JÜLICH
FORSCHUNGSZENTRUM

System Monitoring and Job Reports with LLview

November 24, 2017 | Carsten Karbach

Part I: System Monitoring

November 24, 2017 | Carsten Karbach

JÜLICH FORSCHUNGSZENTRUM

Motivation

- Is my job running?
- When will it start?
- How is the current load?
- How is my job placed?

```
Id                      Owner      Submitted   ST PRI Class        Running On
----------------------- ---------- ----------- -- --- ------------ -----------
juqueen1c1.192822.0     curioni    11/7  11:22 I  50  systemall
juqueen1c1.192824.0     curioni    11/7  11:23 I  50  systemall
juqueen1c1.192825.0     curioni    11/7  11:23 I  50  systemall
juqueen1c1.192835.0     curioni    11/7  13:37 I  50  systemall
juqueen4c1.82115.0      hhh150     11/4  19:13 I  50  m016
juqueen1c1.191375.0     hdu180     11/4  23:25 I  50  m008
juqueen1c1.192826.0     jzam1159   11/7  11:27 I  50  m008
juqueen1c1.192844.0     jhpc0902   11/7  14:00 I  50  m008
juqueen2c1.190024.0     pra08801   11/3  04:31 I  50  m004
juqueen4c1.81936.0      jias1304   11/3  09:49 I  50  m004
juqueen4c1.81937.0      jias1304   11/3  09:49 I  50  m004
juqueen4c1.81938.0      jias1304   11/3  09:49 I  50  m004
juqueen4c1.81939.0      jias1304   11/3  09:50 I  50  m004
juqueen2c1.190675.0     jzam0420   11/5  13:42 I  50  m004
juqueen1c1.190354.10    hwu081     11/3  12:22 I  50  m002
juqueen2c1.190857.0     jiff4605   11/6  10:50 I  50  m002
juqueen2c1.190971.0     hhh073     11/7  10:11 I  50  m002
juqueen2c1.190972.0     grs30007   11/7  10:12 I  50  m002
juqueen3c1.79205.0      jinb3309   11/7  11:56 I  50  m002
juqueen1c1.192846.0     jzam0435   11/7  14:18 I  50  m002
juqueen4c1.80786.3      hbo273     10/28 11:49 I  50  m001
juqueen2c1.189769.2     hbo381     11/1  23:48 I  50  m001
juqueen1c1.190640.0     hch02r     11/4  07:03 I  50  m001
juqueen1c1.190643.0     hch02r     11/4  07:03 I  50  m001
juqueen1c1.190642.0     hch02r     11/4  07:03 I  50  m001
juqueen1c1.190641.0     hch02r     11/4  07:03 I  50  m001
juqueen1c1.190647.0     hch02r     11/4  07:03 I  50  m001
juqueen1c1.190646.0     hch02r     11/4  07:03 I  50  m001
juqueen1c1.190645.0     hch02r     11/4  07:03 I  50  m001
juqueen1c1.190644.0     hch02r     11/4  07:03 I  50  m001
juqueen1c1.190648.0     hch02r     11/4  07:03 I  50  m001
juqueen3c1.78904.7      jikp0501   11/5  09:39 I  50  m001
juqueen2c1.190615.0     hhh045     11/5  09:55 I  50  m001
juqueen2c1.190617.0     hhh045     11/5  09:58 I  50  m001
juqueen4c1.82192.0      hhh045     11/5  14:42 I  50  m001
juqueen4c1.82197.0      hhh045     11/5  15:48 I  50  m001
juqueen2c1.190713.0     jikp0403   11/5  16:27 I  50  m001
juqueen2c1.190715.0     jikp0403   11/5  16:35 I  50  m001
juqueen2c1.190716.0     jikp0403   11/5  16:41 I  50  m001
juqueen1c1.191700.0     talahde    11/5  16:47 I  50  m001
juqueen2c1.190718.0     jikp0403   11/5  16:48 I  50  m001
juqueen2c1.190723.0     jikp0403   11/5  16:56 I  50  m001
juqueen1c1.191701.0     talahde    11/5  16:59 I  50  m001
juqueen2c1.190725.0     jikp0403   11/5  17:11 I  50  m001
--More--
```

JÜLICH FORSCHUNGSZENTRUM

Why system monitoring?

- For users
  - Controlling own running and waiting jobs
  - Planning job submissions
  - Use of idling resources
- For administrators
  - Global overview of system utilization
  - Throughput optimization
  - Batch system configuration optimization
  - Adaptive change of scheduling parameters

$\Rightarrow$ LLview

  - Compact display of all usage data in one window
  - Easy access to system's status data
  - Interactive display for linking information
  - Open Source (BSD-style)
  - Available for all JSC systems

