JÜLICH
FORSCHUNGSZENTRUM

# Advanced System Monitoring in PTP

November 19, 2013 | Wolfgang Frings and Carsten Karbach

# Content


1. PTP System Monitoring – status quo
2. Short term enhancements
3. Long term enhancements

# Part I: System Monitoring – status quo

November 19, 2013 | Wolfgang Frings and Carsten Karbach

JÜLICH FORSCHUNGSZENTRUM

# PTP Monitoring Scope

- job and system monitoring of **large-scale** supercomputers
- examples for large-scale systems monitored with PTP:
  - BG/Q system JUQUEEN (JSC), LoadLeveler, 458K cores
  - Cray XT Jaguar (ORNL), TORQUE+ALPS, 299K cores
  - Kraken (NICS), Moab+PBS, 112K cores
  - Yellowstone (NCAR), LSF, 72K cores
- monitoring of **multiple target systems** in one perspective
- support for **many batch systems** (Grid Engine, LoadLeveler, Open MPI, PBS, Slurm, Torque, LSF)
- overview of the system on a **single screen**
- uniform interface to supercomputers
- based on monitoring application **LLview**

JÜLICH FORSCHUNGSZENTRUM

# PTP Monitoring Perspective

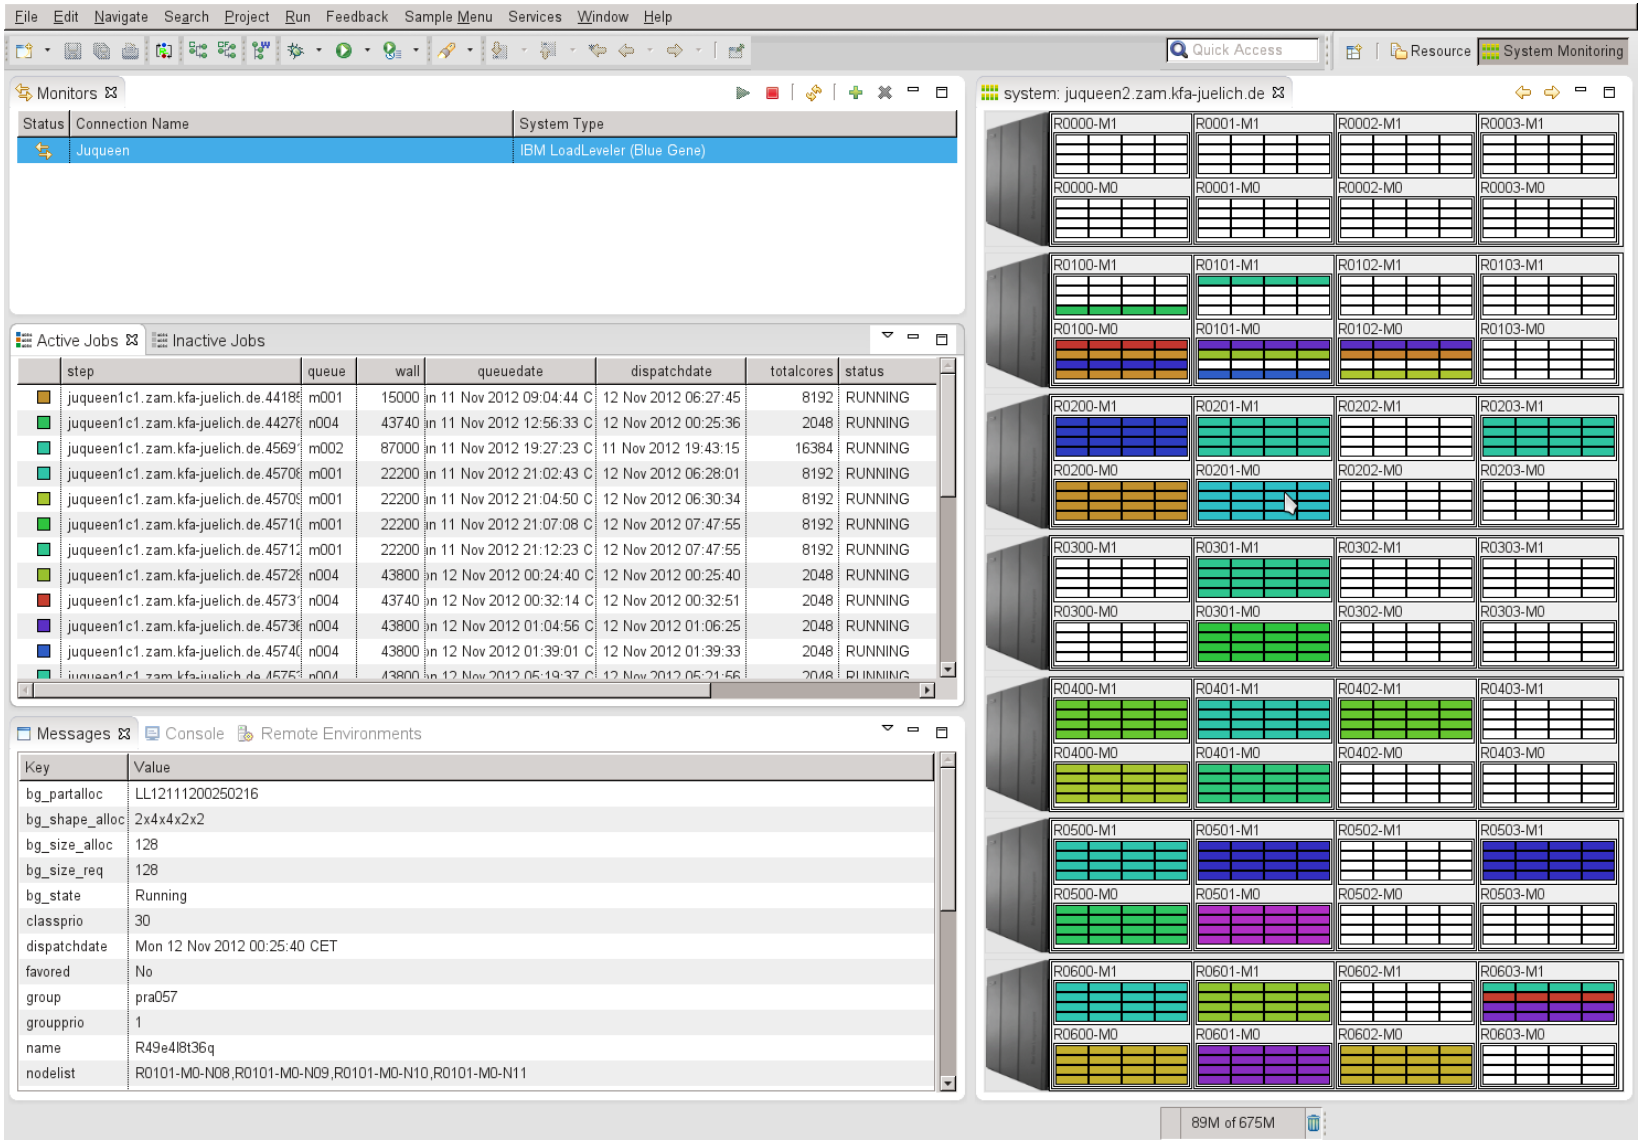

JÜLICH FORSCHUNGSZENTRUM

# Monitoring Views

- **Nodes View** renders target system architecture, maps jobs to compute resources
- **Active Jobs View** lists running jobs
- **Inactive Jobs View** lists queued jobs
- **Monitoring View** selects active target system, starts/stops monitoring
- **Message View** shows message of the day

# User interaction

- **job management**: cancel, get output/error
- **filtering**: show only user jobs, flexible filtering dialog
- linking information: click on job → highlight its nodes in Nodes View, show detail information in Message View
- change level of detail

