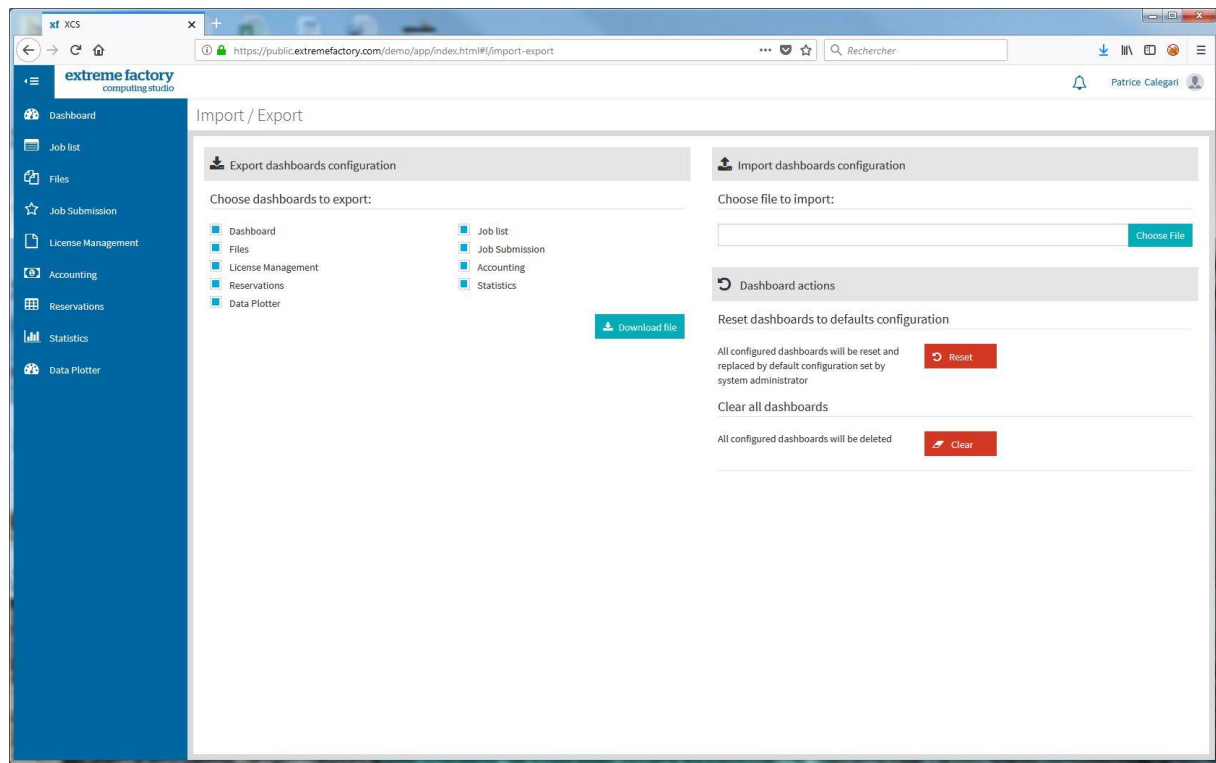
Example 3: 6 components with edited theme

The screenshot displays the XCS user dashboard interface, which is divided into several functional areas:

- Left Sidebar:** A navigation menu with options: Files, Jobs, Accounting, Reservations, Job list, Submit Job, Statistics, License Management, and Dashboard (highlighted).
- Applications Panel:** A configuration area for STAR-CCM+ (POD) with fields for Project (Demo), Cluster (XFPPU (8) (active)), Name (fluid), Version (12.06.011), Number of cores (2048), Queue (Normal), Input File (/home/pmartin/Upload/sae_step4/sae_step4.sim), and Compress result file? (Yes .zip). A "Submit form" button is at the bottom right.
- Job List Table:** A table with columns: ID, Name, Application, Status, Status..., Created, Updated, Start date, and Work. It lists various jobs such as fluid, Shuttle..., crash, wings, crash2, Job170..., and Boat.
- File Manager:** A view of the file system for user pmartin, showing a directory structure with folders like Jobs, Upload, DeepLinesData, EnsigntData, Opti, Pamcrash, and StarviewData, each with its size and date.
- Disk Usage:** A pie chart showing 1.20 TB total space, with 0.85 TB used and 0.35 TB free.
- Job Status:** A bar chart showing the count of jobs in different states: Done (134), Failed (13), Terminated (21), Terminating (1), Pending (2), and Active (1).
- Links and Downloads:** A list of links including Quick User Guide (pdf file), XRV client (Less for Windows), XRV client (for Linux), and Atos home page.

