science + computing

A Bull Group Company

Optimization of optimization runs, or, how do I get my results faster?

2014-11-07

Jan Wender

HPC Software and Services

Who I Am

LinkedIn

Jan Wender

HPC Services and Software, science + computing ag: Helping you to run your cluster most efficiently

Stuttgart Area, Germany | Information Technology and Services

Current	**HPC Services and Software, Team Leader** at **science + computing ag** **Senior IT Consultant** at **science + computing ag**
Past	Assistant at EU-Project LINGUA Research Assistant at Universität Trier Student assistant at Universität Trier see all
Education	Universität Trier Technische Universität Carolo-Wilhelmina zu Braunschweig

Many thanks to FAU Erlangen for offering
the RRZE Icon Set with CC-BY-SA 3.0 license
https://creativecommons.org/licenses/by-sa/3.0/de/

science + computing

A Bull Group Company

How to Scale?

http://www.flickr.com/photos/xjrlokix/4379281690/

http://www.flickr.com/photos/didbygraham/1217435486/

Scale Up

Bigger is Better!

Scale Out

If one is not enough, take two

When Scale-Out goes wrong...

http://www.flickr.com/photos/joiseyshowaa/2402764792/

Workstation vs. Cluster

High-End-Workstation

- 1 Node, 2 CPUs
- per CPU
  - 4 – 8 (18) Cores
  - Freq. up to 3.5 GHz (2.3 GHz)
- 8 – 16 (36) Cores
- bis 512 GB RAM

Cluster

- Several Nodes (z.B. 5)
- per CPU
  - 10 or 12 Cores
  - Freq. up to 2.6 GHz
- 100 – 120 Cores
- bis 1280/2560 GB RAM