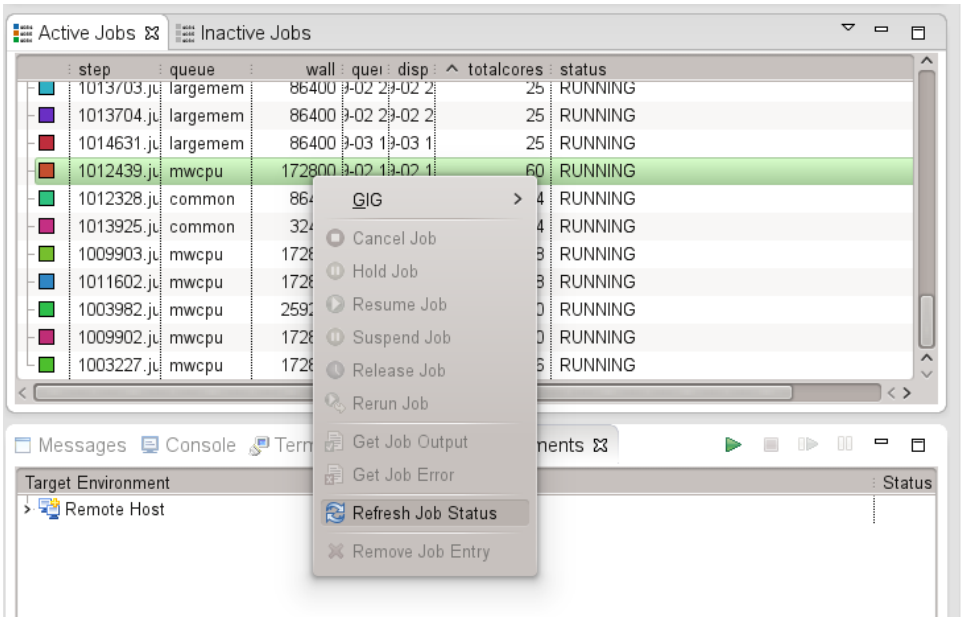


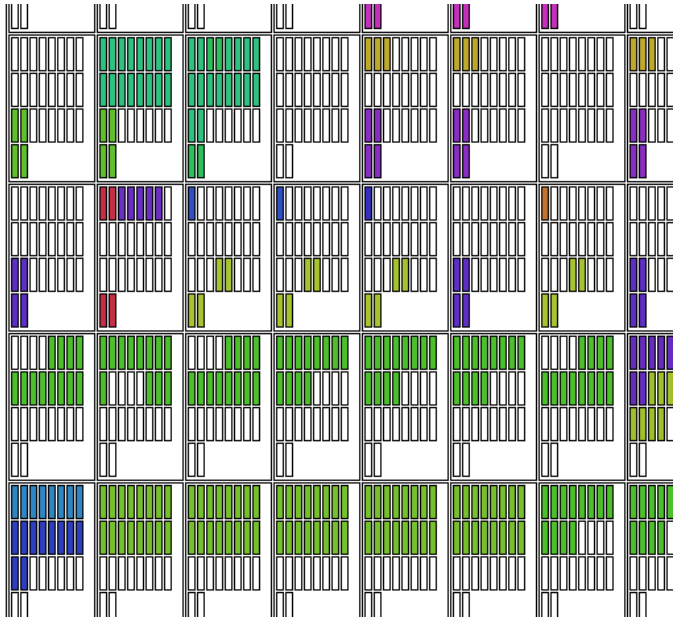

# Part II: Short term enhancements

November 19, 2013 | Wolfgang Frings and Carsten Karbach

JÜLICH FORSCHUNGSZENTRUM

# Job selection

- raised in bug 403060
- allow selection of **multiple** jobs
- keep selected job selected until it is de-selected
- mark entire **connected area** of each job

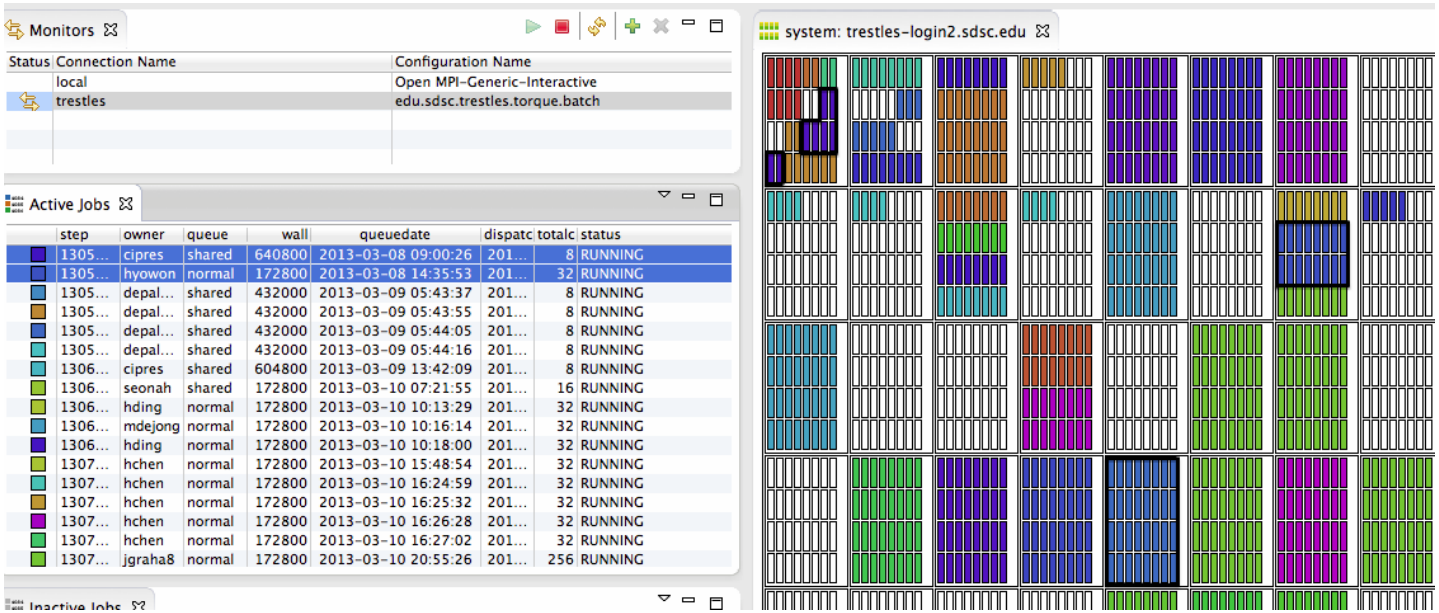


Source: `https://bugs.eclipse.org/bugs/attachment.cgi?id=228316`

JÜLICH FORSCHUNGSZENTRUM

# Improved job localization

- **adjustable** minimum rectangle size
- currently set to 7 px
- ensure, that rectangle width/height are at least 7 px large, if possible take more space