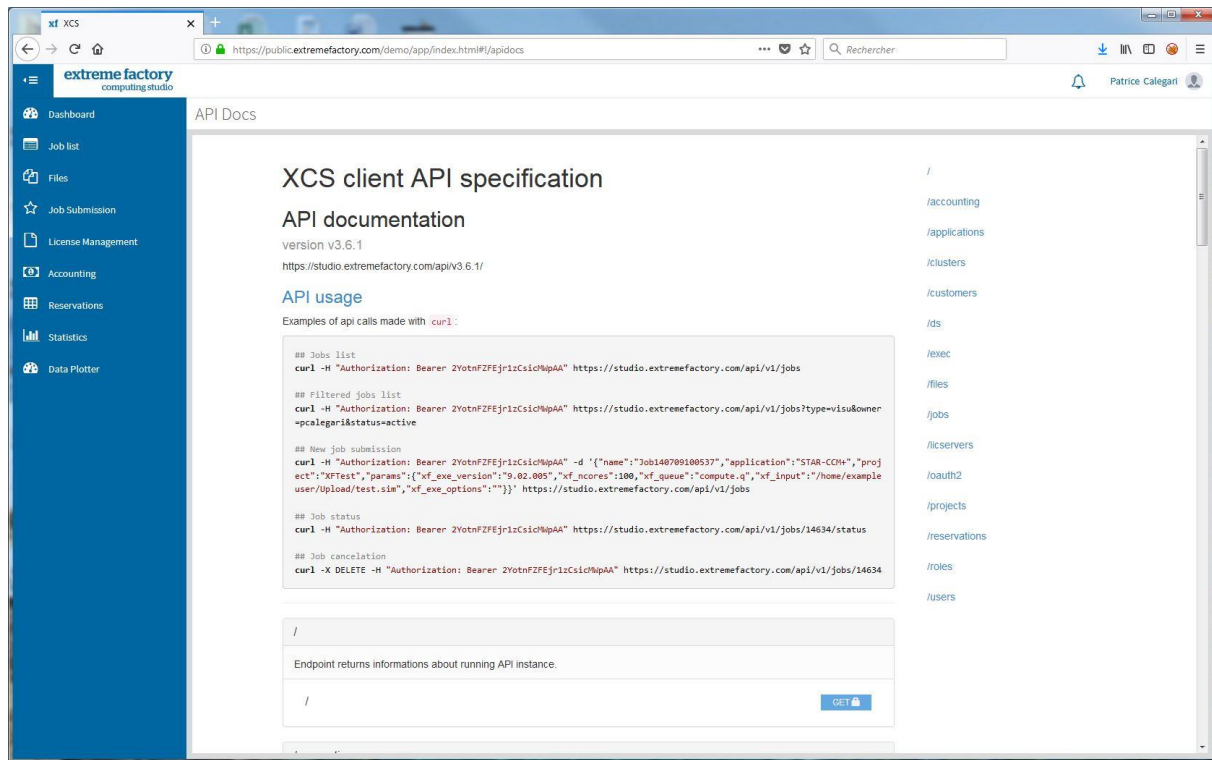
XCS dashboard main menu

import/export dashboards

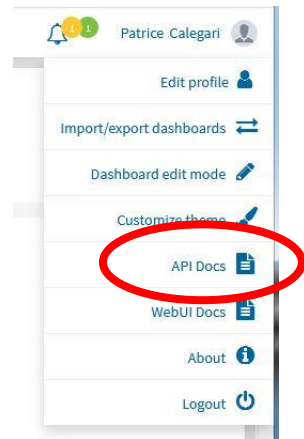


XCS dashboard main menu

REST API documentation



The screenshot shows a web browser window displaying the XCS dashboard. The main content area is titled "XCS client API specification" and "API documentation", version v3.6.1. It includes a section for "API usage" with examples of curl commands for various endpoints like /jobs, /jobs?type=visu&owner=pcalegari&status=active, and /jobs/14634/status. A sidebar on the left contains navigation links: Dashboard, Job list, Files, Job Submission, License Management, Accounting, Reservations, Statistics, and Data Plotter. The top right of the dashboard shows the user's name "Patrice Calegari" and a notification bell icon.



The user profile menu is displayed, showing the user's name "Patrice Calegari" and a notification bell icon. The menu items are: Edit profile, Import/export dashboards, Dashboard edit mode, Customize theme, API Docs (highlighted with a red circle), WebUI Docs, About, and Logout.

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

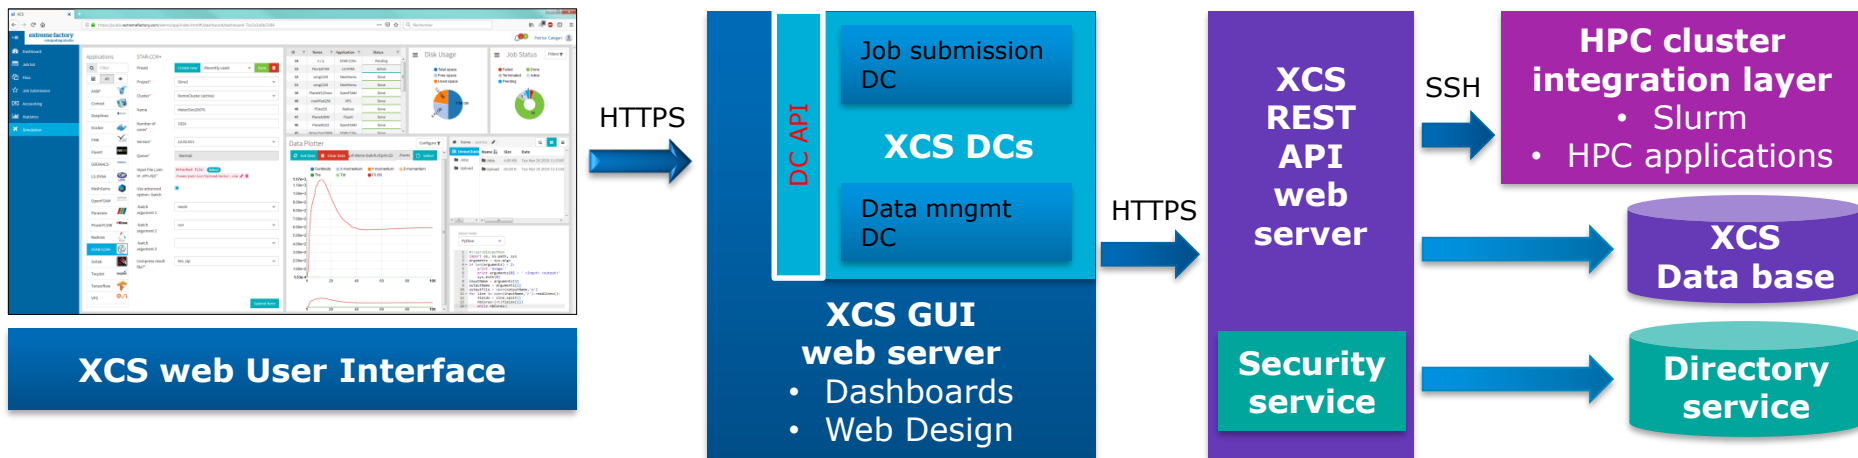
- The solution should be used in existing environments without modifying them