Structure of a HPC Cluster

Small Cluster

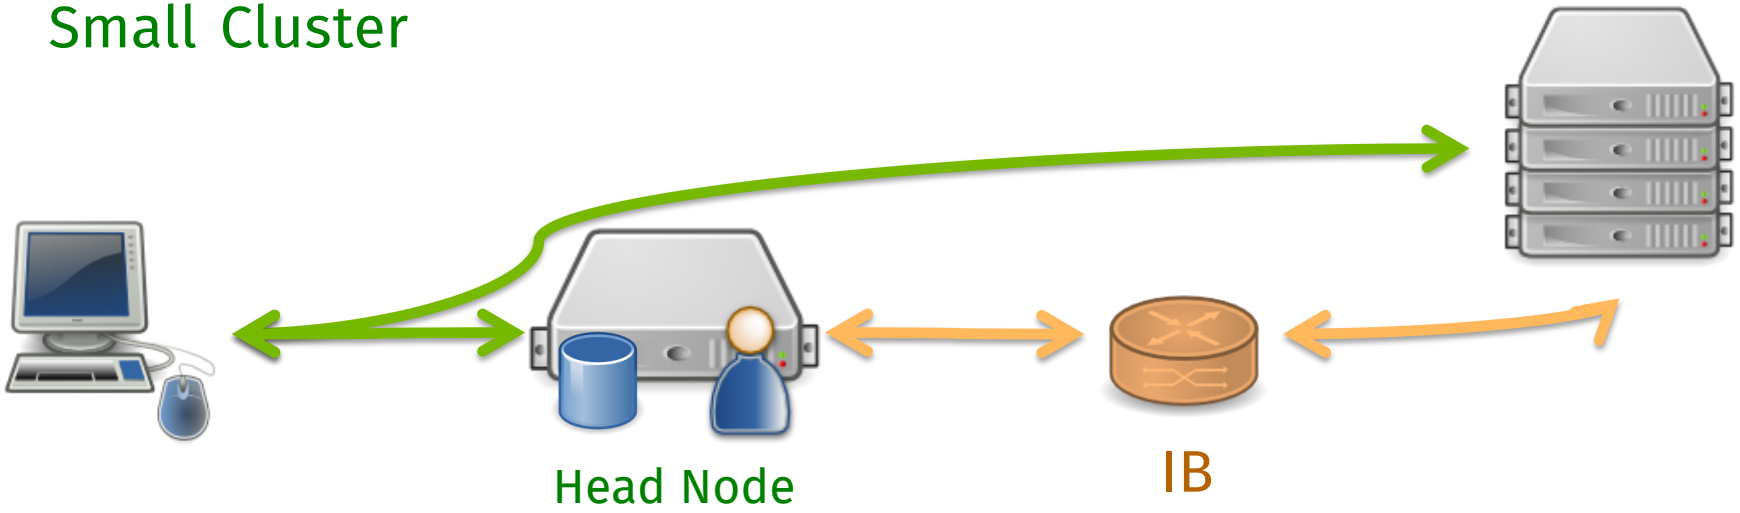

- Small sites
- ~5 Compute Nodes
- Headnode including Storage, also Login Node

Structure of a HPC Cluster

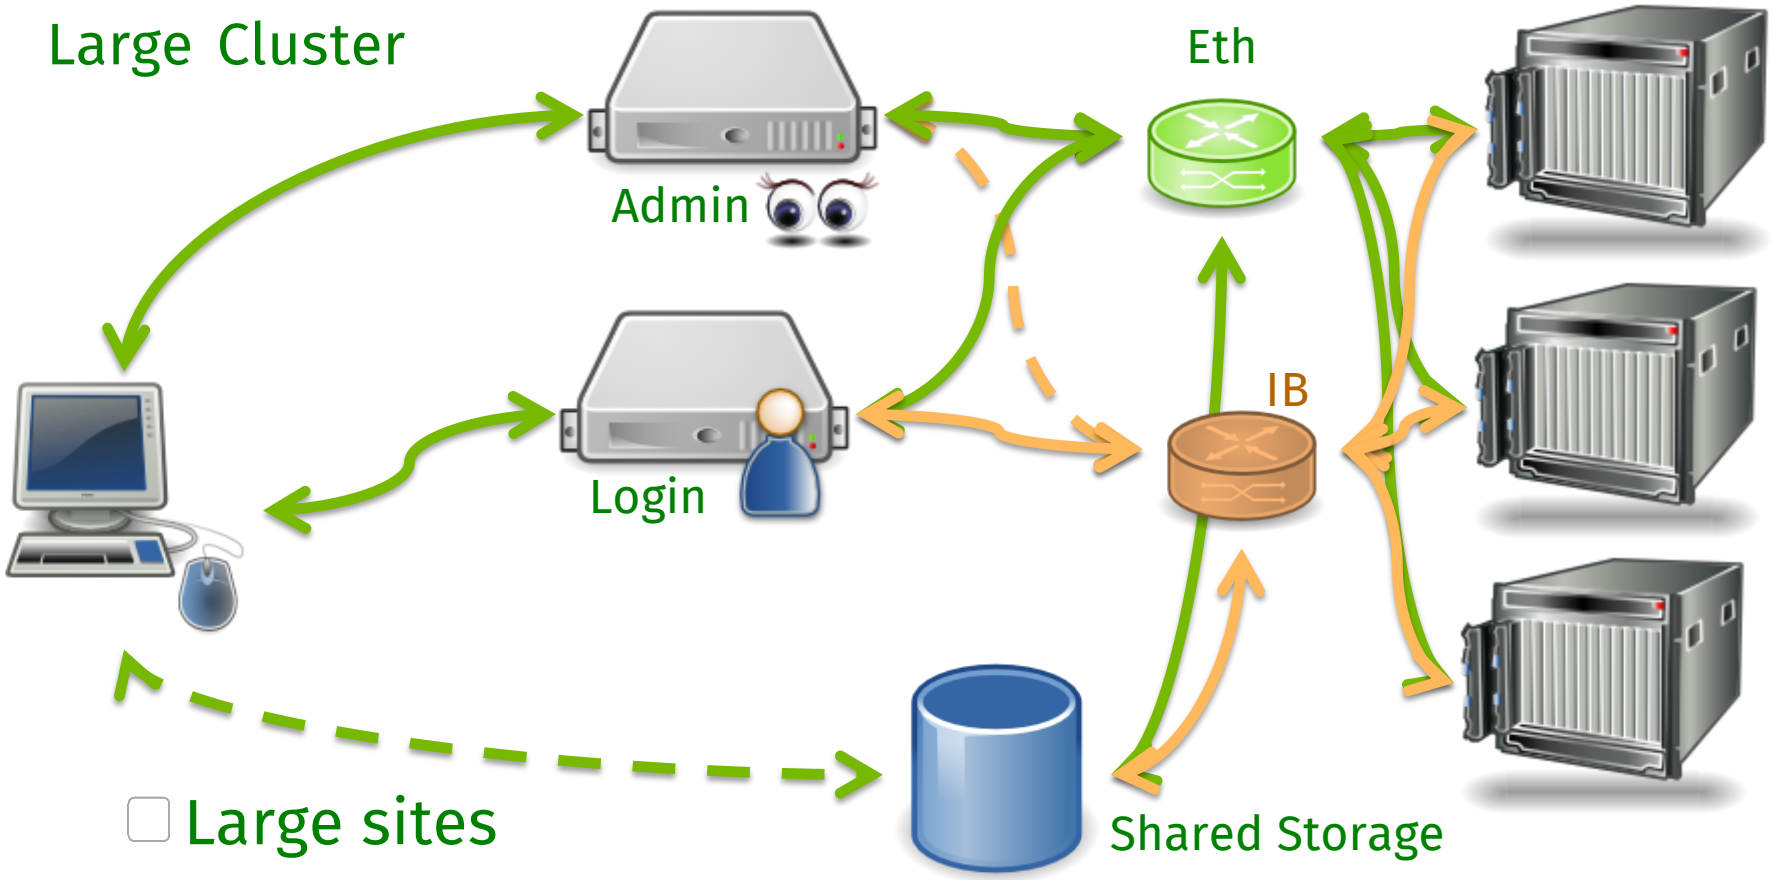