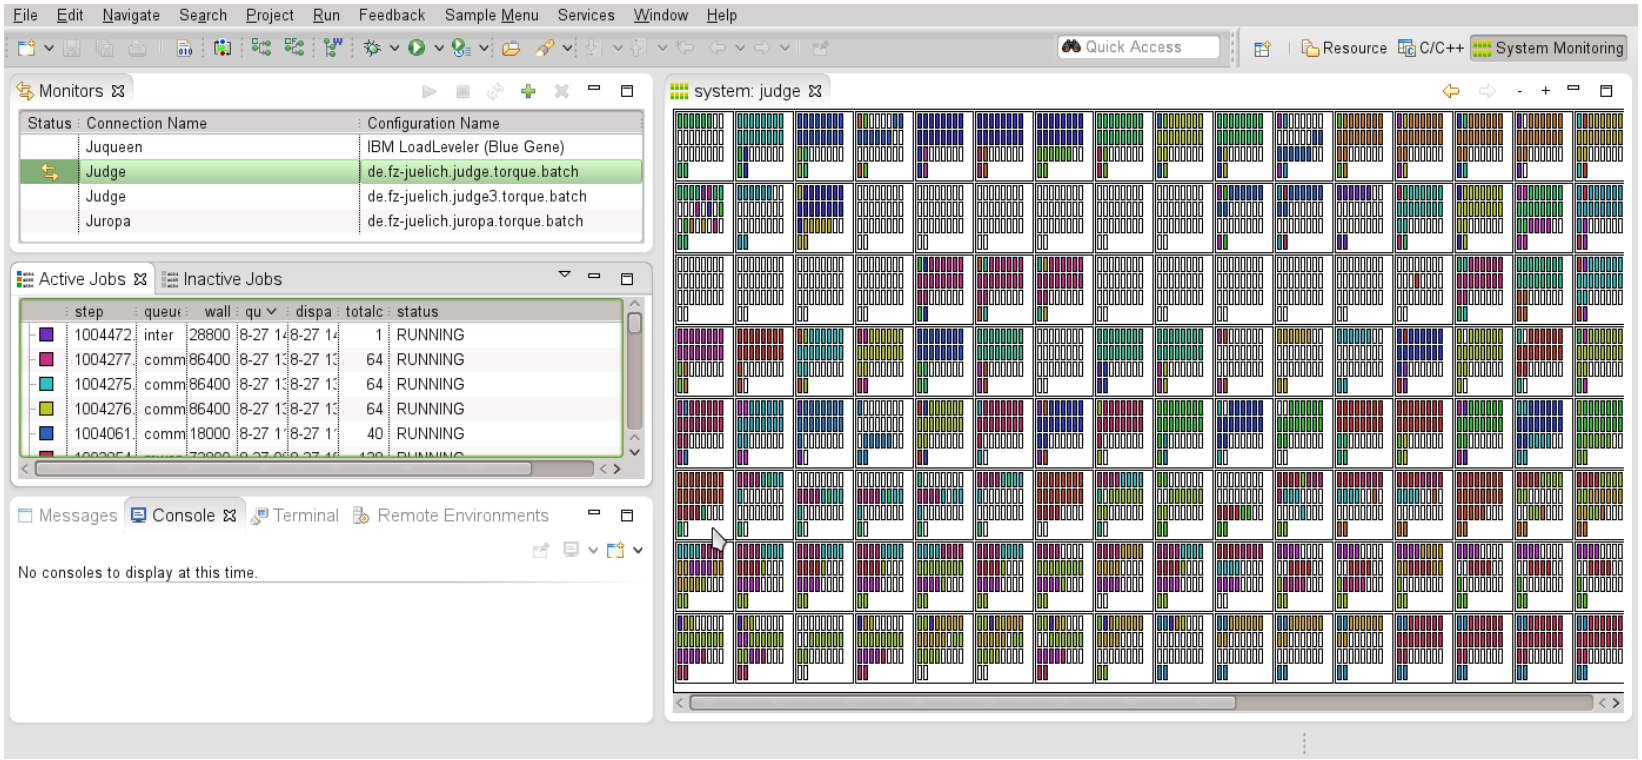

JÜLICH FORSCHUNGSZENTRUM

# Improved job localization

- **adjustable** minimum rectangle size
- currently set to 7 px
- ensure, that rectangle width/height are at least 7 px large, if possible take more space

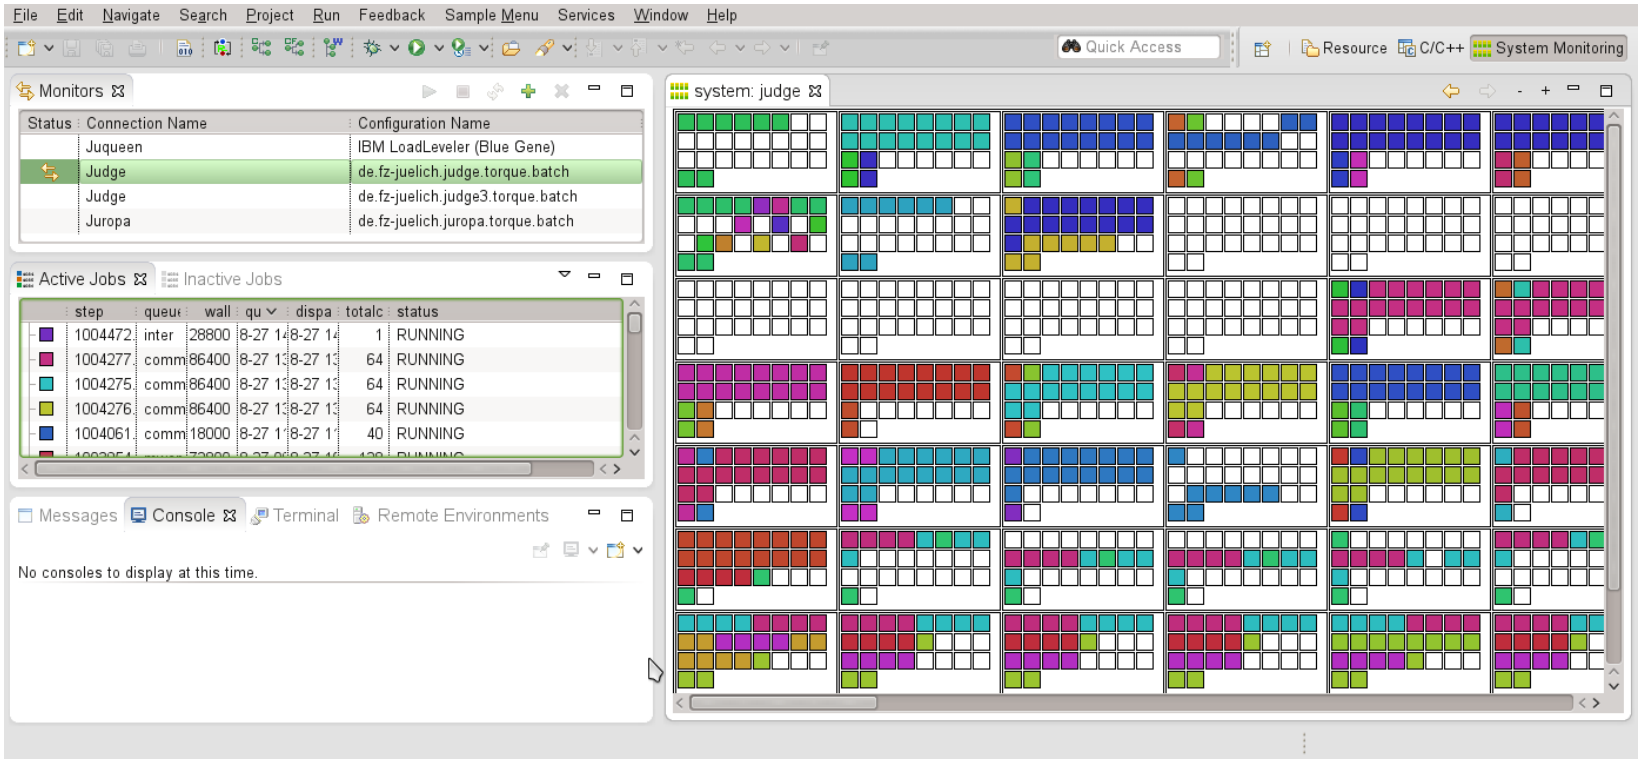

# Improved job localization

- **adjustable** minimum rectangle size
- currently set to 7 px
- ensure, that rectangle width/height are at least 7 px large, if possible take more space