Keep all the intelligence in the REST API server

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

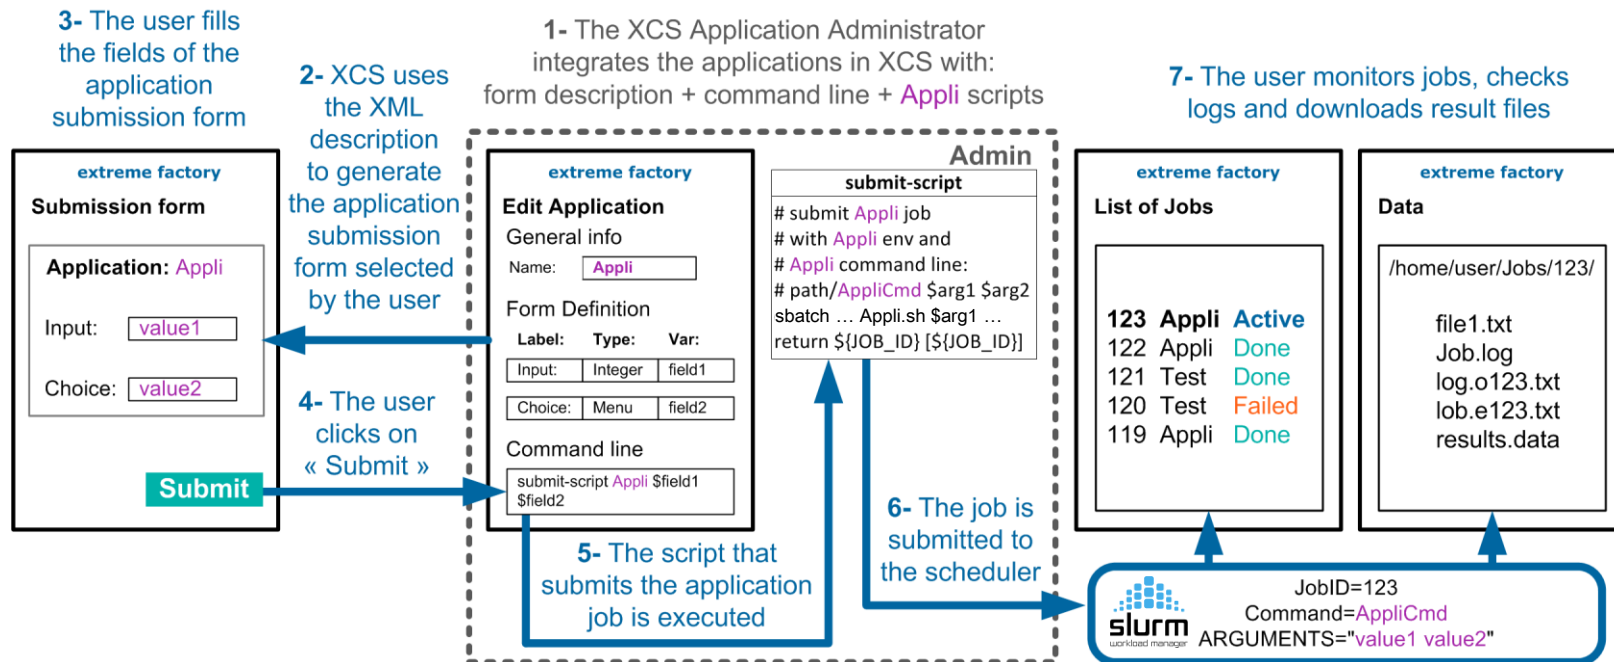
XCS architecture

current v3



DC = Dashboard Component

Slurm job submission workflow with XCS



XCS application administrator dashboard

HPC application general information

The screenshot displays the XCS application administrator dashboard. The left sidebar contains navigation options: Admin Tables, Clusters, Directory Services, Customers, Projects, Users, Applications (selected), Application list, Licenses, Roles, Accounting, Credit Formula, Credits, Dashboard Presets, and Statistics. The main content area is titled 'Application Administration' and is divided into two sections: 'Applications list' and 'Edit application: Fluent'.

Applications list:

Application Name	Type	Id
Fluent	Computation	id: 5
Glgears	Visualization	id: 2
OpenFOAM	Computation	id: 8
Paraview	Visualization	id: 3
STAR-CCM+	Computation	id: 6
TensorFlow	Computation	id: 9
Tensorflow DL	Computation	id: 4
Test MPI	Computation	id: 1

Edit application: Fluent (General information tab):

- Name: Fluent
- Description: Fluid Dynamics application from ANSYS
- Type: Computation
- Supported file extensions: sim, X (with an 'Add extension' button)
- Maximum wall time: 1000:00:00
- Job name: Choose the way of job name creation * (radio buttons for 'Defined by user' and 'Generated using pattern')
- Check job name uniqueness:
- Software name: FluentDemo
- Command: `xf_run -s ${application} -v ${version} -i ${input} -q ${queue} -n ${ncores} -j ${jobName} -- ${precision}`
- Availability: (empty field)

At the bottom right, there are buttons for 'Delete', 'Save as copy', and 'Save'.

XCS application administrator dashboard

HPC application form definition

The screenshot displays the XCS application administrator dashboard. The left sidebar contains a navigation menu with the following items: Admin Tables, Clusters, Directory Services, Customers, Projects, Users, Applications (highlighted), Application list, Licenses, Roles, Accounting, Credit Formula, Credits, Dashboard Presets, and Statistics. The main content area is titled 'Application Administration' and is divided into two sections: 'Applications list' and 'Edit application: Fluent'. The 'Applications list' section shows a table of applications with columns for name, type, and ID. The 'Edit application: Fluent' section is currently on the 'Forms definition' tab, which is divided into 'Form components' and 'Form pages'. The 'Form components' section lists various input types such as checkbox, radio, options, options_ds, file, directory, file_and_directory, integer, decimal, number, area, text, text_without_space, and comment. The 'Form pages' section shows a form for 'STAR-CCM+' with fields for Version, Number of cores, Queue, Input File (with an attached file), and Precision. At the bottom right of the dashboard, there are buttons for 'Delete', 'Save as copy', and 'Save'.

Application Name	Type	ID
Fluent	Computation	id: 5
Glxgears	Visualization	id: 2
OpenFOAM	Computation	id: 8
Paraview	Visualization	id: 3
STAR-CCM+	Computation	id: 6
TensorFlow	Computation	id: 9
Tensorflow DL	Computation	id: 4
Test MPI	Computation	id: 1

3

BEM

Bull Efficiency Manager

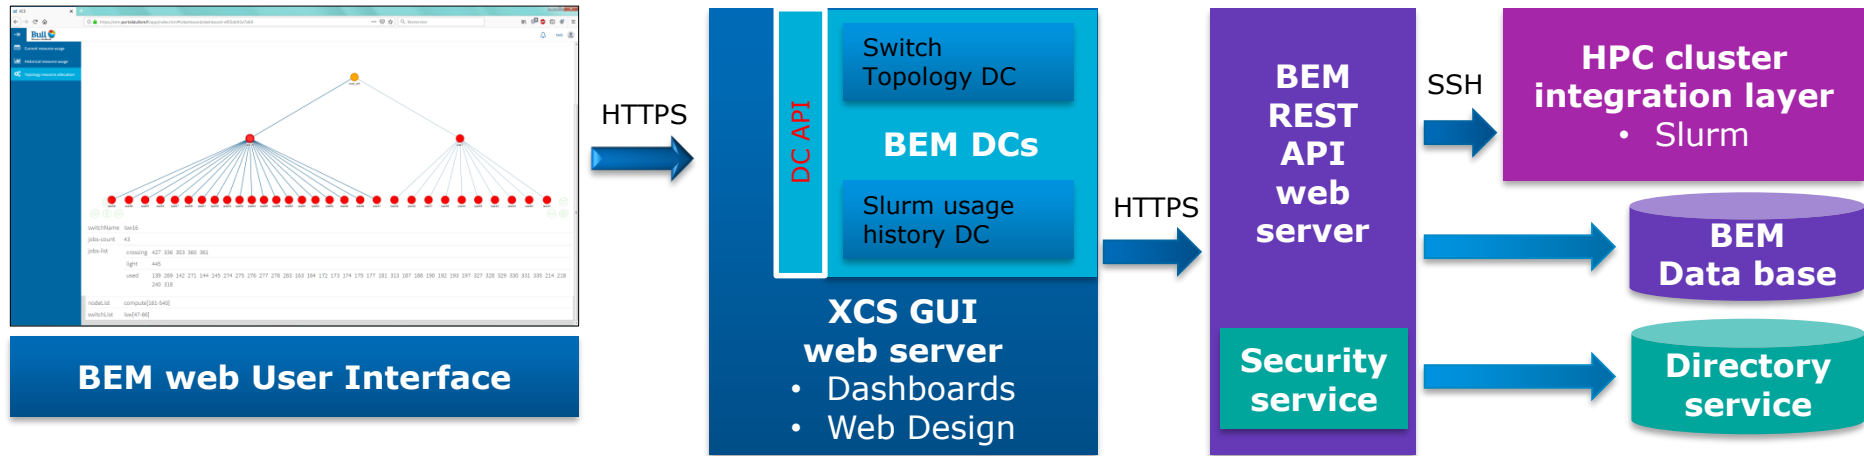
Bull Efficiency Manager (BEM)