- Large sites
  - Thousands of Nodes
- Dedicated Nodes
- Dedicated Storage

Energy Consumption

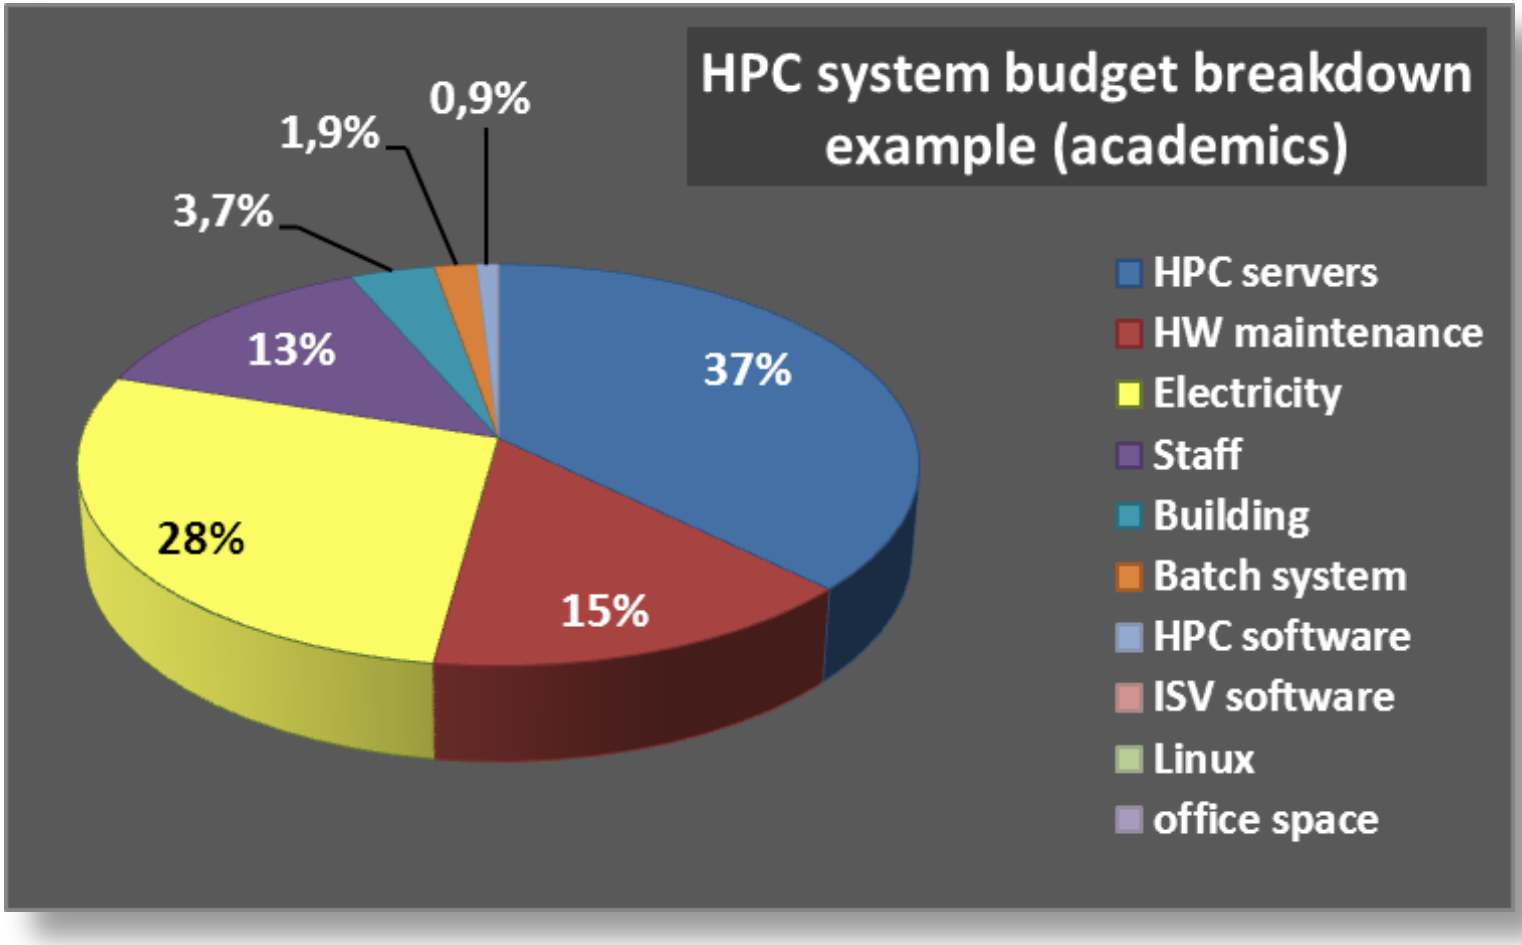

Cooling & Power Usage Effectiveness

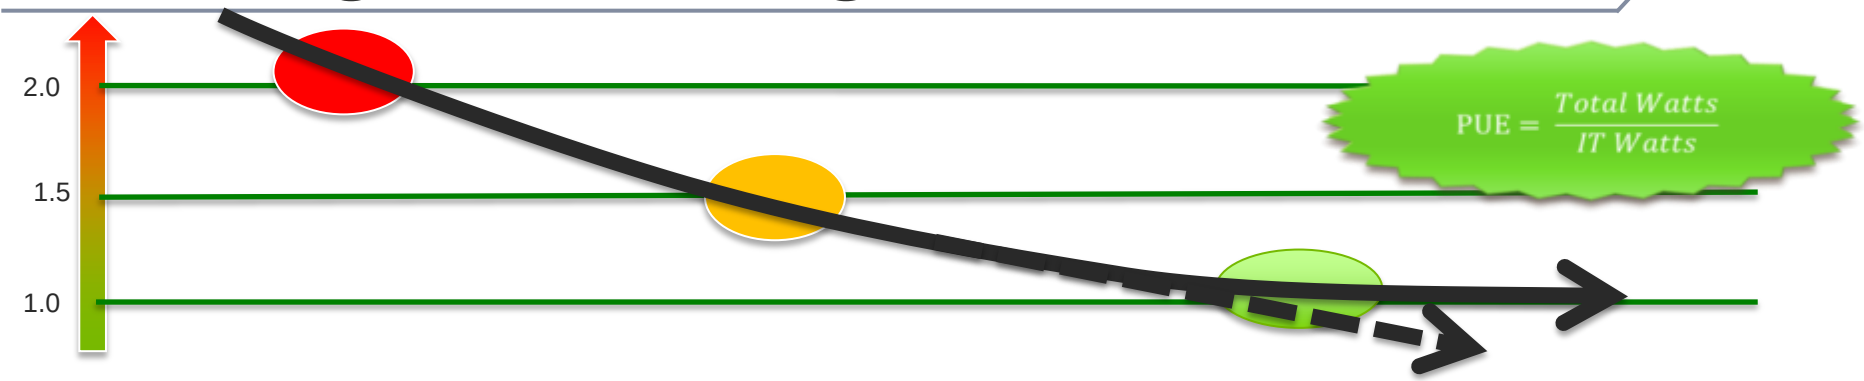

Air-cooled
10(-20) kW/rack