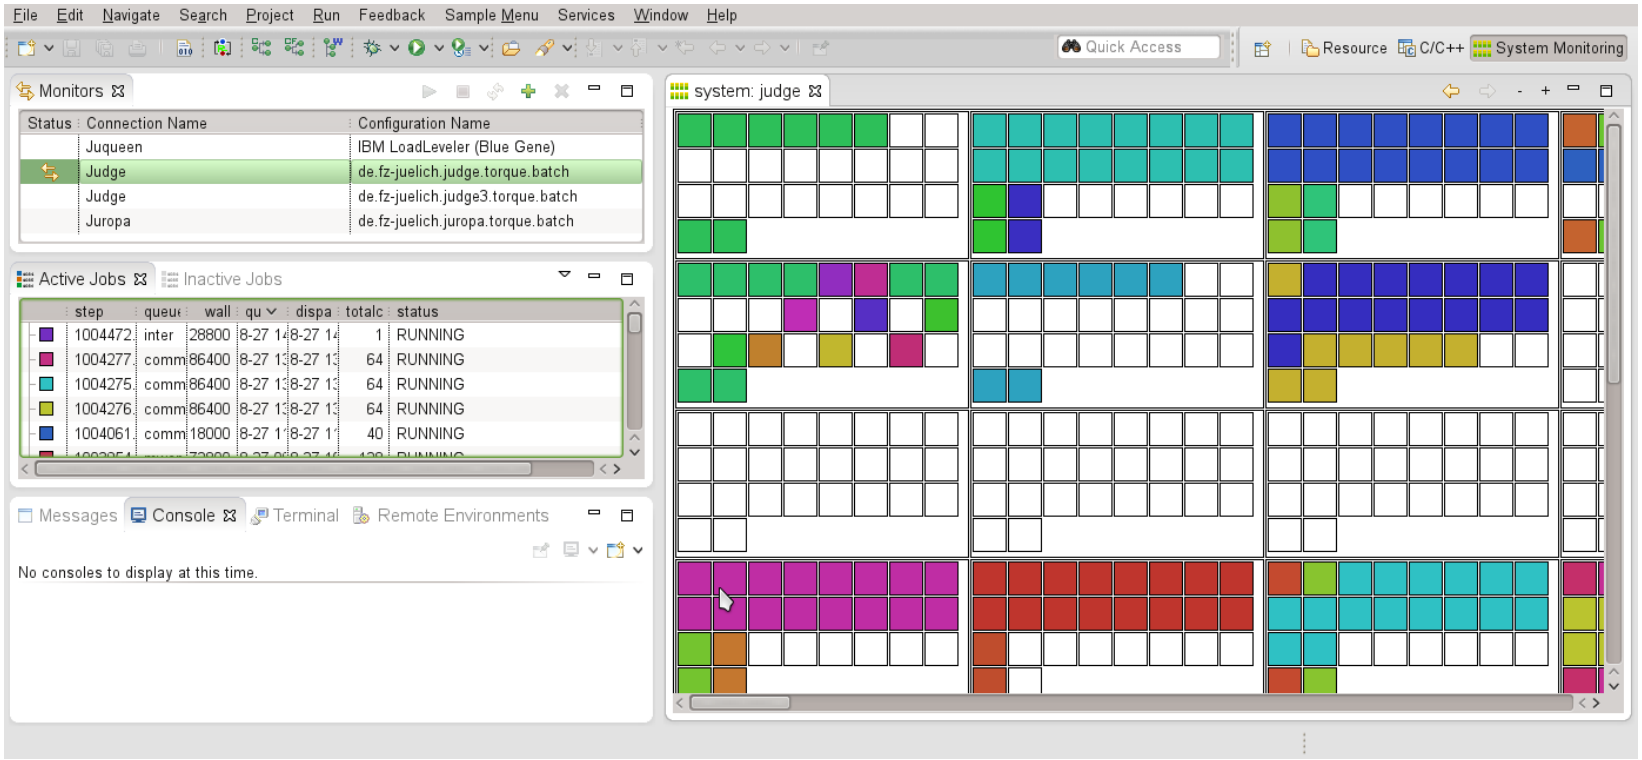

JÜLICH FORSCHUNGSZENTRUM

# Provide customized LML layouts

- get familiar with system architectures of supported target systems
- map system **topology** into LML layout
- great potential in **customized layouts**: level of detail, job filtering, showing node names

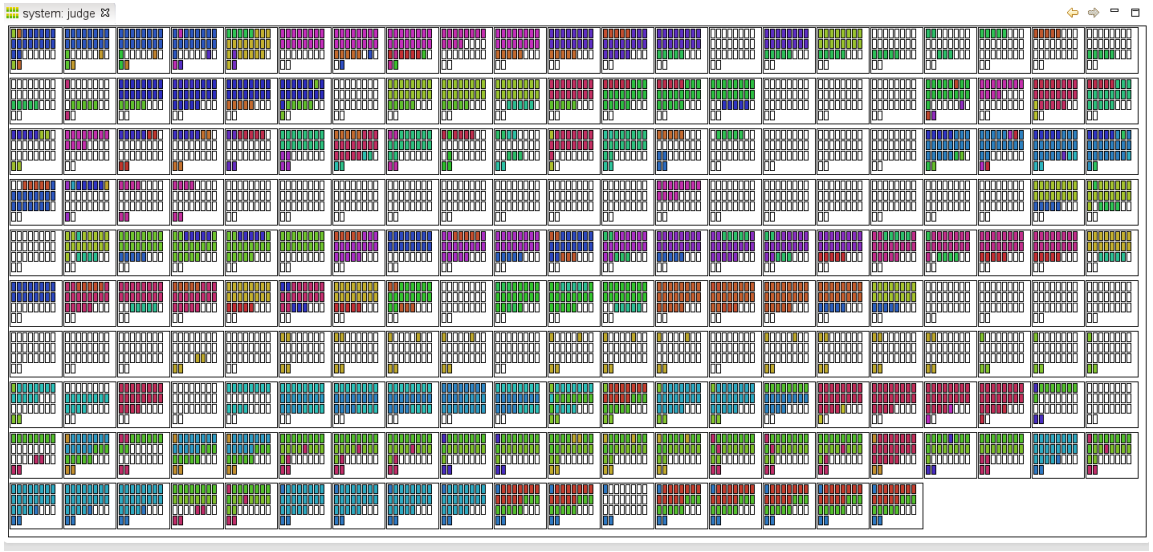


- define own grid

JÜLICH FORSCHUNGSZENTRUM

# Provide customized LML layouts

- get familiar with system architectures of supported target systems
- map system **topology** into LML layout
- great potential in **customized layouts**: level of detail, job filtering, showing node names

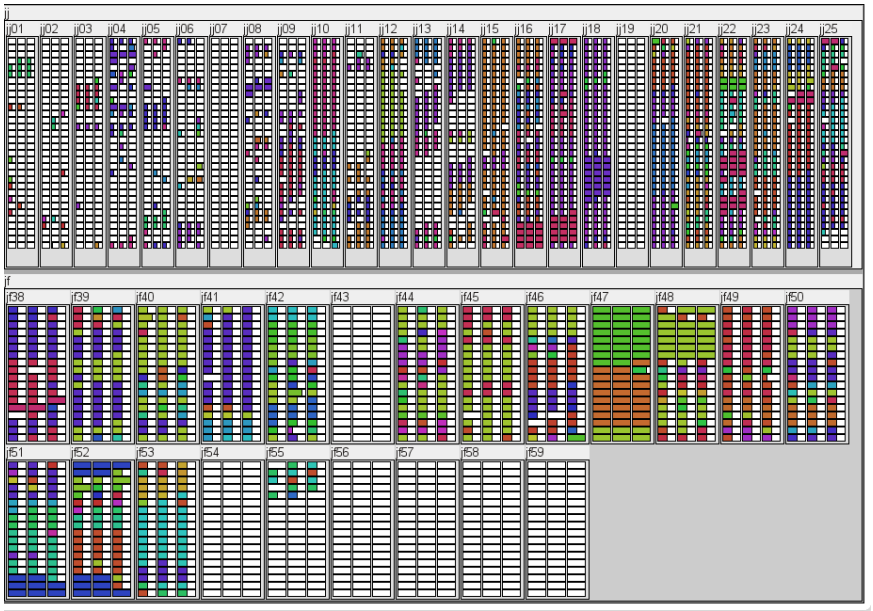

- hierarchy
- level of detail

JÜLICH FORSCHUNGSZENTRUM

# Job handling and additional job list

## Completed jobs

- jobs submitted externally disappear when completed
- if batch system does not list them, the job entries are lost
- **idea**: keep track of user's running jobs, which are removed on update