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

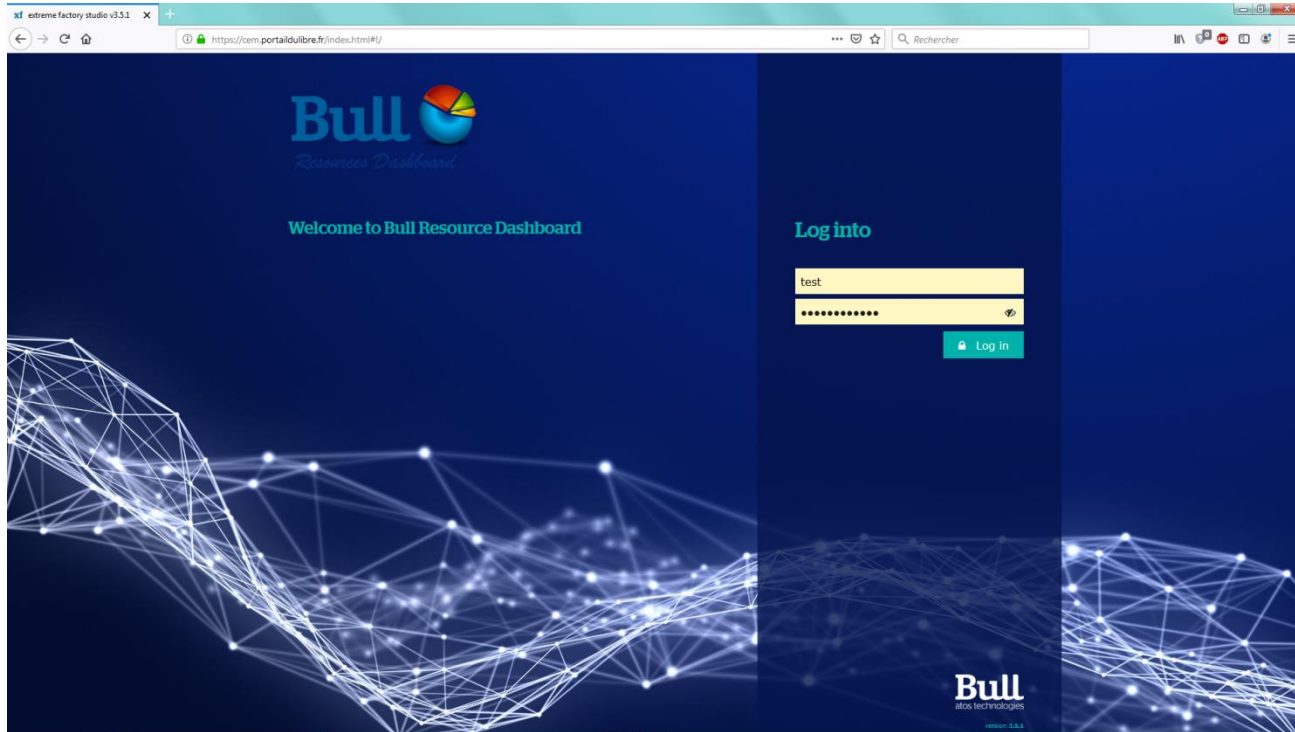
current v3



DC = Dashboard Component

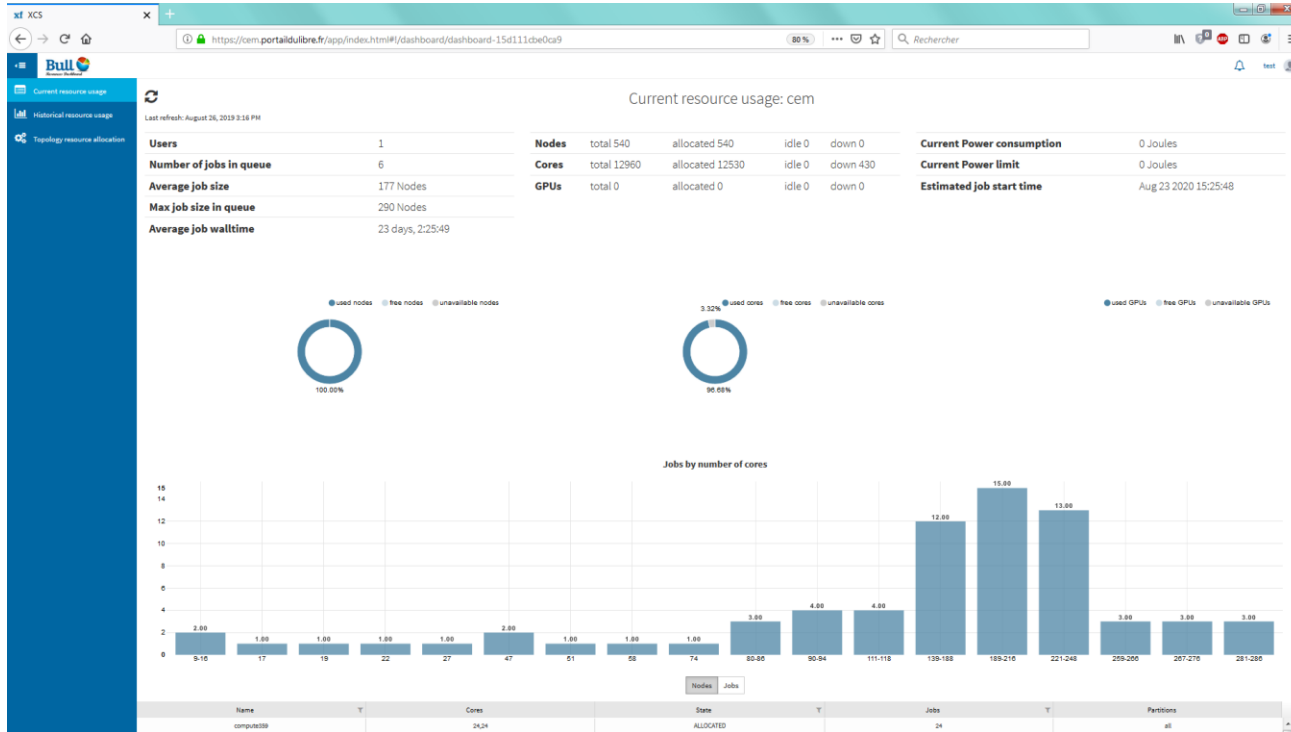
BEM

Login Page



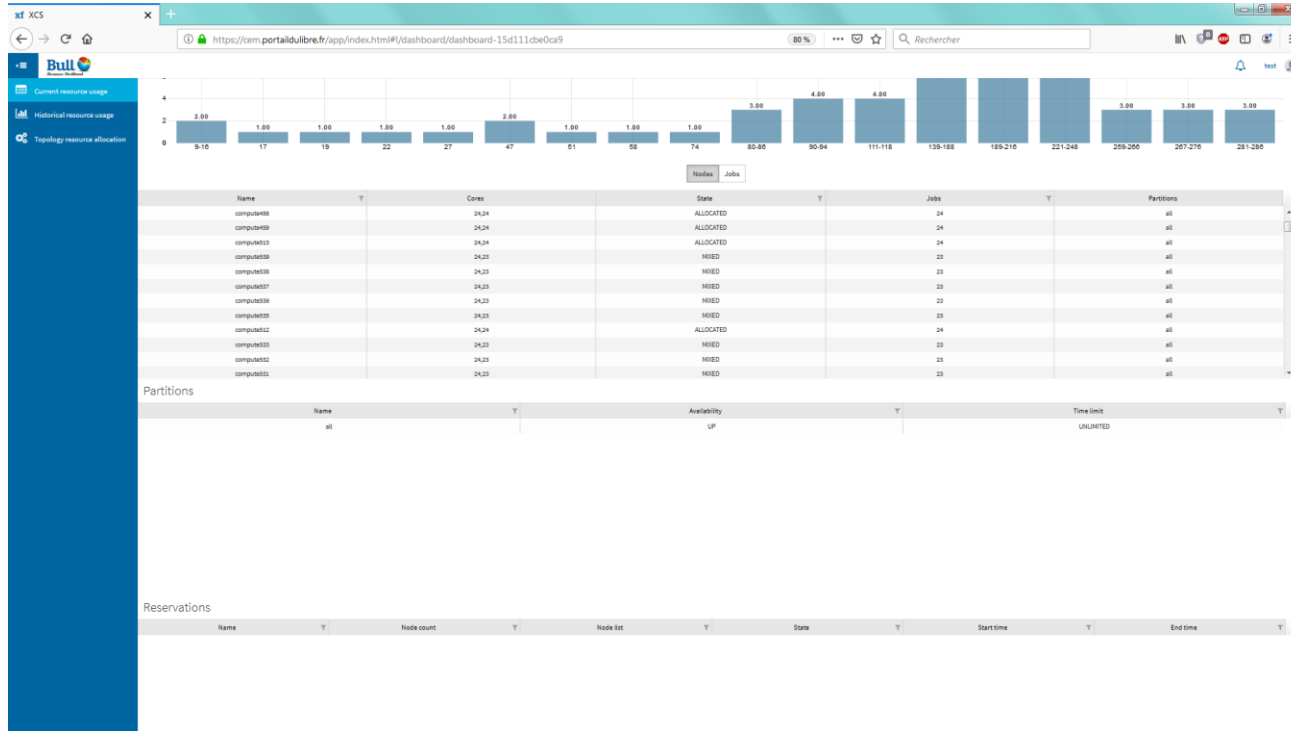
BEM

Current resource usage 1/3