Room 20° C
A/C water 7-12° C

PUE ≥ 1.7

Water-cooled doors
40 kW/rack
Room 20°C
Water 7-12°C

PUE ~ 1.4

Direct-Liquid-cooling
80 kW/rack
Room up to 30°C
Water Up to 40 °C

PUE < 1.1

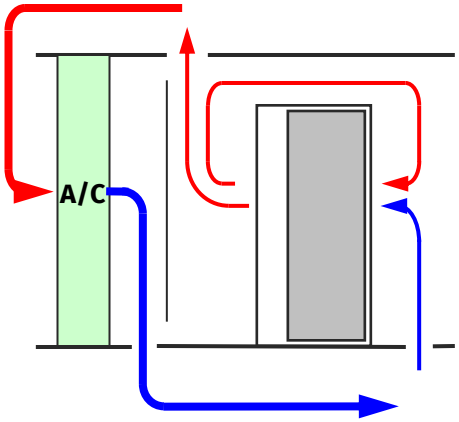

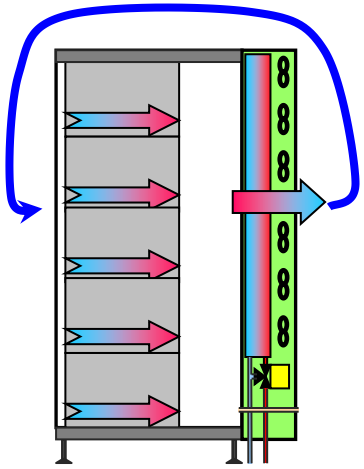

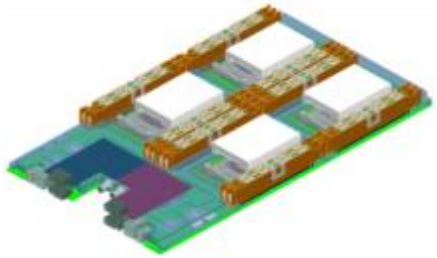

Co-generation

optiSLang Batch-Jobs

Access to the Cluster

- Generic Interfaces
- optiSLang Node does not need to be in the cluster
- Connection into the cluster via ssh