JÜLICH FORSCHUNGSZENTRUM

LLview

→ Visualizes supercomputer status on a single screen

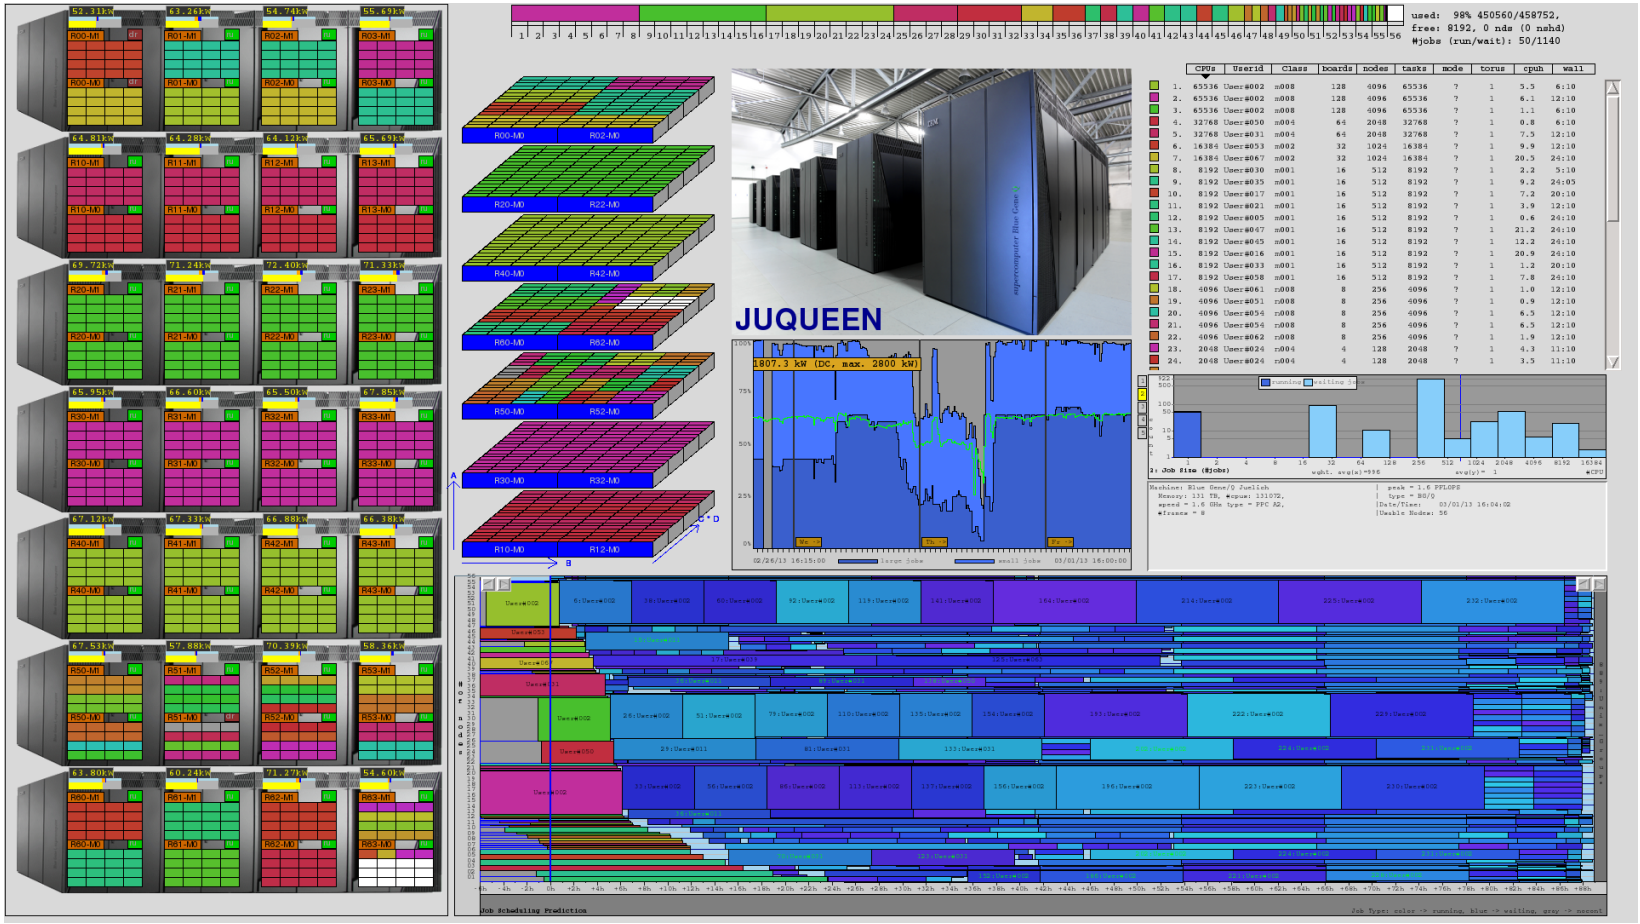

Source: Screenshot LLview for JUQUEEN

JÜLICH FORSCHUNGSZENTRUM

LLview Example: JURECA

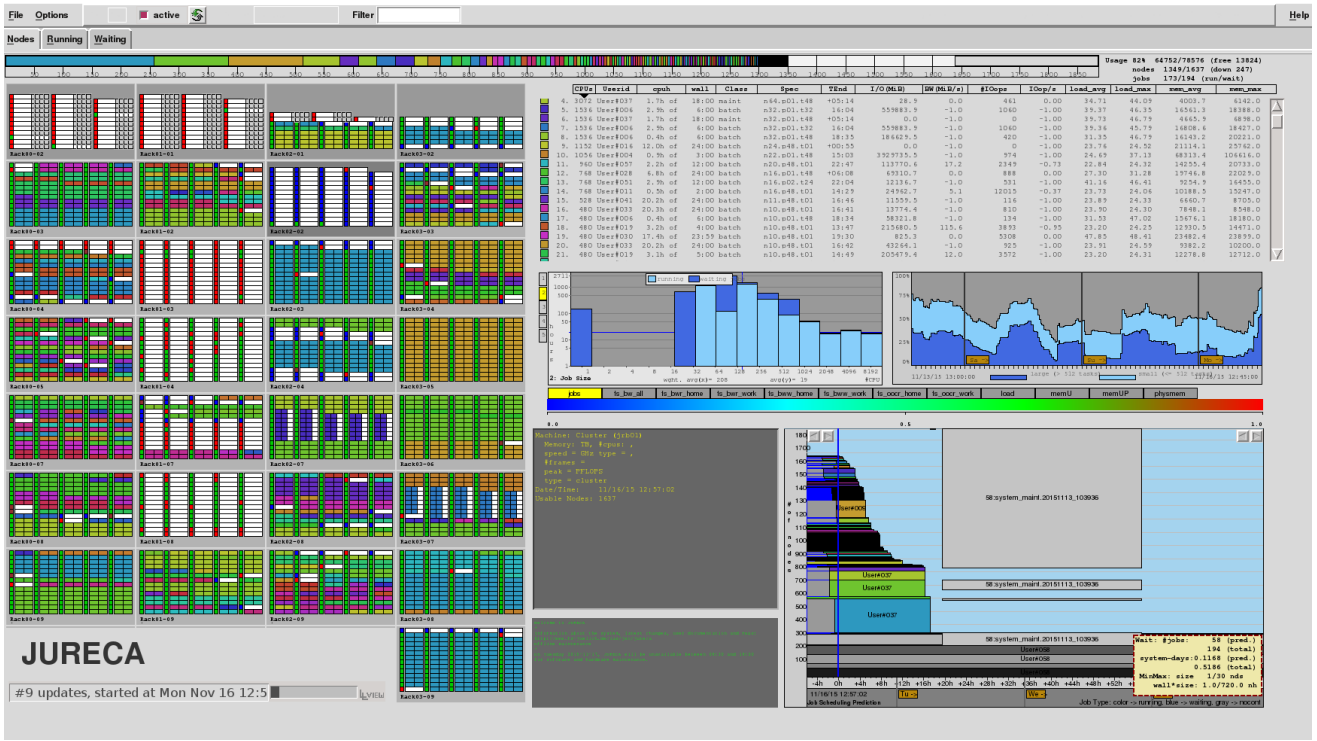

Source: Screenshot LLview for JURECA

- High fragmentation
- Heterogeneous
- Batch System: SLURM
- Ongoing development

LLview Architecture

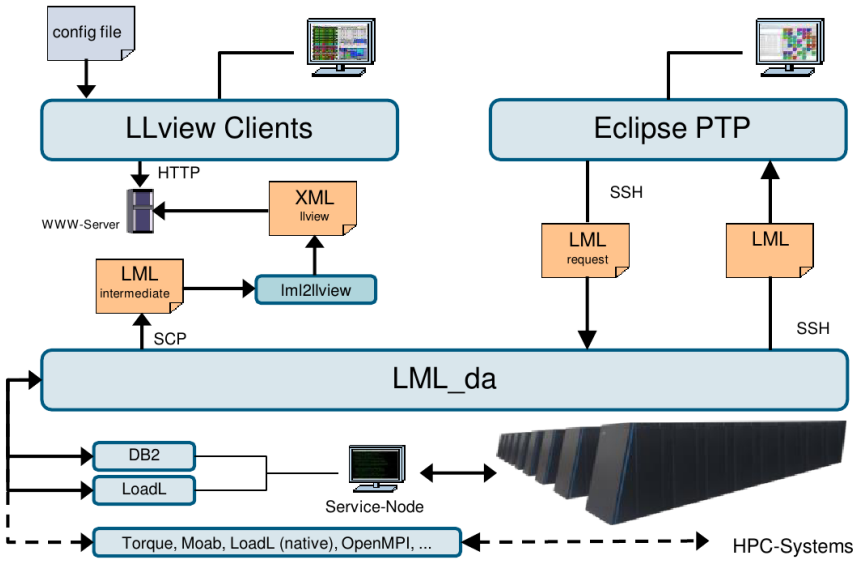

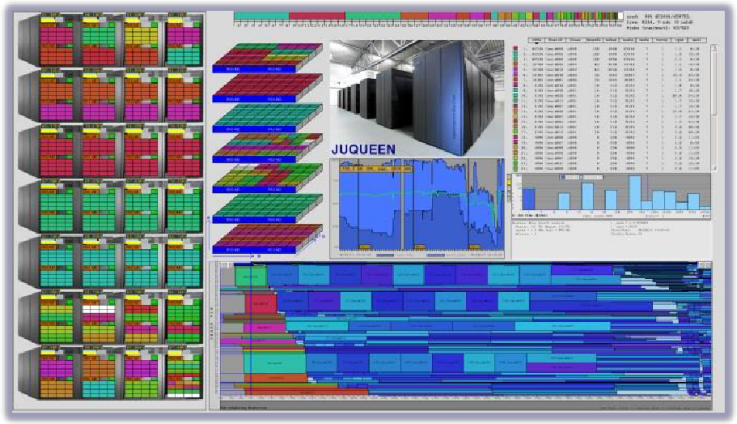

- Client-Server architecture, LML_da as backend
- Clients: Perl-Tk, PTP, Webinterface
- Platform independent: works on Windows, Mac and Linux
- Wide range of supported batch systems,
  minimal effort for extension
- Minor performance impact on monitored system,
  only **central batch** system is queried

JÜLICH FORSCHUNGSZENTRUM

PTP – Parallel Tools Platform

What is PTP?