## New job list

- currently: active and inactive jobs
- better: submitted, active and completed jobs

# Part III: Long term enhancements

November 19, 2013 | Wolfgang Frings and Carsten Karbach

JÜLICH FORSCHUNGSZENTRUM

# Caching LML files

- **multiple users** on the same target system
- currently each user triggers separate LML_DA workflow
- **cache** LML file in public directory (e.g. /tmp), use LML cache as data source

## Current workflow

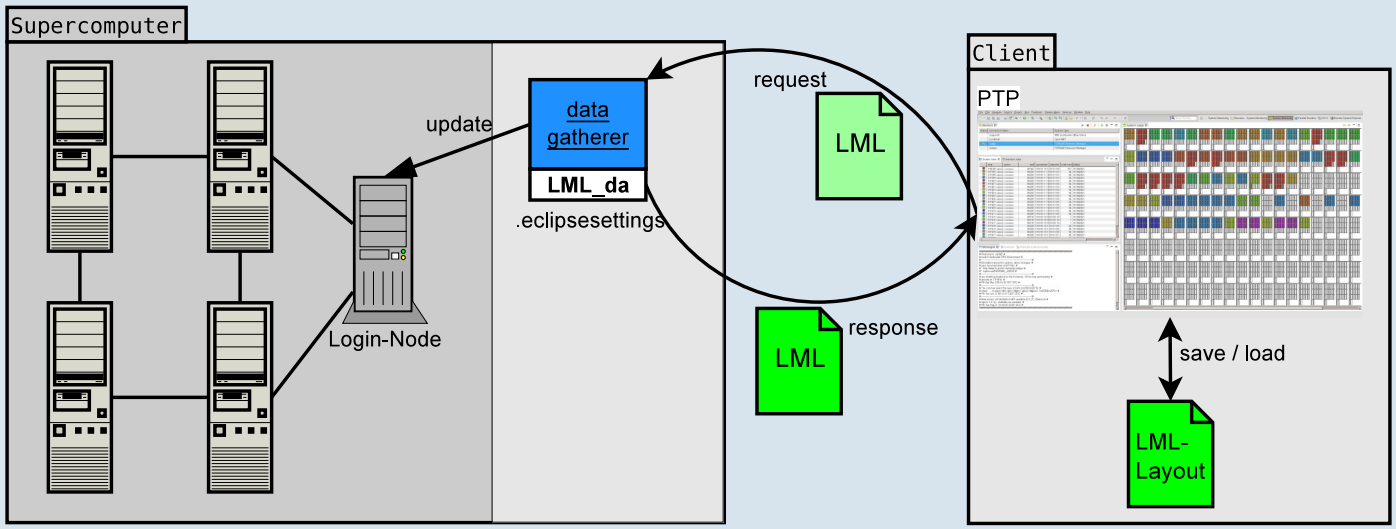

JÜLICH FORSCHUNGSZENTRUM

# Caching LML files

- **multiple users** on the same target system
- currently each user triggers separate LML_DA workflow
- **cache** LML file in public directory (e.g. /tmp), use LML cache as data source

## Cache workflow

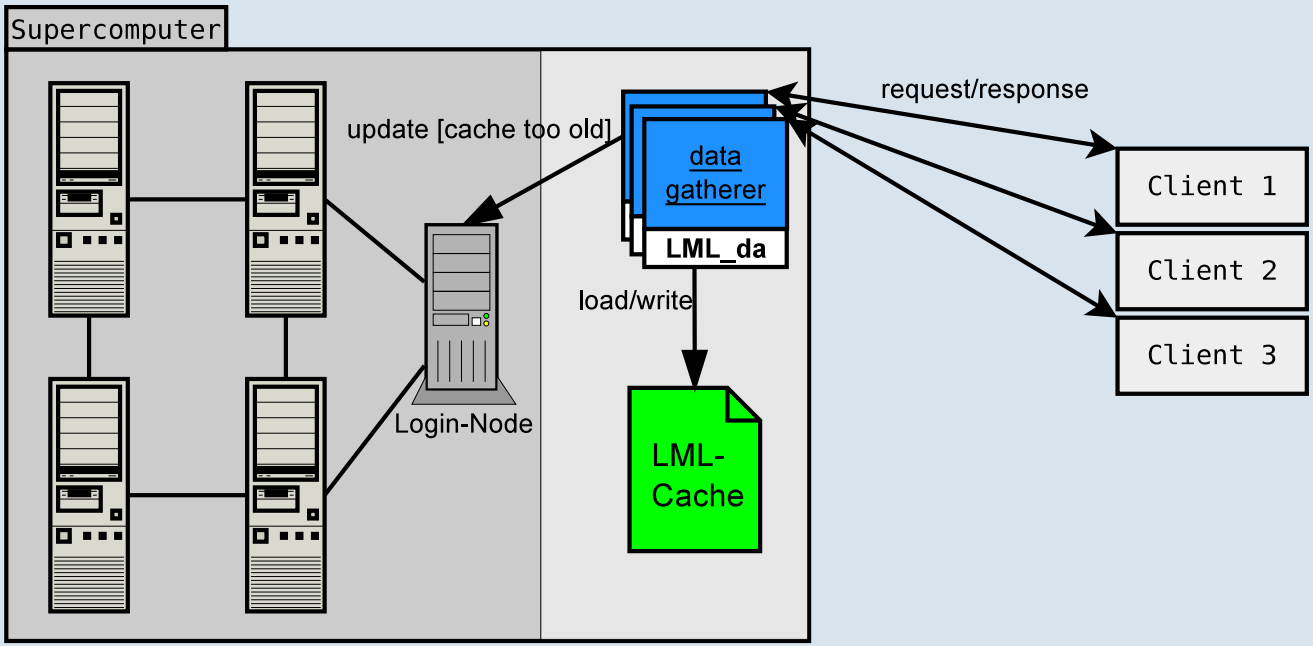

JÜLICH FORSCHUNGSZENTRUM

# New monitoring views

- derived from LLview, new monitoring types are possible
- fast overview on system statistics, history and prediction
- data description is **already included** in LML
- todo: data generation and visualization for new diagrams