optiSLang node with cluster access

- osl client in optiSLang integated

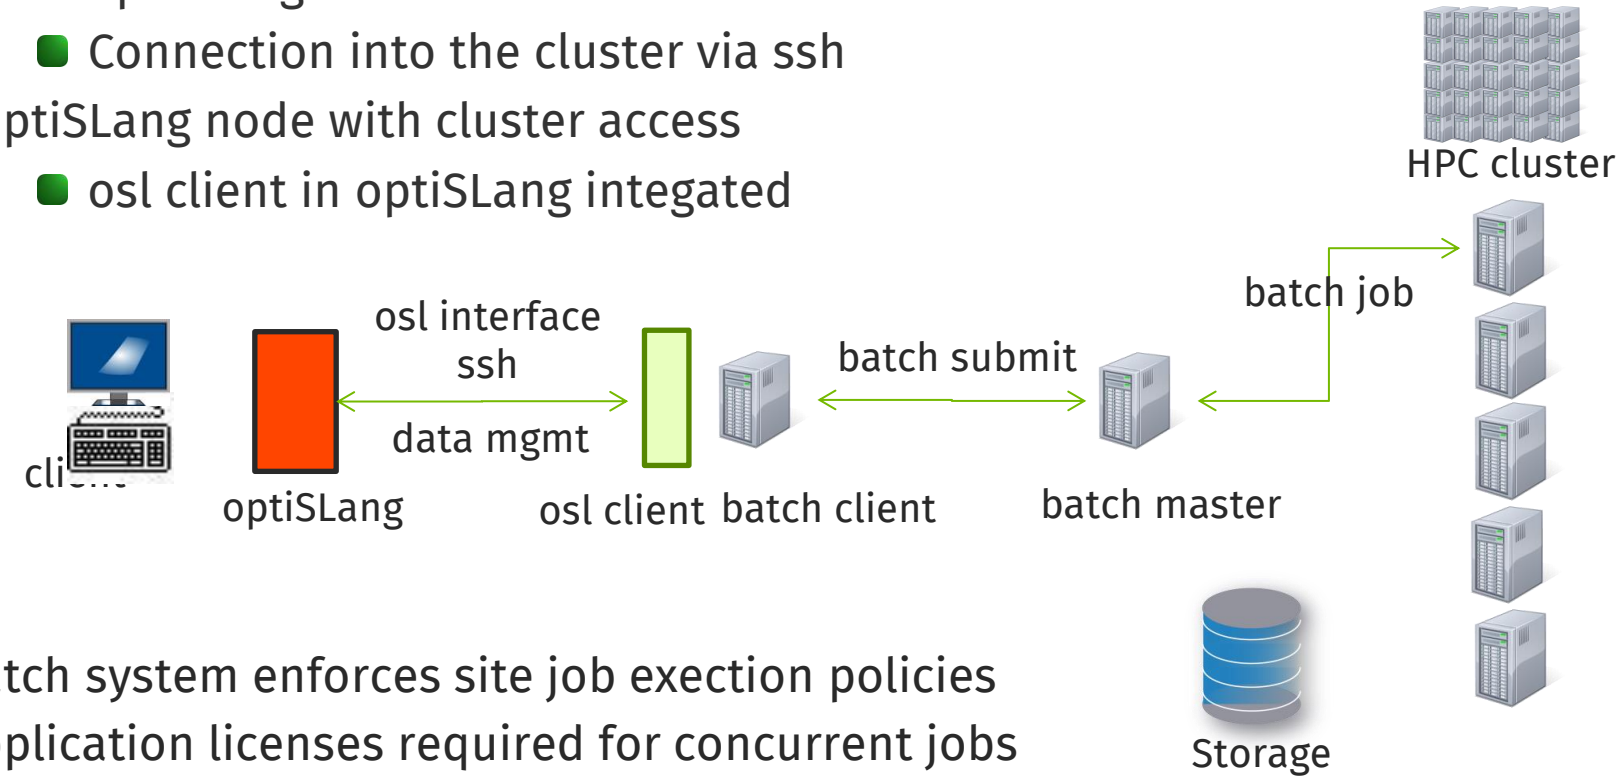

Batch system enforces site job exection policies

Application licenses required for concurrent jobs

Batch Policies

Possible configurations for prioritization of optiSLang jobs

- Examples LSF:
  - SLAs: allow a number of optiSLang jobs to run concurrently when there are jobs available
  - Guarantee a number of slots for optiSLang jobs
  - Preempt lower priority jobs for high priority optiSLang jobs
  - Give equal or higher priority to jobs of different users

```
Begin ServiceClass
NAME       = optiSLangJob
PRIORITY   = 10
GOALS      = [velocity 11 timeWindow (1:9:00-5:17:00)]
End ServiceClass

Begin ServiceClass
NAME           = slaOptimizeLong
GOALS          = [ guarantee ]
ACCESS_CONTROL = QUEUES[ long ]
AUTO_ATTACH    = y
DESCRIPTION    = Guaranteed slots
End ServiceClass
```

HPC in the Cloud

- Access via Portal
- Remote Visualisierung
- Benefits
  - Provider operated
  - High Availability
  - Flexible Usage
  - Pay per use
- Challenges
  - Data Transfer
  - Application Licensing
  - Security
- Utilization?