- **IDE** for parallel application development
- Based on **Eclipse**
- **Open-source** project
- Developers: IBM, U.Oregon, UTK, Heidelberg University, NCSA, UIUC, JSC, ...

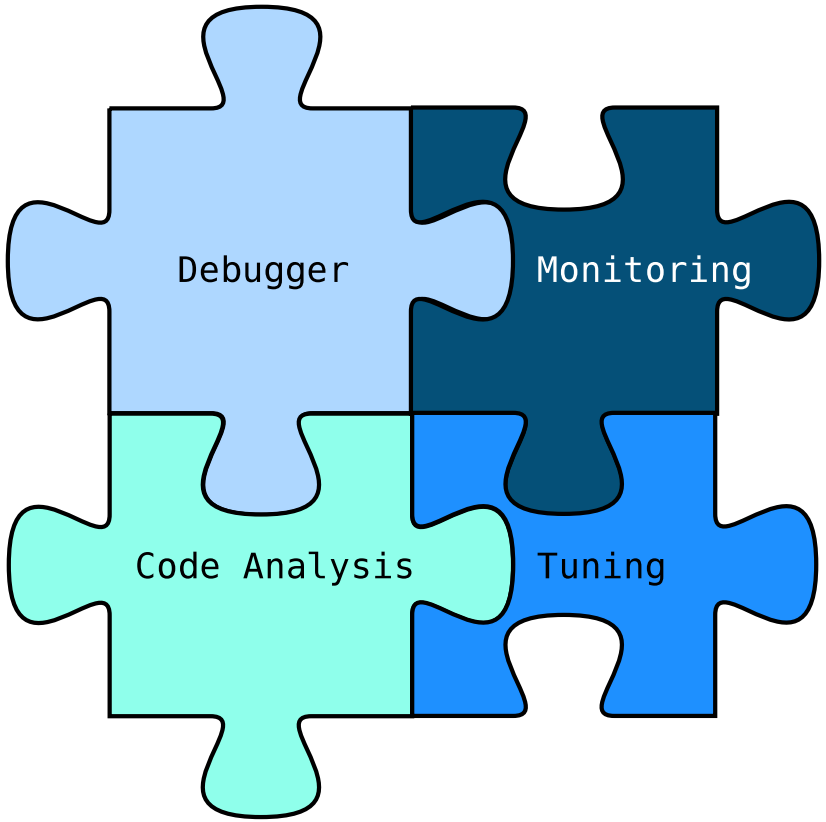

- **JSC-PTP tutorials** → `http://www.fz-juelich.de/SharedDocs/Downloads/IAS/JSC/EN/PTP/JSCPTPJuqueen.html`
- **PTP Download** → `http://www.eclipse.org/downloads/eclipse-packages/`, *Eclipse for Parallel Application Developers*

JÜLICH FORSCHUNGSZENTRUM

How to start the LLview client? I

- Four options to start LLview, sorted by effort to get started

Option 1: via SSH

```
ssh -X karbach@jureca
llview
```

Option 2: Webinterface

- Screenshots of LLview updated every minute (static)
- Link (static): `https://llview.fz-juelich.de/LLweb/juqueen/Image.html`
- Link (dynamic SVG): `https://llview.fz-juelich.de/LLweb/juqueen/svg/`
- Access secured by JSC webservice accounts
  - register at dispatch
  - request access to LLview via jsc-dispatch@fz-juelich.de

JÜLICH FORSCHUNGSZENTRUM

How to start the LLview client? II

Option 3: VNC

- start VNC server on JURECA with
  `vncserver -profile vis`
- tunnel VNC traffic to local system
- start VNC viewer
- click on LLview links
- detailed step by step guide here

Option 4: Local installation

- Download and install the LLview client locally

Part II: Job Reports

November 24, 2017 | Carsten Karbach

JÜLICH FORSCHUNGSZENTRUM

Job Reports

What can I get?

- Detailed reports on active and finished jobs
- Job summary and time-based diagrams of many job metrics such as load, memory usage, I/O
- Use cases: development, production checks, job health check, light-weight performance analysis

Details

- Includes metrics for CPUs and GPUs
- No instrumentation needed, data is retrieved from IPMI once per minute
- Development in progress

JÜLICH FORSCHUNGSZENTRUM

Access

- login: `https://llview.fz-juelich.de/LLweb/jureca/jobreport/login.php`
- Authenticate with webservice account, register here
- Only available for JURECA

JÜLICH FORSCHUNGSZENTRUM

User Jobs

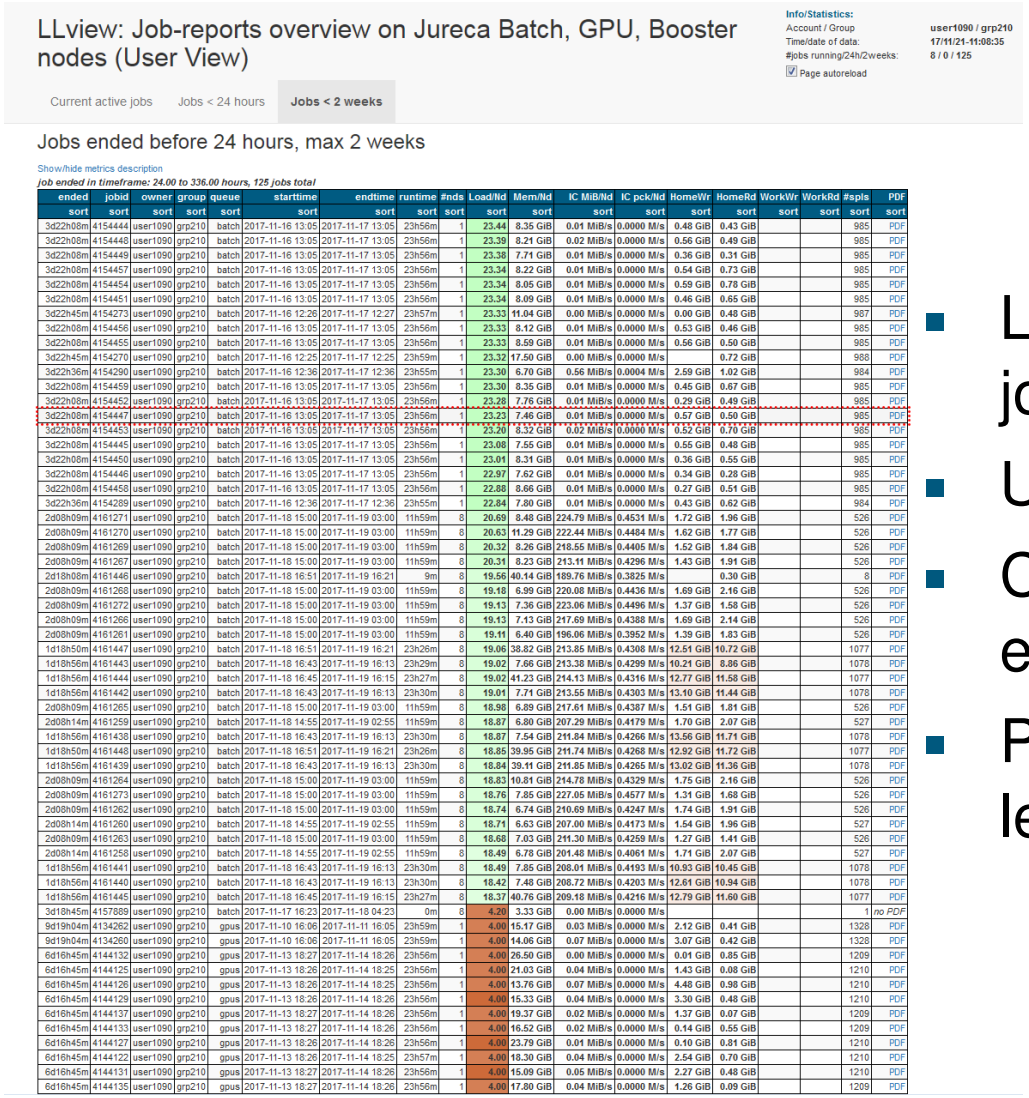

LLview: Job-reports overview on Jureca Batch, GPU, Booster nodes (User View)