JÜLICH FORSCHUNGSZENTRUM

# Histograms

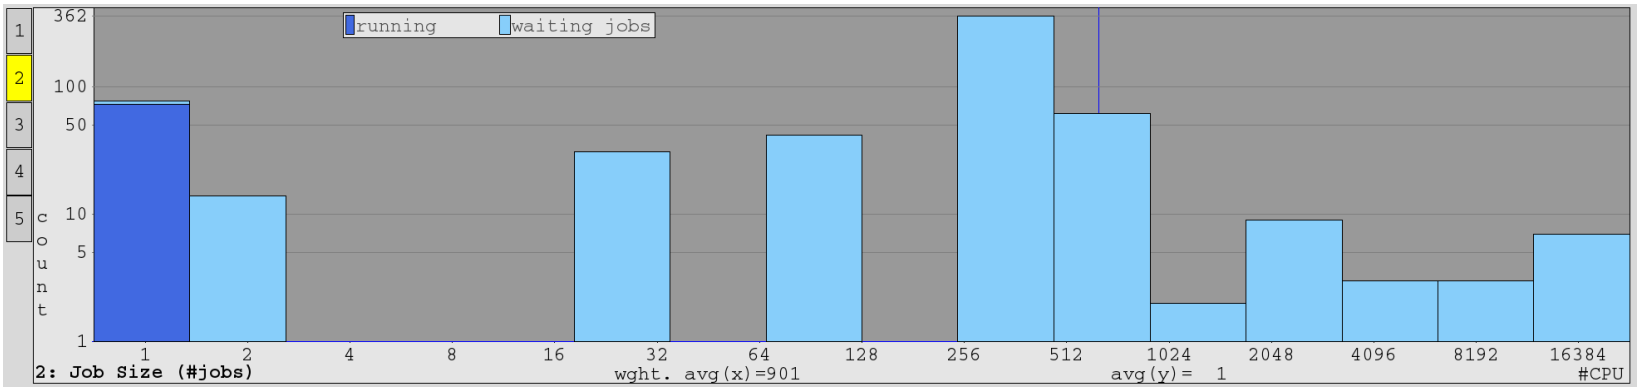


- statistical data rendered as histograms
- visualization of job parameter distribution: queue, size, waiting time

JÜLICH FORSCHUNGSZENTRUM

# Load history

- **usage history** of the target system (e.g. last 3 days)
- extendable for power/memory/accelerator usage
- requires **LML log**, switch to **stateful** server

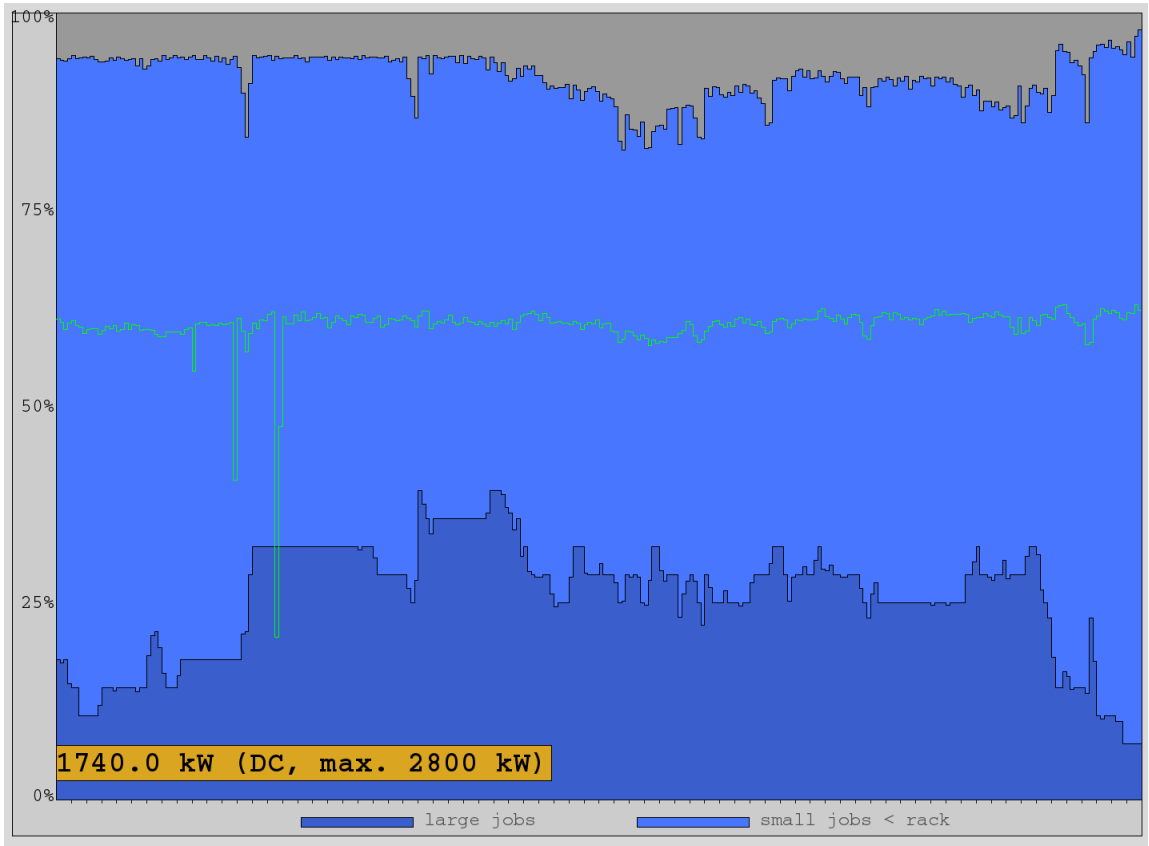

JÜLICH FORSCHUNGSZENTRUM

# Prediction diagram

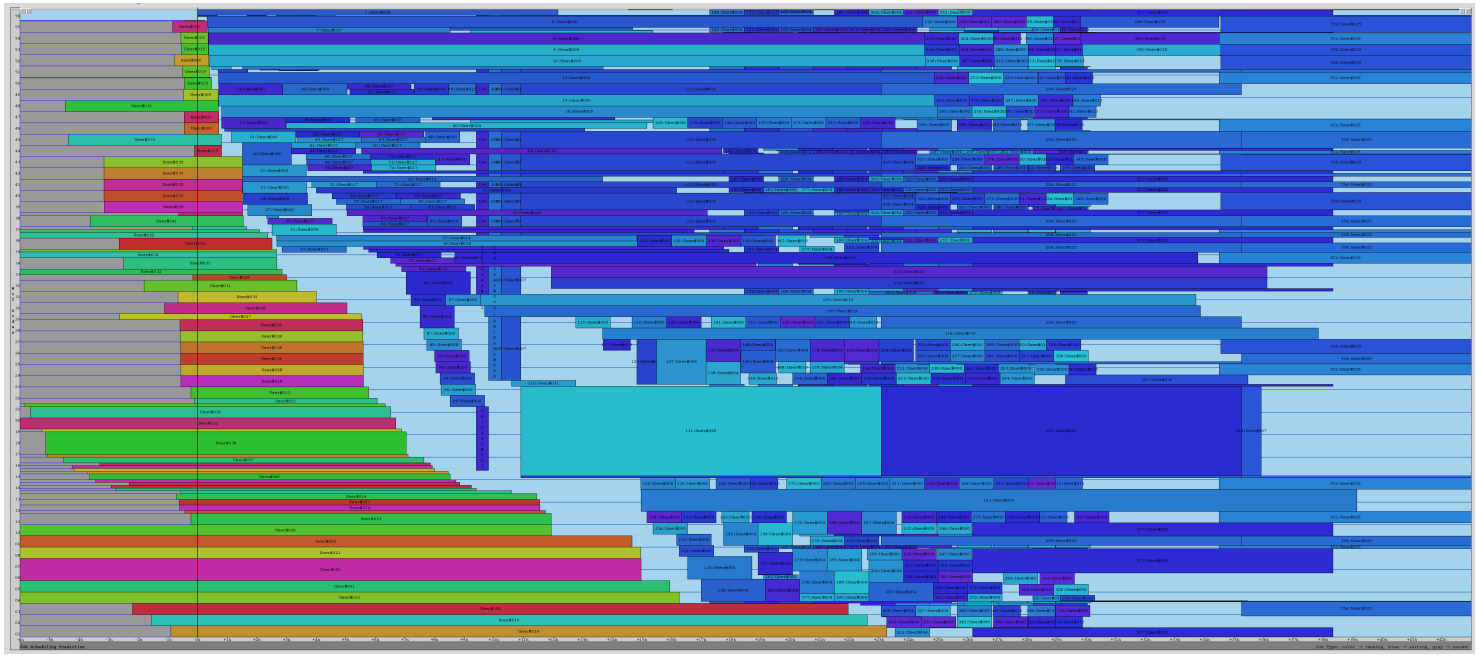

- Gantt chart visualization of **future workload**
- each rectangle represents one job,
  x-axis $\rightarrow$ time, y-axis $\rightarrow$ resources
- requires **JuFo** integration (see next slides) for simulation of future schedule