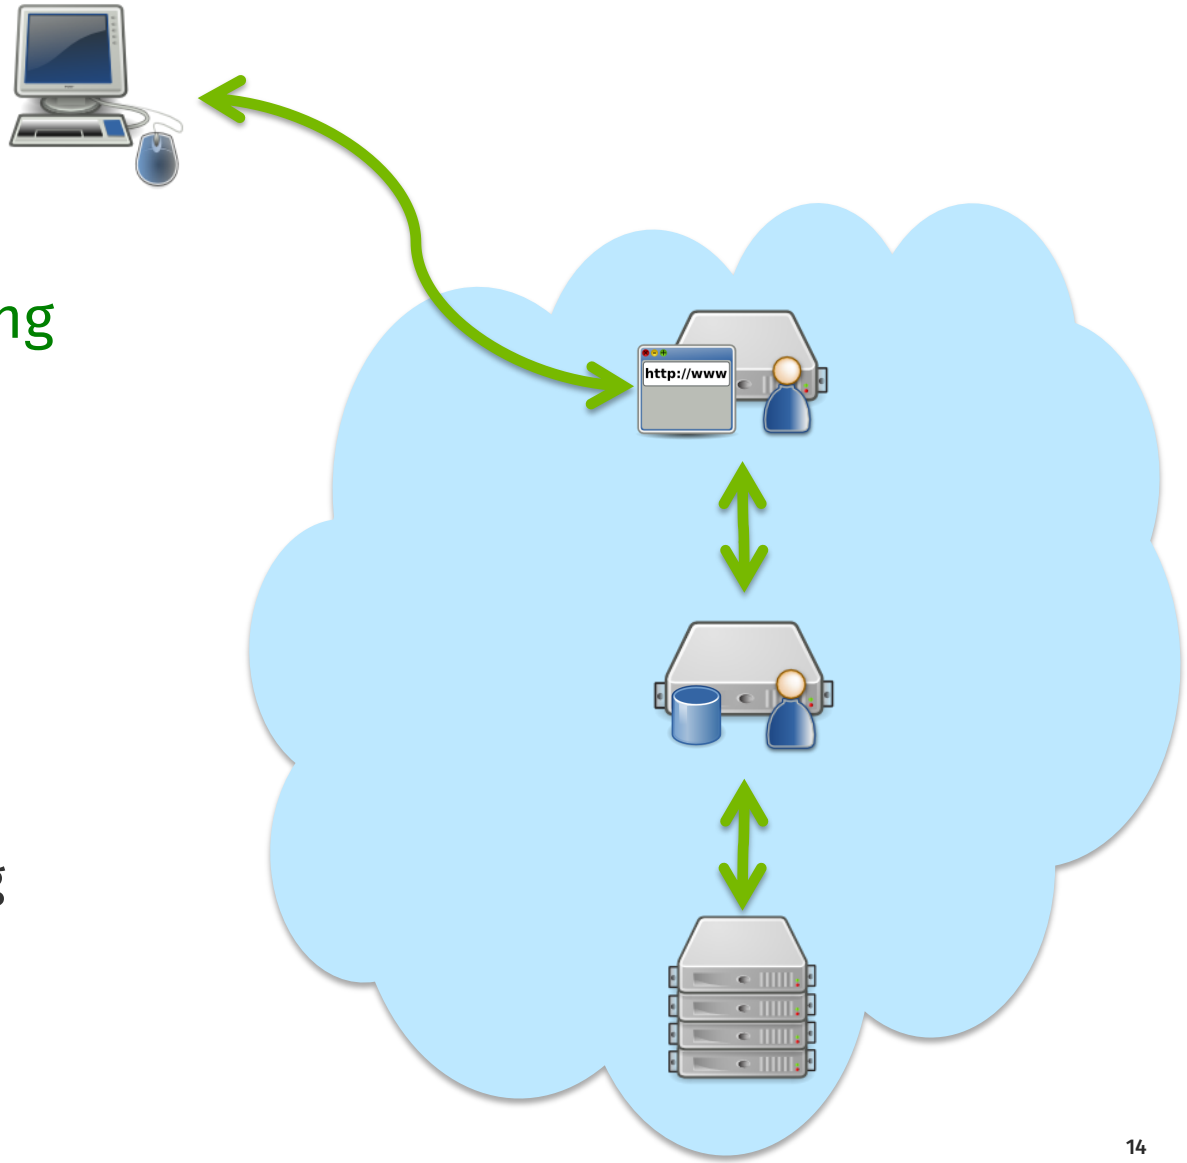

Bull's HPC on Demand offer

extreme factory

- performing compute infrastructure based on Bullx solutions
- Setup and operated by Bull HPC experts
- High level of service with total security
- Web portal access
- A quick and easy response to your innovation needs
- No requirement for heavy investment

To compute smart !

An operational scientific ‘cloud’

Portal

Security

Project management

Data management

Batch scheduling

2D / 3D Visualization

Numerical simulation applications commercial / open source

HPC Middleware : schedulers, clustered file systems, parallel mathematical libraries, compilers, remote visualization