Info/Statistics:
Account / Group: user1090 / grp210
Time/date of data: 17/11/21-11:08:35
#jobs running/24h/2weeks: 8 / 0 / 125
Page autoreload

Current active jobs | Jobs < 24 hours | Jobs < 2 weeks

Jobs ended before 24 hours, max 2 weeks

Show/hide metrics description

job ended in timeframe: 24.00 to 336.00 hours, 125 jobs total

ended	jobid	owner	group	queue	starttime	endtime	runtime	#nds	Load/Nd	Mem/Nd	IC MiB/Nd	IC pck/Nd	HomeWr	HomeRd	WorkWr	WorkRd	#spls	PDF
3d22h08m	4154444	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.44	8.35 GiB	0.01 MiB/s	0.0000 M/s	0.48 GiB	0.43 GiB			985	PDF
3d22h08m	4154448	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.39	8.21 GiB	0.02 MiB/s	0.0000 M/s	0.56 GiB	0.49 GiB			985	PDF
3d22h08m	4154449	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.38	7.71 GiB	0.01 MiB/s	0.0000 M/s	0.36 GiB	0.31 GiB			985	PDF
3d22h08m	4154457	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.34	8.22 GiB	0.01 MiB/s	0.0000 M/s	0.54 GiB	0.73 GiB			985	PDF
3d22h08m	4154454	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.34	8.05 GiB	0.01 MiB/s	0.0000 M/s	0.59 GiB	0.78 GiB			985	PDF
3d22h08m	4154451	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.34	8.09 GiB	0.01 MiB/s	0.0000 M/s	0.46 GiB	0.65 GiB			985	PDF
3d22h45m	4154273	user1090	grp210	batch	2017-11-16 12:26	2017-11-17 12:27	23h57m	1	23.33	11.04 GiB	0.00 MiB/s	0.0000 M/s	0.00 GiB	0.48 GiB			987	PDF
3d22h08m	4154456	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.33	8.12 GiB	0.01 MiB/s	0.0000 M/s	0.53 GiB	0.46 GiB			985	PDF
3d22h08m	4154455	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.33	8.59 GiB	0.01 MiB/s	0.0000 M/s	0.56 GiB	0.50 GiB			985	PDF
3d22h45m	4154270	user1090	grp210	batch	2017-11-16 12:25	2017-11-17 12:25	23h59m	1	23.32	17.50 GiB	0.00 MiB/s	0.0000 M/s		0.72 GiB			988	PDF
3d22h36m	4154290	user1090	grp210	batch	2017-11-16 12:36	2017-11-17 12:36	23h55m	1	23.30	6.70 GiB	0.56 MiB/s	0.0004 M/s	2.59 GiB	1.02 GiB			984	PDF
3d22h08m	4154459	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.30	8.35 GiB	0.01 MiB/s	0.0000 M/s	0.45 GiB	0.67 GiB			985	PDF
3d22h08m	4154452	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.28	7.76 GiB	0.01 MiB/s	0.0000 M/s	0.29 GiB	0.49 GiB			985	PDF
3d22h08m	4154447	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.23	7.46 GiB	0.01 MiB/s	0.0000 M/s	0.57 GiB	0.50 GiB			985	PDF
3d22h08m	4154453	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.20	8.32 GiB	0.02 MiB/s	0.0000 M/s	0.52 GiB	0.70 GiB			985	PDF
3d22h08m	4154445	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.08	7.55 GiB	0.01 MiB/s	0.0000 M/s	0.55 GiB	0.48 GiB			985	PDF
3d22h08m	4154450	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	23.01	8.31 GiB	0.01 MiB/s	0.0000 M/s	0.36 GiB	0.55 GiB			985	PDF
3d22h08m	4154446	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	22.97	7.62 GiB	0.01 MiB/s	0.0000 M/s	0.34 GiB	0.28 GiB			985	PDF
3d22h08m	4154458	user1090	grp210	batch	2017-11-16 13:05	2017-11-17 13:05	23h56m	1	22.88	8.66 GiB	0.01 MiB/s	0.0000 M/s	0.27 GiB	0.51 GiB			985	PDF
3d22h36m	4154289	user1090	grp210	batch	2017-11-16 12:36	2017-11-17 12:36	23h55m	1	22.84	7.80 GiB	0.01 MiB/s	0.0000 M/s	0.43 GiB	0.62 GiB			984	PDF
2d08h09m	4161271	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	20.69	8.48 GiB	224.79 MiB/s	0.4531 M/s	1.72 GiB	1.96 GiB			526	PDF
2d08h09m	4161270	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	20.63	11.29 GiB	222.44 MiB/s	0.4484 M/s	1.62 GiB	1.77 GiB			526	PDF
2d08h09m	4161269	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	20.32	8.26 GiB	218.55 MiB/s	0.4405 M/s	1.52 GiB	1.84 GiB			526	PDF
2d08h09m	4161267	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	20.31	8.23 GiB	213.11 MiB/s	0.4296 M/s	1.43 GiB	1.91 GiB			526	PDF
2d18h08m	4161446	user1090	grp210	batch	2017-11-18 16:51	2017-11-19 16:21	9m	8	19.56	40.14 GiB	189.76 MiB/s	0.3825 M/s		0.30 GiB			8	PDF
2d08h09m	4161268	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	19.18	6.99 GiB	220.08 MiB/s	0.4436 M/s	1.69 GiB	2.16 GiB			526	PDF
2d08h09m	4161272	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	19.13	7.36 GiB	223.06 MiB/s	0.4496 M/s	1.37 GiB	1.58 GiB			526	PDF
2d08h09m	4161266	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	19.13	7.13 GiB	217.69 MiB/s	0.4388 M/s	1.69 GiB	2.14 GiB			526	PDF
2d08h09m	4161261	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	19.11	6.40 GiB	196.06 MiB/s	0.3952 M/s	1.39 GiB	1.83 GiB			526	PDF
1d18h50m	4161447	user1090	grp210	batch	2017-11-18 16:51	2017-11-19 16:21	23h26m	8	19.06	38.82 GiB	213.85 MiB/s	0.4308 M/s	12.51 GiB	10.72 GiB			1077	PDF
1d18h56m	4161443	user1090	grp210	batch	2017-11-18 16:43	2017-11-19 16:13	23h29m	8	19.02	7.66 GiB	213.38 MiB/s	0.4299 M/s	10.21 GiB	8.86 GiB			1078	PDF
1d18h56m	4161444	user1090	grp210	batch	2017-11-18 16:45	2017-11-19 16:15	23h27m	8	19.02	41.23 GiB	214.13 MiB/s	0.4316 M/s	12.77 GiB	11.58 GiB			1077	PDF
1d18h56m	4161442	user1090	grp210	batch	2017-11-18 16:43	2017-11-19 16:13	23h30m	8	19.01	7.71 GiB	213.55 MiB/s	0.4303 M/s	13.10 GiB	11.44 GiB			1078	PDF
2d08h09m	4161265	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	18.98	6.89 GiB	217.61 MiB/s	0.4387 M/s	1.51 GiB	1.81 GiB			526	PDF
2d08h14m	4161259	user1090	grp210	batch	2017-11-18 14:55	2017-11-19 02:55	11h59m	8	18.87	6.80 GiB	207.29 MiB/s	0.4179 M/s	1.70 GiB	2.07 GiB			527	PDF
1d18h56m	4161438	user1090	grp210	batch	2017-11-18 16:43	2017-11-19 16:13	23h30m	8	18.87	7.54 GiB	211.84 MiB/s	0.4266 M/s	13.56 GiB	11.71 GiB			1078	PDF
1d18h50m	4161448	user1090	grp210	batch	2017-11-18 16:51	2017-11-19 16:21	23h26m	8	18.85	39.95 GiB	211.74 MiB/s	0.4268 M/s	12.92 GiB	11.72 GiB			1077	PDF
1d18h56m	4161439	user1090	grp210	batch	2017-11-18 16:43	2017-11-19 16:13	23h30m	8	18.84	39.11 GiB	211.85 MiB/s	0.4265 M/s	13.02 GiB	11.36 GiB			1078	PDF
2d08h09m	4161264	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	18.83	10.81 GiB	214.78 MiB/s	0.4329 M/s	1.75 GiB	2.16 GiB			526	PDF
2d08h09m	4161273	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	18.76	7.85 GiB	227.05 MiB/s	0.4577 M/s	1.31 GiB	1.68 GiB			526	PDF
2d08h09m	4161262	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	18.74	6.74 GiB	210.69 MiB/s	0.4247 M/s	1.74 GiB	1.91 GiB			526	PDF
2d08h14m	4161260	user1090	grp210	batch	2017-11-18 14:55	2017-11-19 02:55	11h59m	8	18.71	6.63 GiB	207.00 MiB/s	0.4173 M/s	1.54 GiB	1.96 GiB			527	PDF
2d08h09m	4161263	user1090	grp210	batch	2017-11-18 15:00	2017-11-19 03:00	11h59m	8	18.68	7.03 GiB	211.30 MiB/s	0.4259 M/s	1.27 GiB	1.41 GiB			526	PDF
2d08h14m	4161258	user1090	grp210	batch	2017-11-18 14:55	2017-11-19 02:55	11h59m	8	18.49	6.78 GiB	201.48 MiB/s	0.4061 M/s	1.71 GiB	2.07 GiB			527	PDF
1d18h56m	4161441	user1090	grp210	batch	2017-11-18 16:43	2017-11-19 16:13	23h30m	8	18.49	7.85 GiB	208.01 MiB/s	0.4193 M/s	10.93 GiB	10.45 GiB			1078	PDF
1d18h56m	4161440	user1090	grp210	batch	2017-11-18 16:43	2017-11-19 16:13	23h30m	8	18.42	7.48 GiB	208.72 MiB/s	0.4203 M/s	12.61 GiB	10.94 GiB			1078	PDF
1d18h56m	4161445	user1090	grp210	batch	2017-11-18 16:45	2017-11-19 16:15	23h27m	8	18.37	40.76 GiB	209.18 MiB/s	0.4216 M/s	12.79 GiB	11.60 GiB			1077	PDF
3d18h45m	4157889	user1090	grp210	batch	2017-11-17 16:23	2017-11-18 04:23	0m	8	4.20	3.33 GiB	0.00 MiB/s	0.0000 M/s					1	*no PDF*
9d19h04m	4134262	user1090	grp210	gpus	2017-11-10 16:06	2017-11-11 16:05	23h59m	1	4.00	15.17 GiB	0.03 MiB/s	0.0000 M/s	2.12 GiB	0.41 GiB			1328	PDF
9d19h04m	4134260	user1090	grp210	gpus	2017-11-10 16:06	2017-11-11 16:05	23h59m	1	4.00	14.06 GiB	0.07 MiB/s	0.0000 M/s	3.07 GiB	0.42 GiB			1328	PDF
6d16h45m	4144132	user1090	grp210	gpus	2017-11-13 18:27	2017-11-14 18:26	23h56m	1	4.00	26.50 GiB	0.00 MiB/s	0.0000 M/s	0.01 GiB	0.85 GiB			1209	PDF
6d16h45m	4144125	user1090	grp210	gpus	2017-11-13 18:26	2017-11-14 18:25	23h56m	1	4.00	21.03 GiB	0.04 MiB/s	0.0000 M/s	1.43 GiB	0.08 GiB			1210	PDF
6d16h45m	4144126	user1090	grp210	gpus	2017-11-13 18:26	2017-11-14 18:25	23h56m	1	4.00	13.76 GiB	0.07 MiB/s	0.0000 M/s	4.48 GiB	0.98 GiB			1210	PDF
6d16h45m	4144129	user1090	grp210	gpus	2017-11-13 18:26	2017-11-14 18:26	23h56m	1	4.00	15.33 GiB	0.04 MiB/s	0.0000 M/s	3.30 GiB	0.48 GiB			1210	PDF
6d16h45m	4144137	user1090	grp210	gpus	2017-11-13 18:27	2017-11-14 18:26	23h56m	1	4.00	19.37 GiB	0.02 MiB/s	0.0000 M/s	1.37 GiB	0.07 GiB			1209	PDF
6d16h45m	4144133	user1090	grp210	gpus	2017-11-13 18:27	2017-11-14 18:26	23h56m	1	4.00	16.52 GiB	0.02 MiB/s	0.0000 M/s	0.14 GiB	0.55 GiB			1209	PDF
6d16h45m	4144127	user1090	grp210	gpus	2017-11-13 18:26	2017-11-14 18:26	23h56m	1	4.00	23.79 GiB	0.01 MiB/s	0.0000 M/s	0.10 GiB	0.81 GiB			1210	PDF
6d16h45m	4144122	user1090	grp210	gpus	2017-11-13 18:26	2017-11-14 18:25	23h57m	1	4.00	18.30 GiB	0.04 MiB/s	0.0000 M/s	2.54 GiB	0.70 GiB			1210	PDF
6d16h45m	4144131	user1090	grp210	gpus	2017-11-13 18:27	2017-11-14 18:26	23h56m	1	4.00	15.09 GiB	0.05 MiB/s	0.0000 M/s	2.27 GiB	0.48 GiB			1210	PDF
6d16h45m	4144135	user1090	grp210	gpus	2017-11-13 18:27	2017-11-14 18:26	23h56m	1	4.00	17.80 GiB	0.04 MiB/s	0.0000 M/s	1.26 GiB	0.09 GiB			1209	PDF