JÜLICH FORSCHUNGSZENTRUM

# JuFo – Overview

- **configurable** simulator for global job schedulers for **on-line prediction** of job dispatch dates
- based on analysis of JSC batch systems **Moab** and **Loadleveler**
- **integrated** with monitoring system **LLview**
- **LML** as configuration and communication data format
- **use-cases**:
  - **user** predicts start dates of submitted jobs
  - **administrator** simulates job scheduler performance with various input parameters, verifies scheduling rules

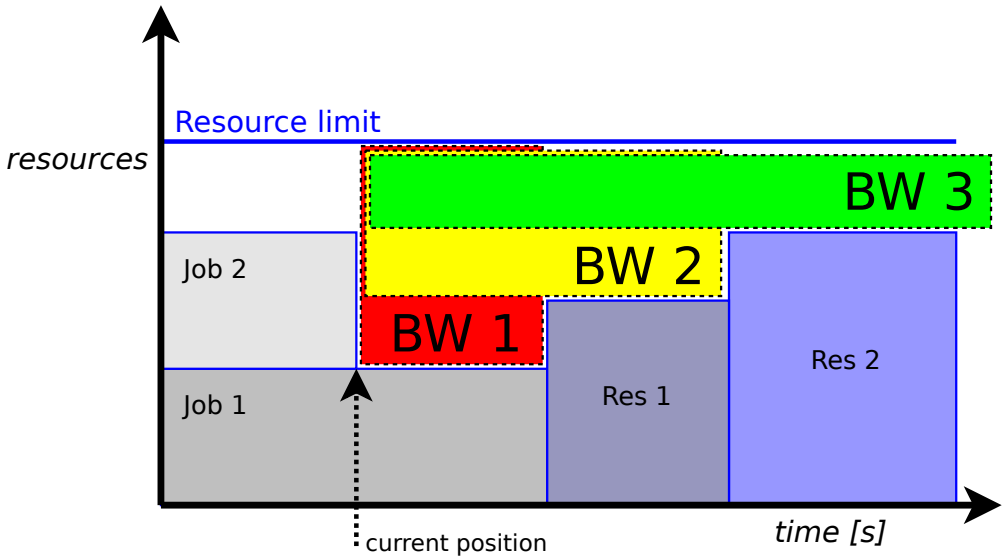

JÜLICH FORSCHUNGSZENTRUM

# JuFo integration

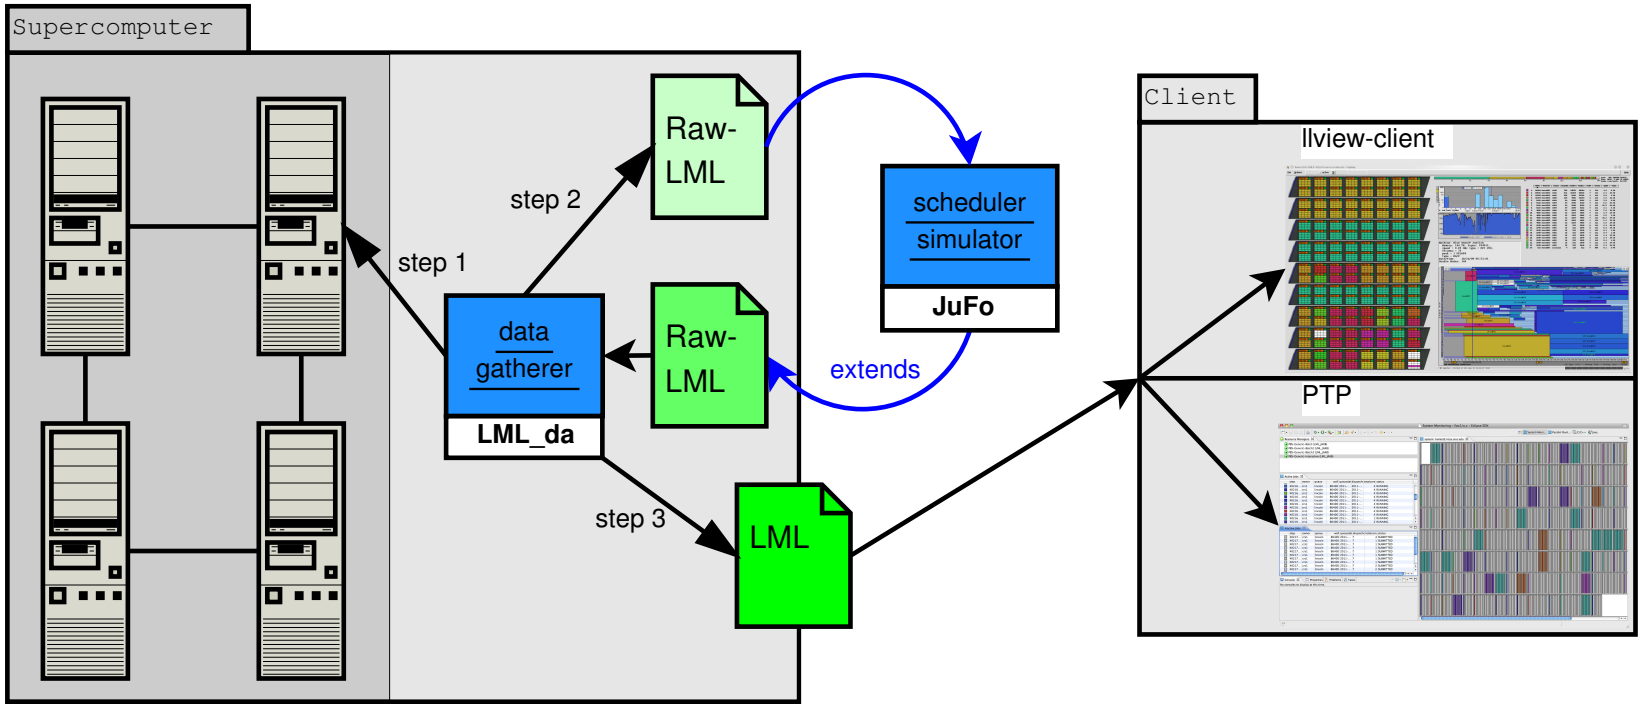


- implemented in C++, additional installation step required
- simulation duration: 1-90 seconds $\Rightarrow$ caching

JÜLICH FORSCHUNGSZENTRUM

# JuFo – Features

- supported **scheduling algorithms**
  - First-Come-First-Served
  - List-Scheduling
  - Backfilling
- available **simulation parameters**
  - generic job **prioritization**
  - advanced **reservations**
  - jobs can request CPUs, GPUs, memory
  - **nodesharing**
  - **queue** constraints
- test framework for evaluating JuFo's accuracy

# Part IV: Conclusion

November 19, 2013 | Wolfgang Frings and Carsten Karbach

JÜLICH FORSCHUNGSZENTRUM

# Conclusion

- PTP provides monitoring system for large-scale supercomputers
- monitoring views: job lists, nodes view
- short term enhancements:
  - adjust job **selection**, multiple jobs
  - simpler **detection** of small jobs
  - create **customized LML layouts**
- long term enhancements:
  - LML file **caching**
  - new monitoring views: histograms, history, prediction
  - integration of **JuFo**