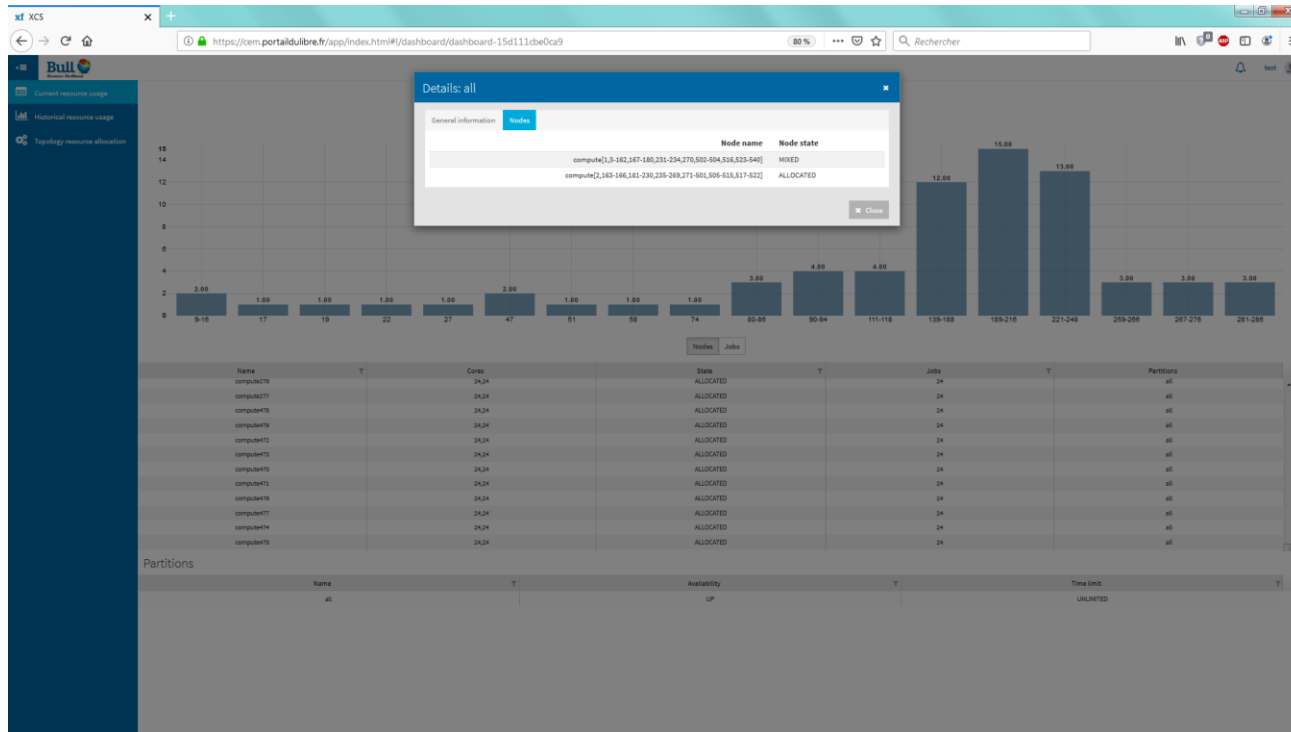
BEM

Current resource usage 2/3



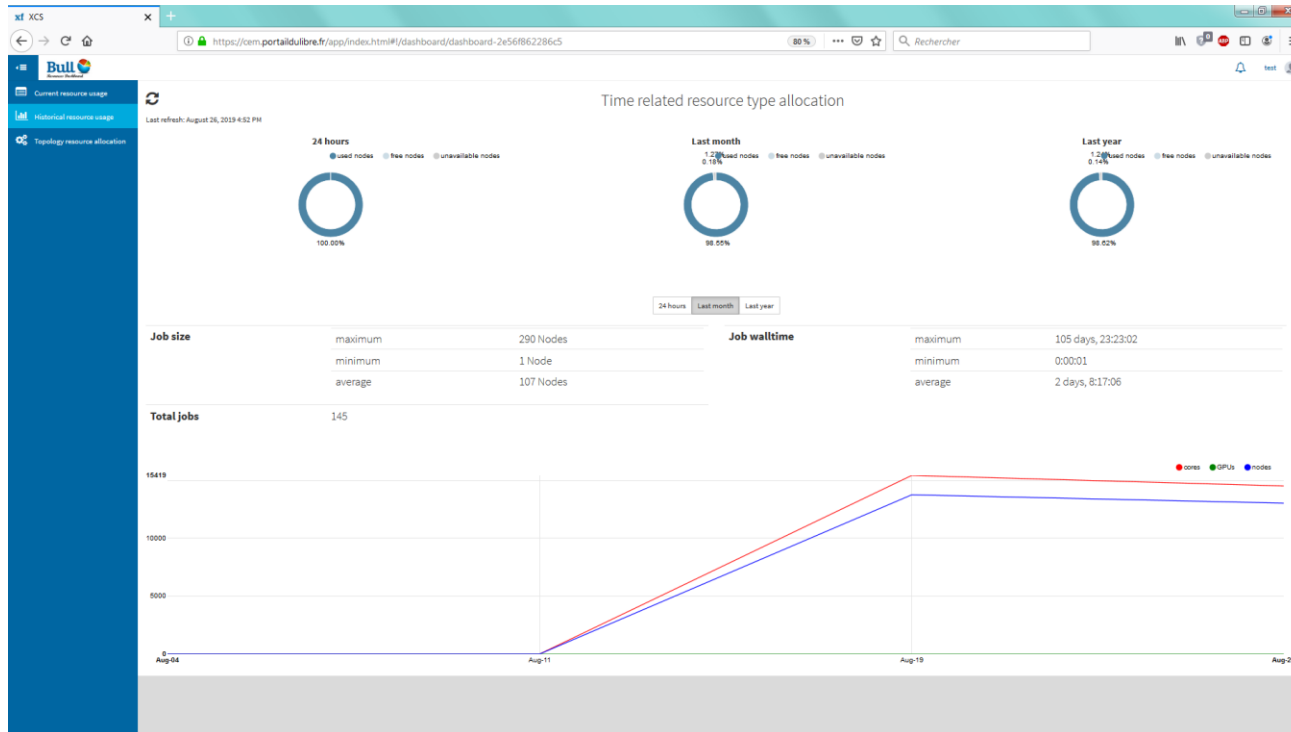
BEM

Current resource usage 3/3



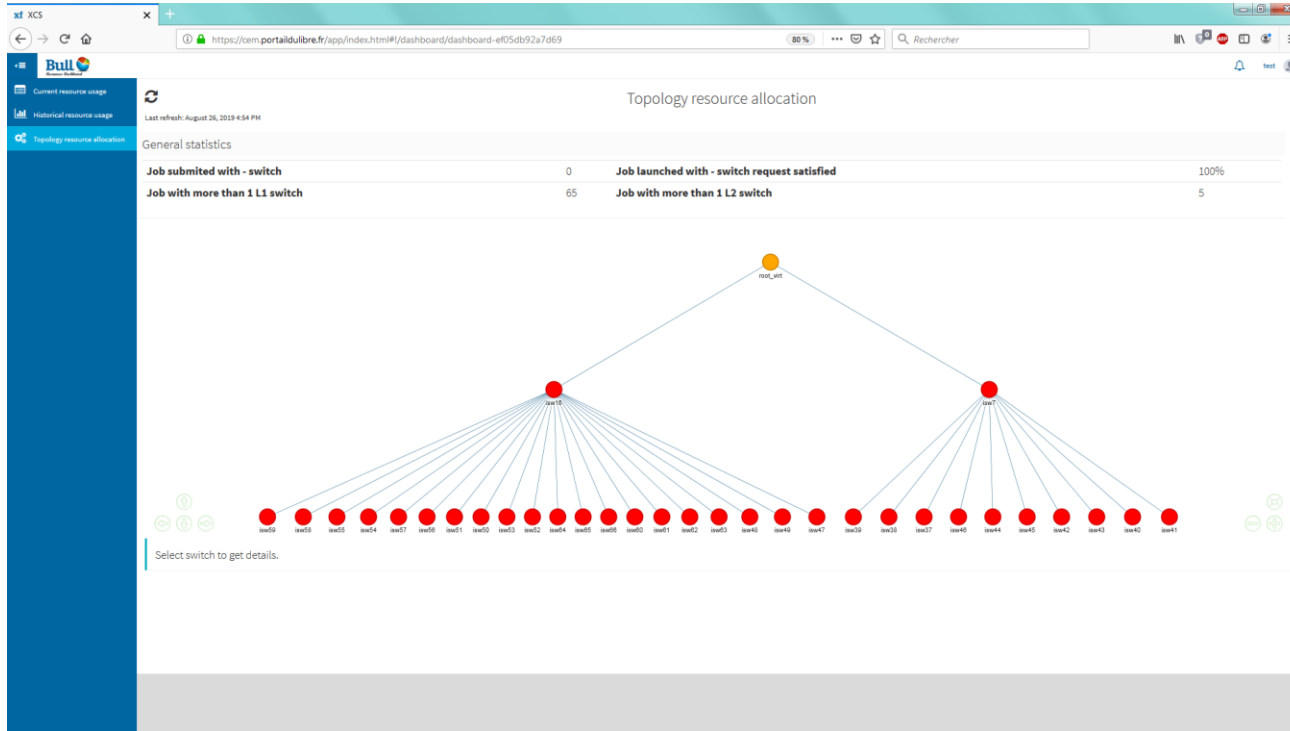
BEM

Historical resource usage



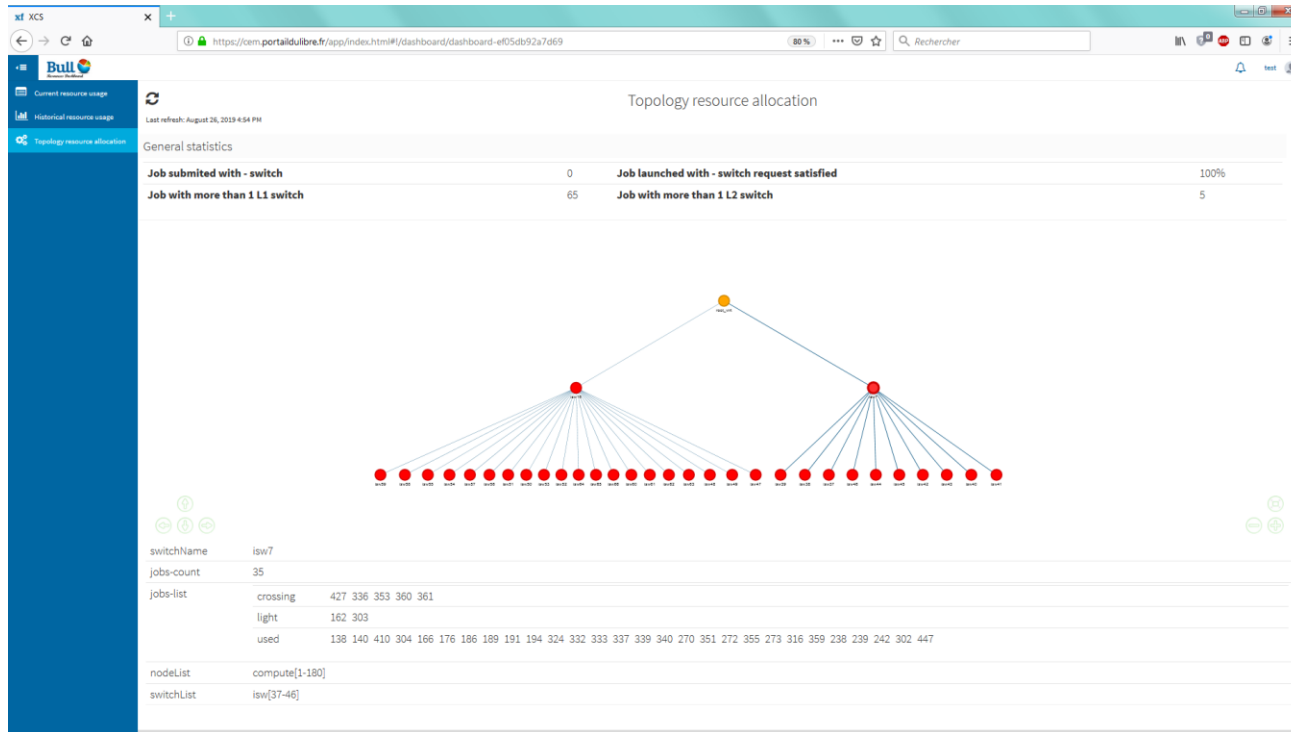
BEM

Topology resource allocation 1/3