- Lists for active and finished jobs
- Updates after 1-5 minutes
- Colormap for quick evaluation of metrics
- PDF column on the right leads to detailed job report

Detailed Job Report as PDF

Job-Info *(Jureca Batch, GPU, Booster nodes)*

Jobid: **4127061** User: **user1009** Group: **grp209** Date/Time of job data: **17/11/09-10:09:02**

Job runtime:	**1h36m**	***Job Performance metrics***	min	avg	max
	--> 80.09% *of wall: 2h00m*	Load (CPU-Nodes):	**0.01**	**41.01**	**47.78**
Job starttime:	**17/11/09-08:32:51**	Memory (CPU-Nodes):	**3103.00**	**3200.28**	**3417.00** *MiB*
Job last timestamp:	**17/11/09-10:09:05** (running)	Interconnect Traffic (in):		**0.00**	**0.68** *MiB/s*
Job endtime (est.):	17/11/09-10:33:01	Interconnect Traffic (out):		**0.03**	**0.50** *MiB/s*
Queue:	**batch**	Interconnect Packets (out):		**22**	**215** *pck/s*
Job Size, #Nodes:	**16** *#data points:* 90	Interconnect Packets (out):		**22**	**215** *pck/s*

Job I/O statistics	Total Data Write	Total Data Read	max. Data rate/Node Write	max. Data rate/Node Read	max. Open-Close Rate/Node
$HOME:	**0.00** MiB	**0.00** MiB	**0.00** MiB/s	**0.00** MiB/s	**1.60** op/s
$WORK:	**54.44** MiB	**5.63** MiB	**0.14** MiB/s	**0.09** MiB/s	**3.78** op/s