Users

extreme factory
stay lean: compute smart

BULL
an atos company

Remote pre-/post-processing

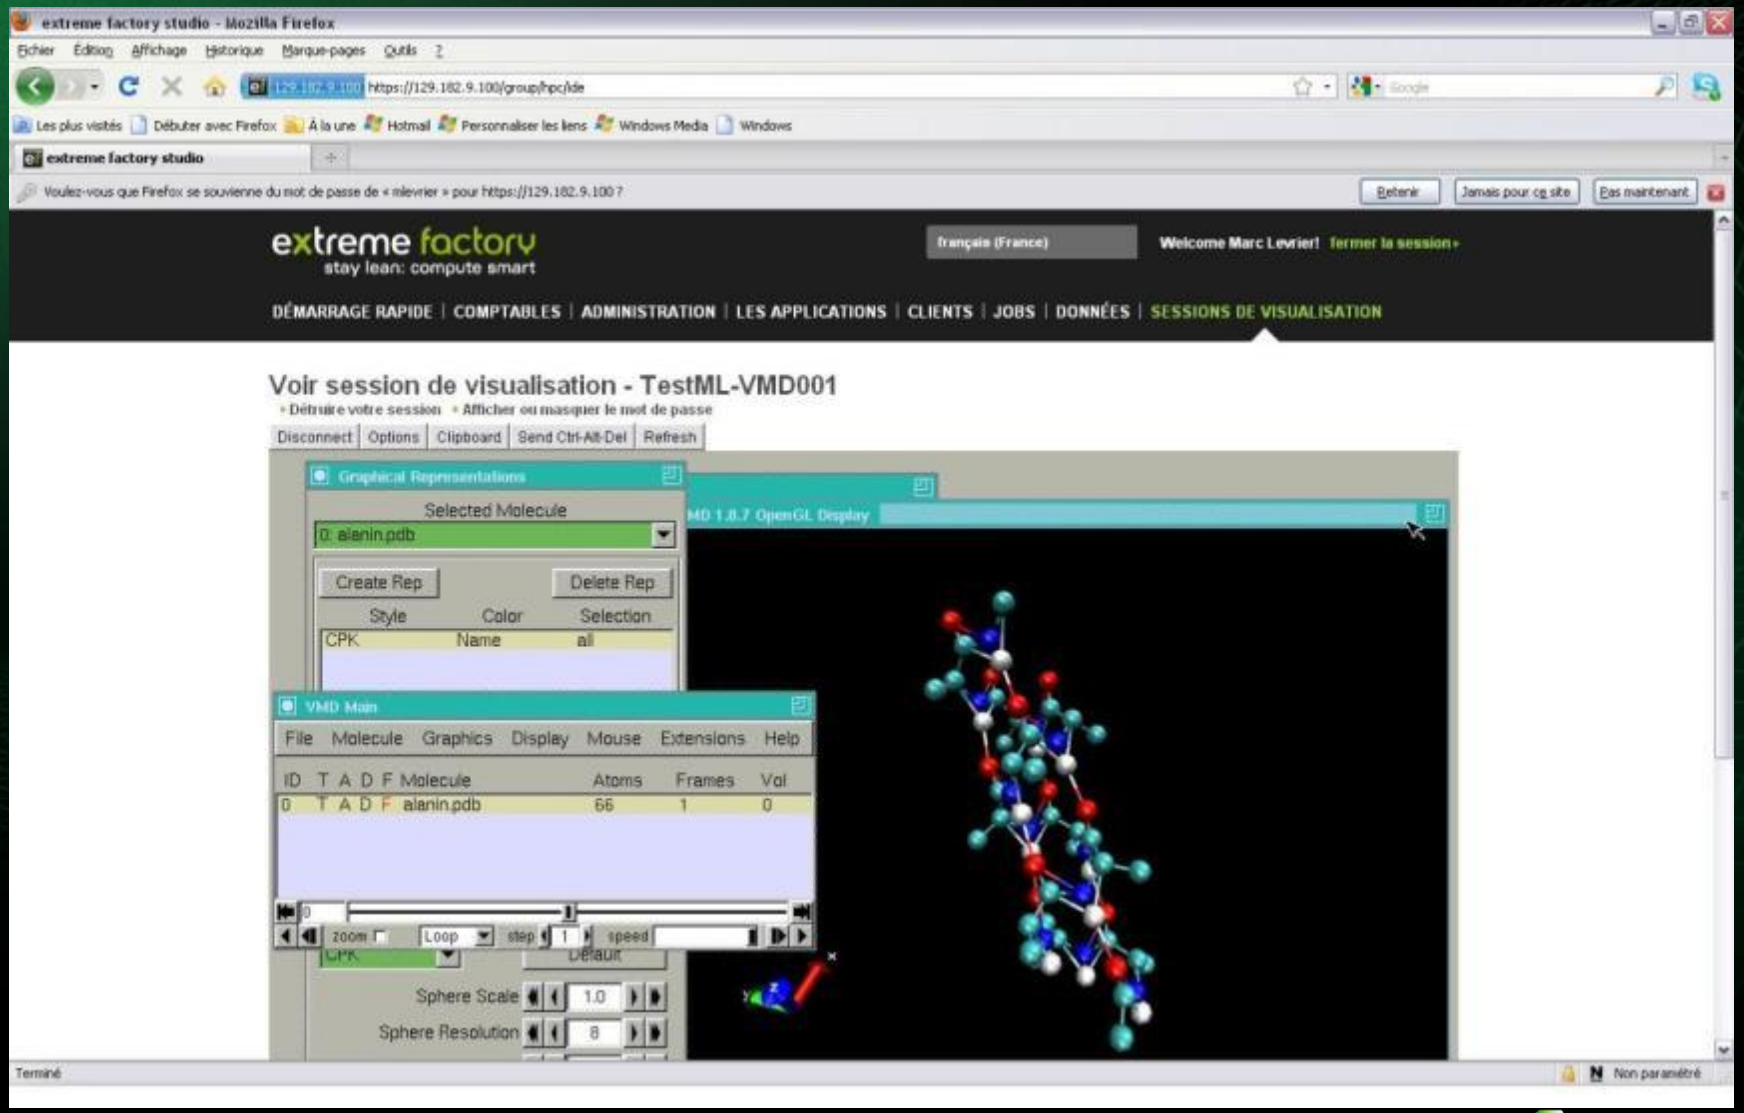

Questions?

Jan Wender
j.wender@science-computing.de
+49 7071 9457257

science + computing
A Bull Group Company

science + computing

A Bull Group Company