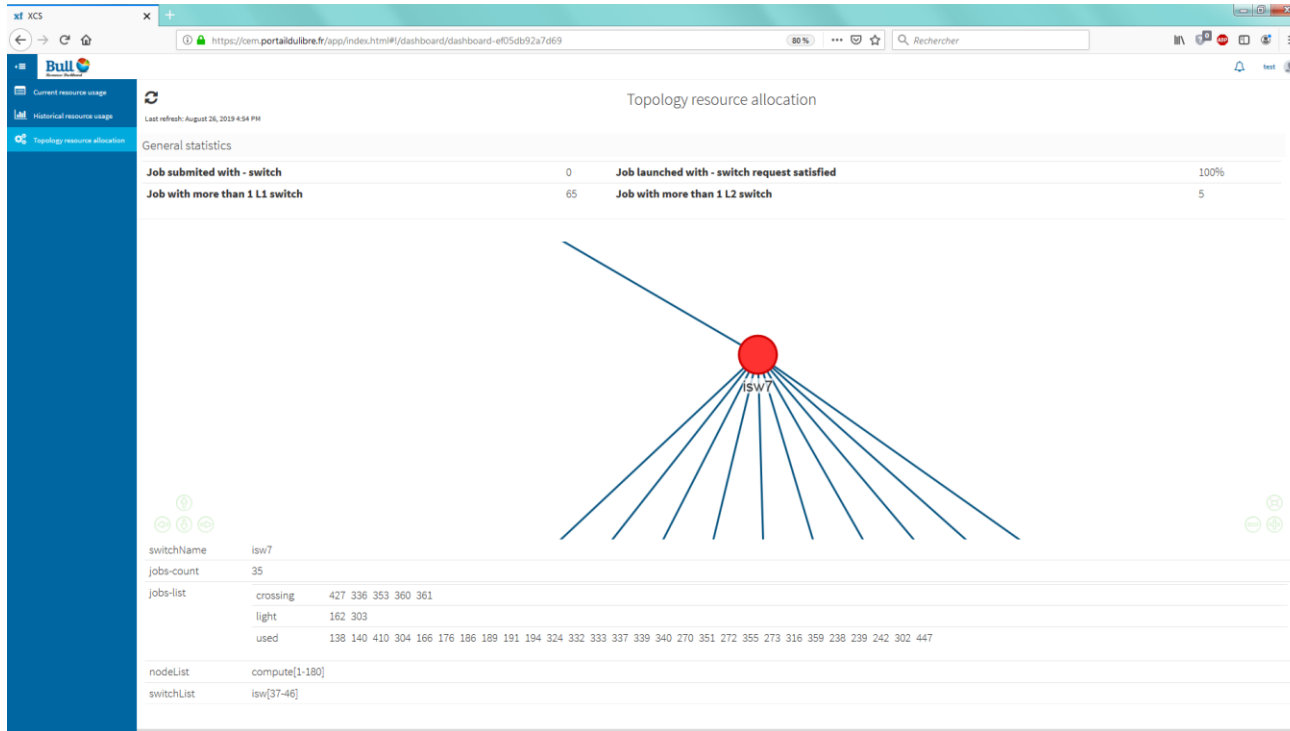
BEM

Current resource usage 2/3



BEM

Current resource usage 3/3



4

Conclusion & Future Work

Conclusions

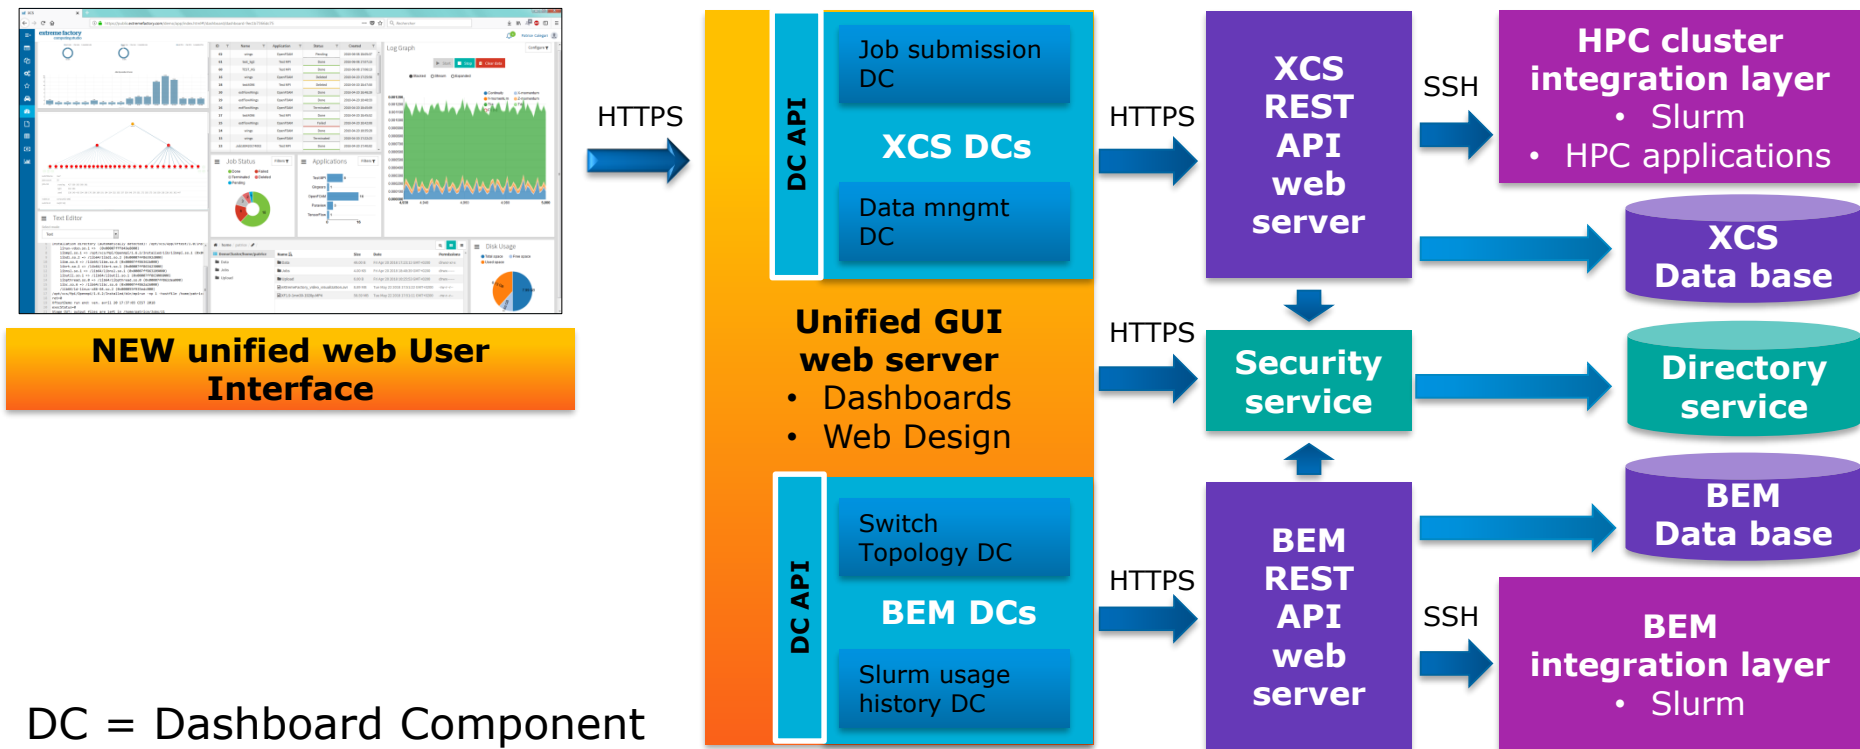
- ▶ XCS is successfully used in production on many sites for several years and it evolves continuously
- ▶ BEM is still under development and the first Minimal Viable Product (MVP) is very promising
- ▶ Mobile devices are becoming a new standard way for doing “everything”, so such a web portal approach 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.)