Job-Usage Overview

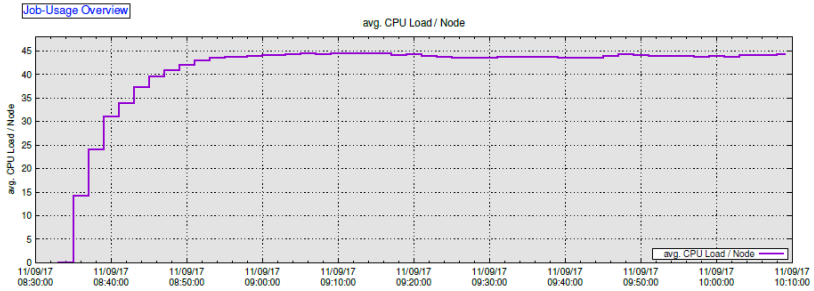

Table of Contents (6 pages)

created by LL view
www.fz-juelich.de/jsc/llview

- Diagrams for all available metrics per job over runtime
- Overall diagrams and diagrams per compute node

JÜLICH FORSCHUNGSZENTRUM

Example job – bad load balancing

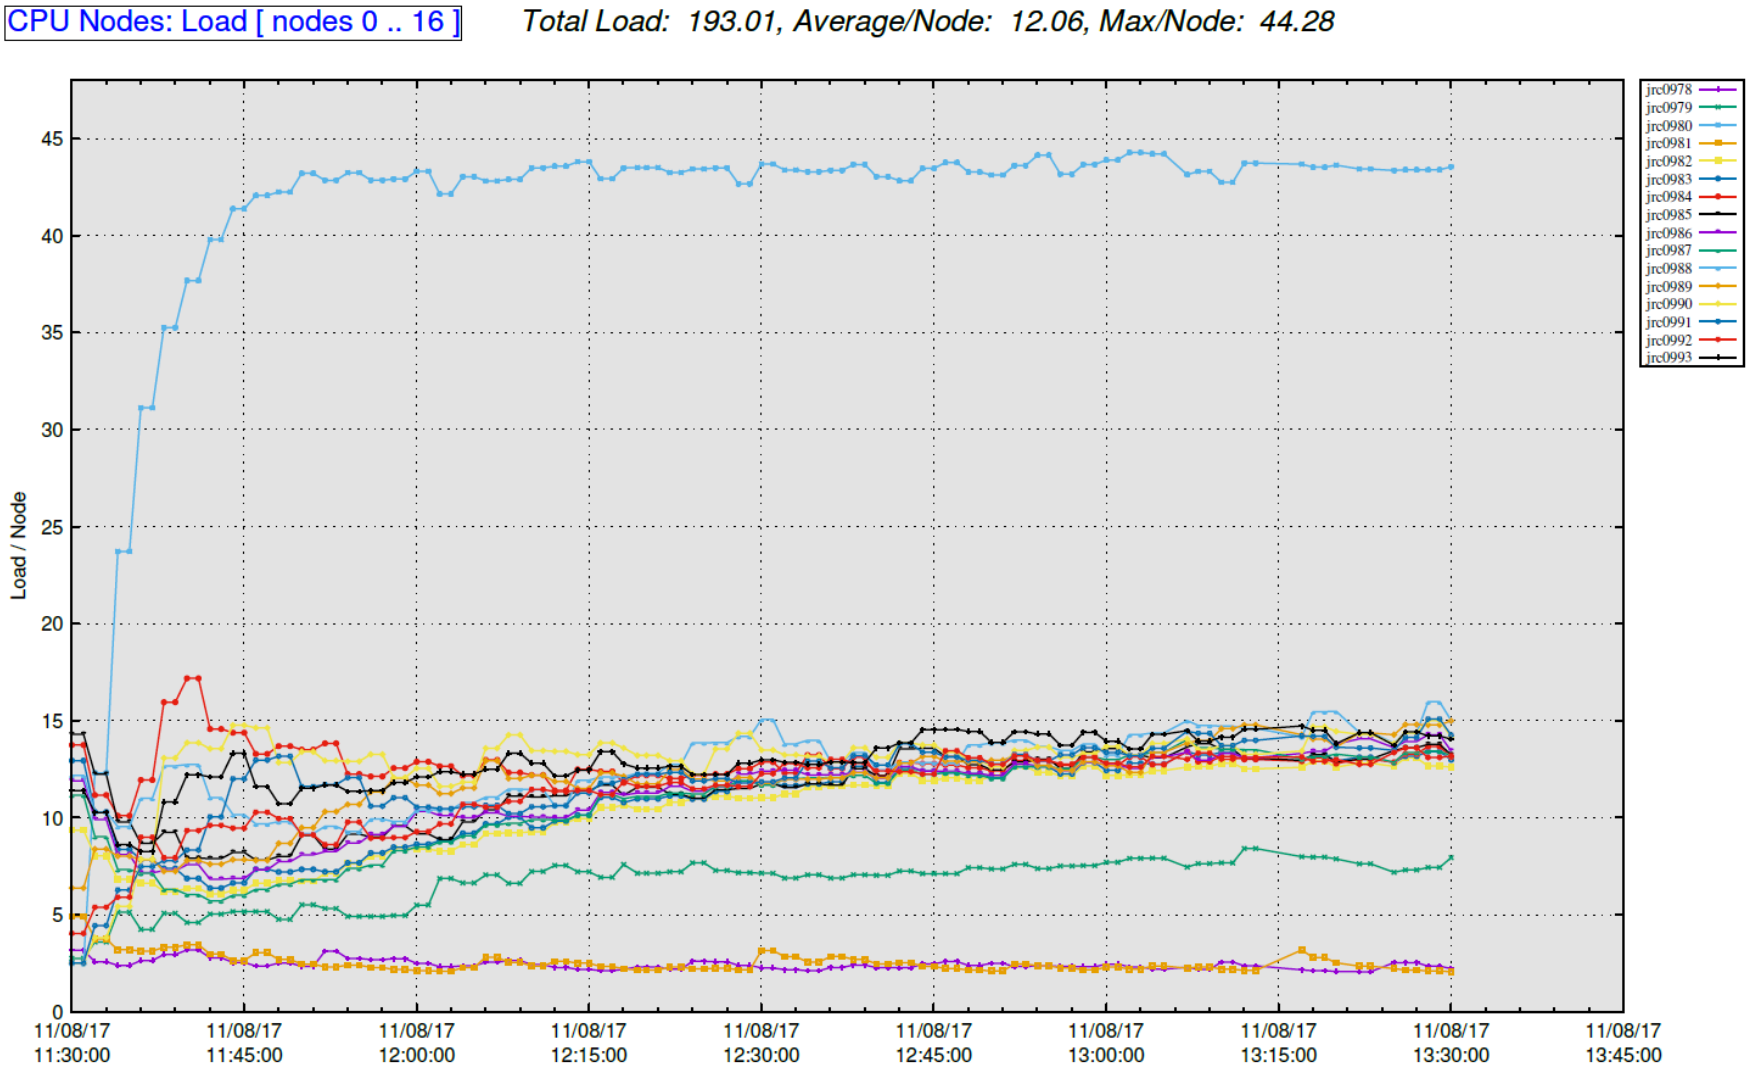

JÜLICH FORSCHUNGSZENTRUM

Example job – good load balancing