JÜLICH FORSCHUNGSZENTRUM

# Your ideas?

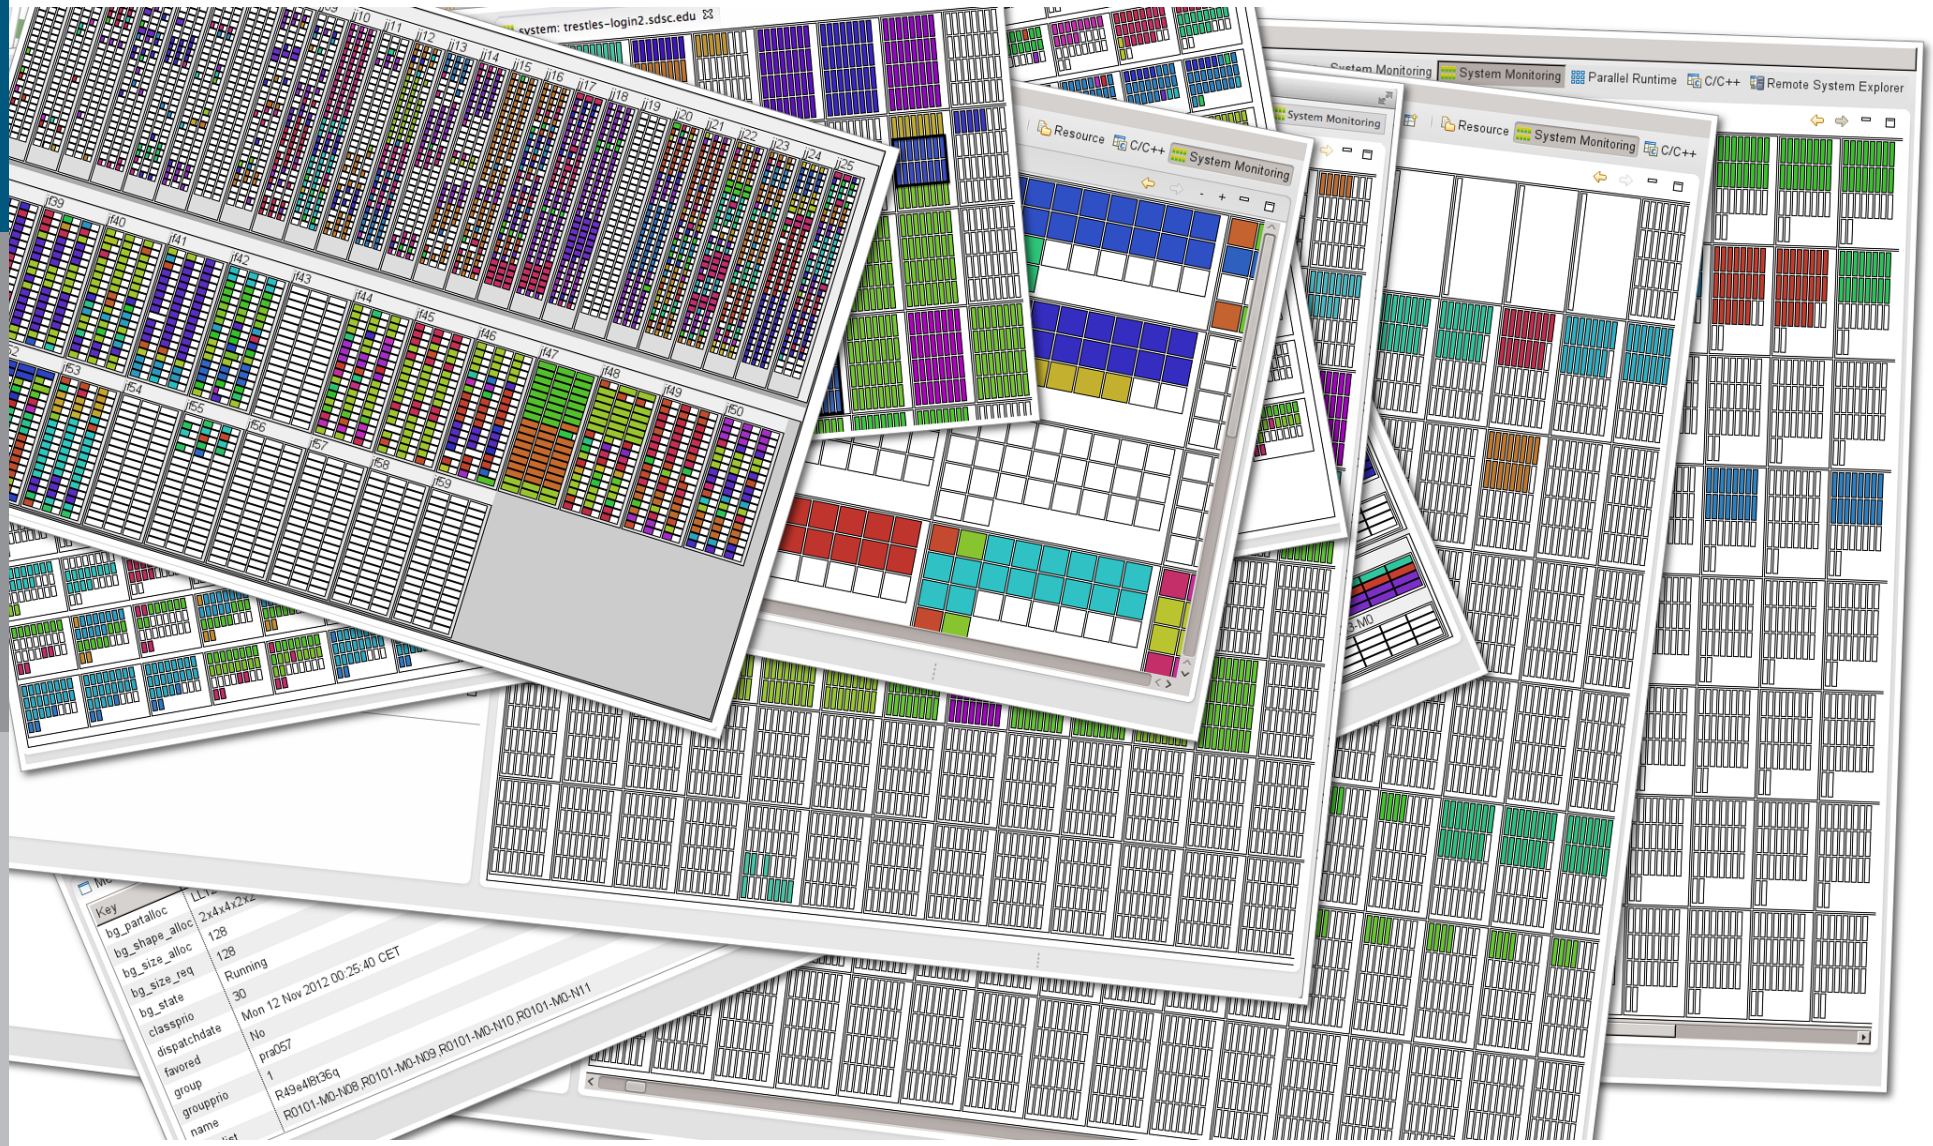

JÜLICH FORSCHUNGSZENTRUM

# Contact

- **E-mail**: c.karbach@fz-juelich.de, w.frings@fz-juelich.de
- **LLview** $\rightarrow$ `http://www.fz-juelich.de/jsc/llview`
- **LML** $\rightarrow$ `http://llview.zam.kfa-juelich.de/LML`