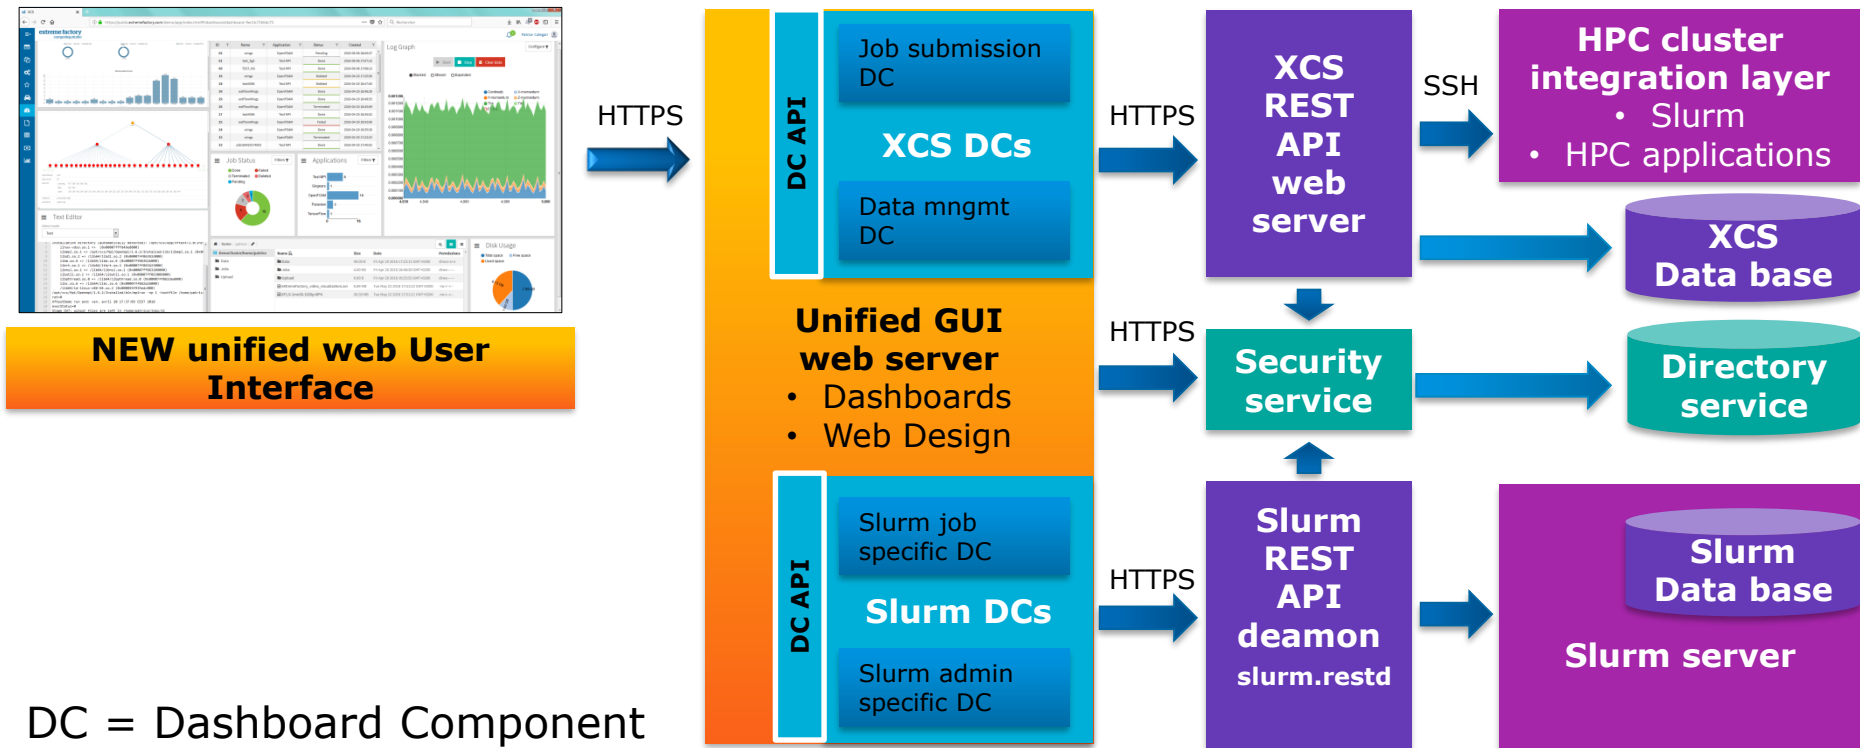
XCS and BEM architecture

Complete solution to be developed in 2020



XCS and Slurm native REST service architecture

Possible evolution...



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, the Atos logo, Atos Syntel, Unify, and Worldline are registered trademarks of the Atos group. May 2019. © 2019 Atos. Confidential information owned by Atos, to be used by the recipient only. This document, or any part of it, may not be reproduced, copied, circulated and/or distributed nor quoted without prior written approval from Atos.

The Atos logo is displayed in a white, bold, sans-serif font. The letter 'o' is stylized with a white circle inside it. The logo is positioned in the bottom right corner of the slide.

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>