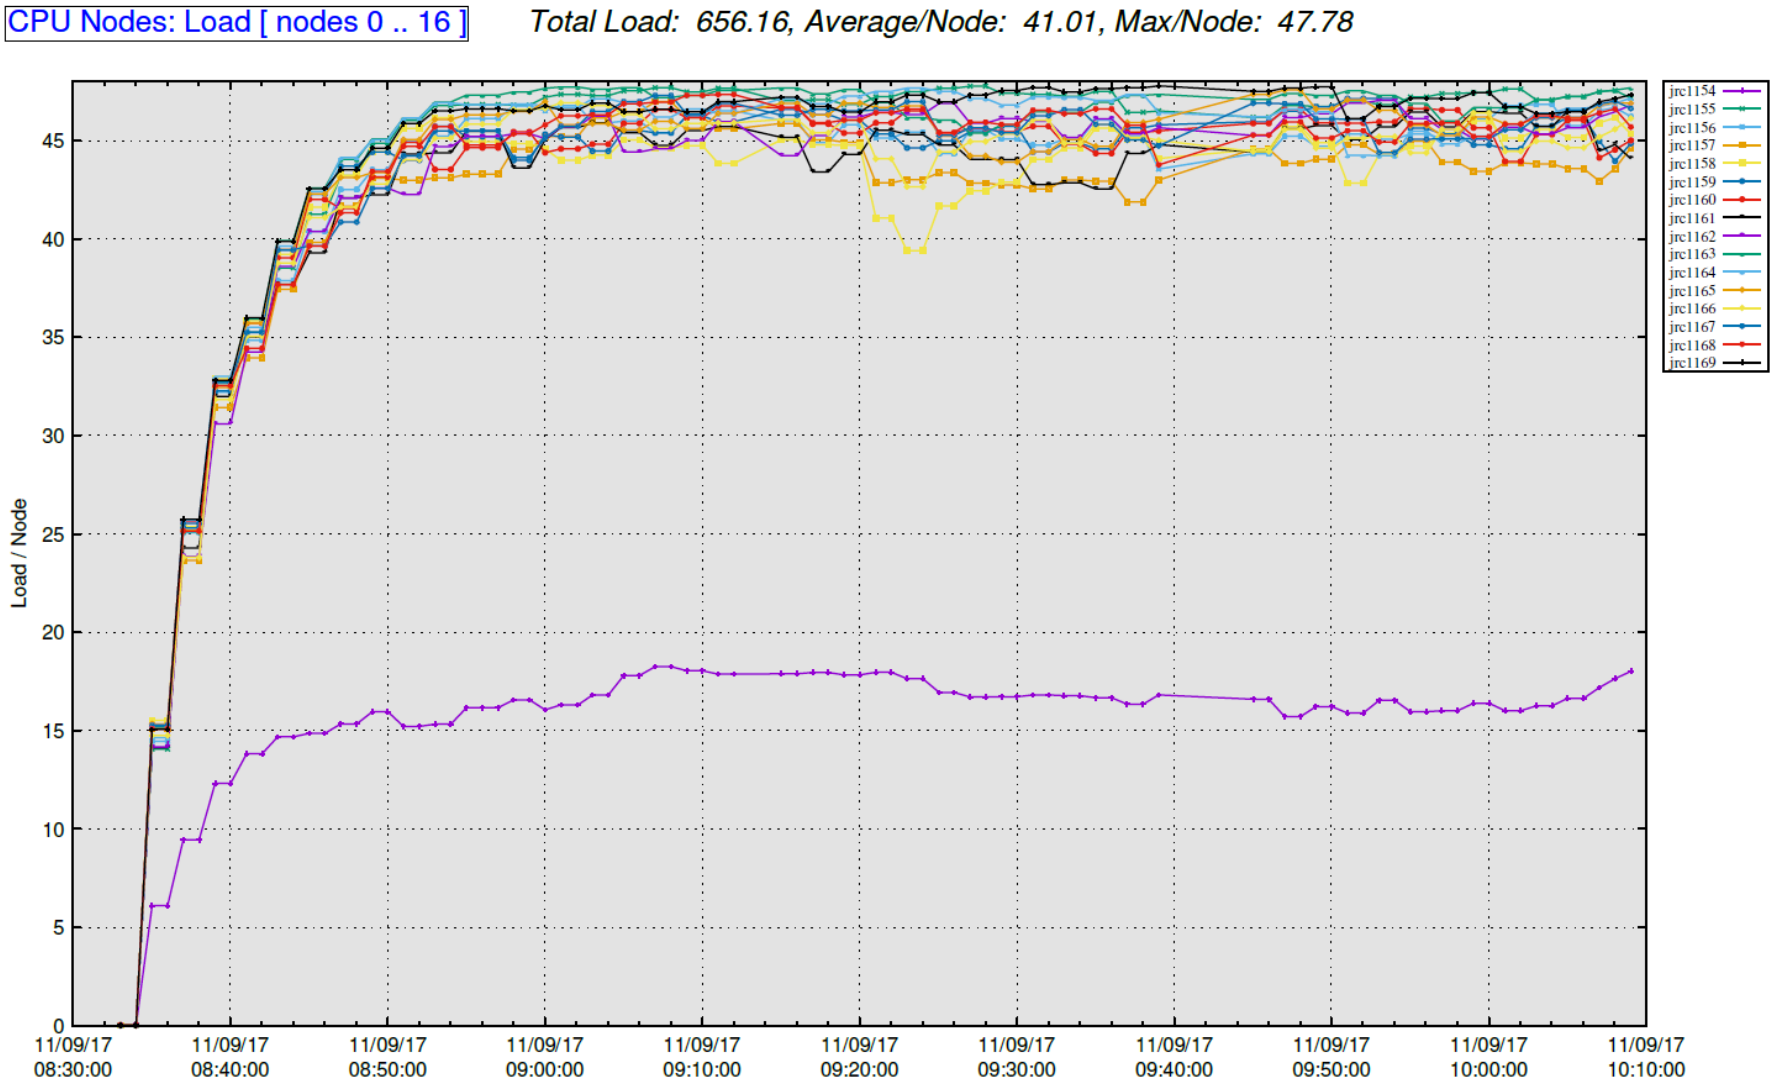

JÜLICH FORSCHUNGSZENTRUM

Contact

- **E-mail**: llview.jsc@fz-juelich.de
- **LLview** → `http://www.fz-juelich.de/jsc/llview`
- **Job Reports** → `https://llview.fz-juelich.de/LLweb/jureca/jobreport/login.php`
- **JSC-PTP tutorials** → `http://www.fz-juelich.de/SharedDocs/Downloads/IAS/JSC/EN/PTP/JSCPTPJuqueen.html`
- **PTP Download** → `http://www.eclipse.org/downloads/eclipse-packages/`

JÜLICH
FORSCHUNGSZENTRUM

Part III: Appendix – LLview Components

November 24, 2017 | Carsten Karbach

Node display

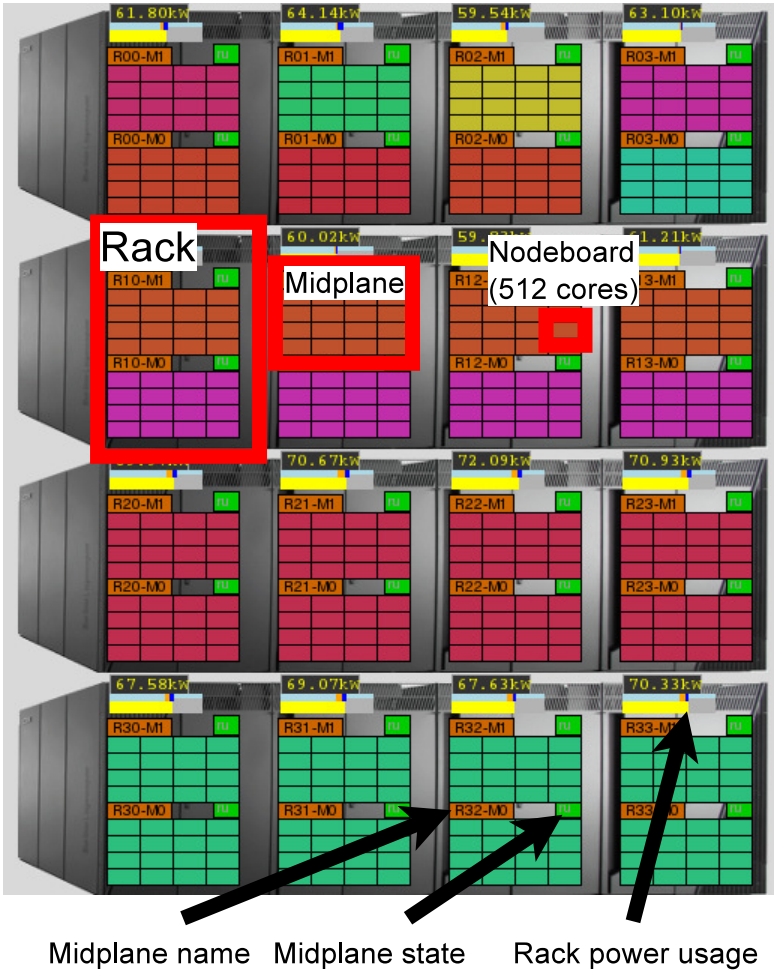

- Compute resources
- Job distribution
- White = Idle, Colored = running
- Node name
- Node status
- Level of detail

JÜLICH FORSCHUNGSZENTRUM

Node display

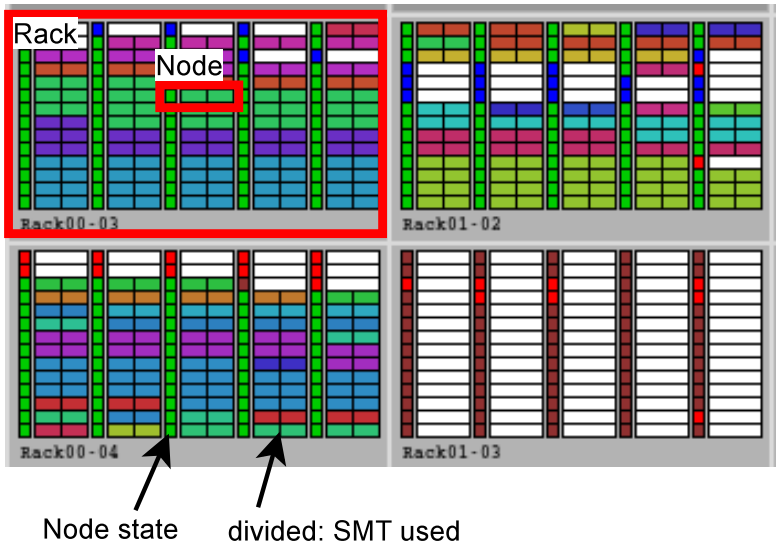

- Compute resources
- Job distribution
- White = Idle, Colored = running
- Node name
- Node status
- Level of detail

JÜLICH FORSCHUNGSZENTRUM

Job list

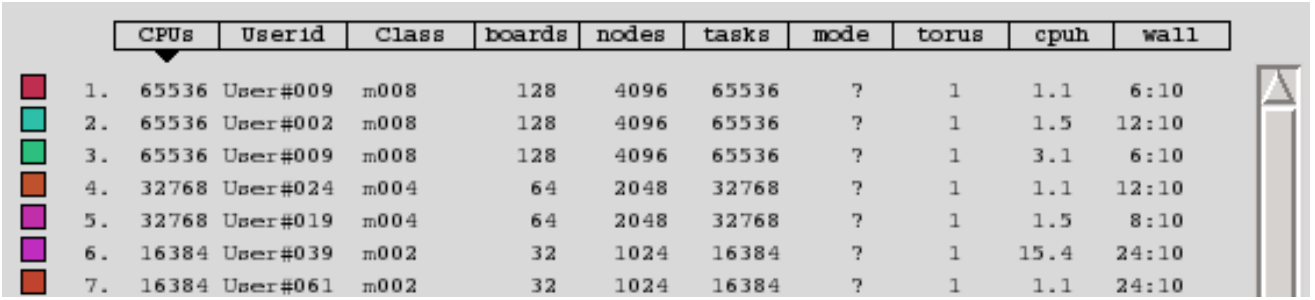

	CPUs	Userid	Class	boards	nodes	tasks	mode	torus	cpuh	wall
1.	65536	User#009	m008	128	4096	65536	?	1	1.1	6:10
2.	65536	User#002	m008	128	4096	65536	?	1	1.5	12:10
3.	65536	User#009	m008	128	4096	65536	?	1	3.1	6:10
4.	32768	User#024	m004	64	2048	32768	?	1	1.1	12:10
5.	32768	User#019	m004	64	2048	32768	?	1	1.5	8:10
6.	16384	User#039	m002	32	1024	16384	?	1	15.4	24:10
7.	16384	User#061	m002	32	1024	16384	?	1	1.1	24:10

- List of running jobs
- Most important attributes per job
- Sort by clicking on the column header
- Identifying color next to each job entry

JÜLICH FORSCHUNGSZENTRUM

Usagebar

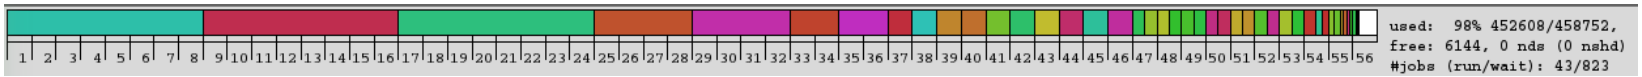

- Summary of system load
- Job size decreases from left to right
- White space shows idling resources
- Unit for JUQUEEN is midplanes,
  for JURECA nodes

JÜLICH FORSCHUNGSZENTRUM

History

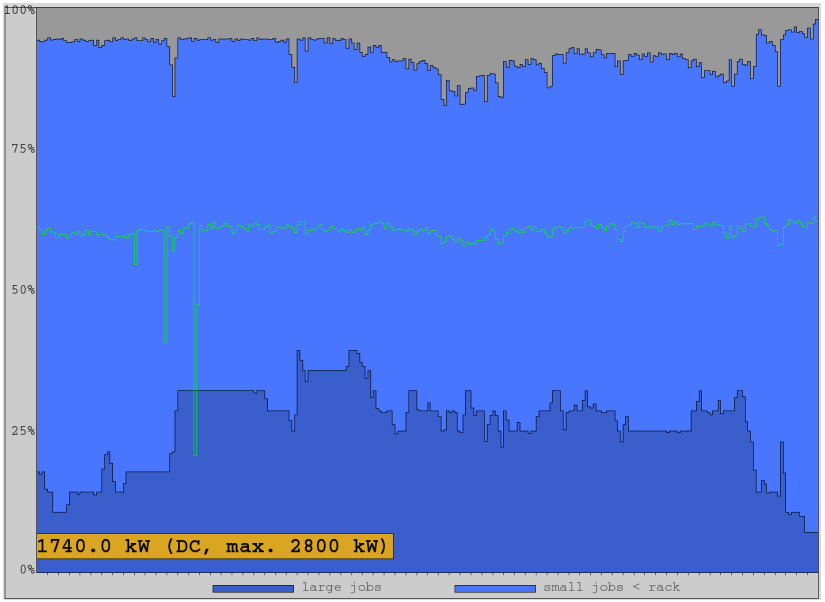

- 3-day load history
- Often divided into small and large jobs
  JUQUEEN (1 midplane), JURECA (512 tasks)
- Mouse-Over for detailed information
- Green line for special history value
  JUQUEEN (power)

JÜLICH FORSCHUNGSZENTRUM

Prediction

- Scheduler prediction based on submitted jobs
- Wall clock limit as job duration
- Blue = predicted job, Colored = running jobs
- Each rectangle one job,
  x-axis = time, y-axis = nodes/midplanes
- Use of idle times
- More transparent scheduling
- JUQUEEN: self-implemented, JURECA: use of SLURM's prediction

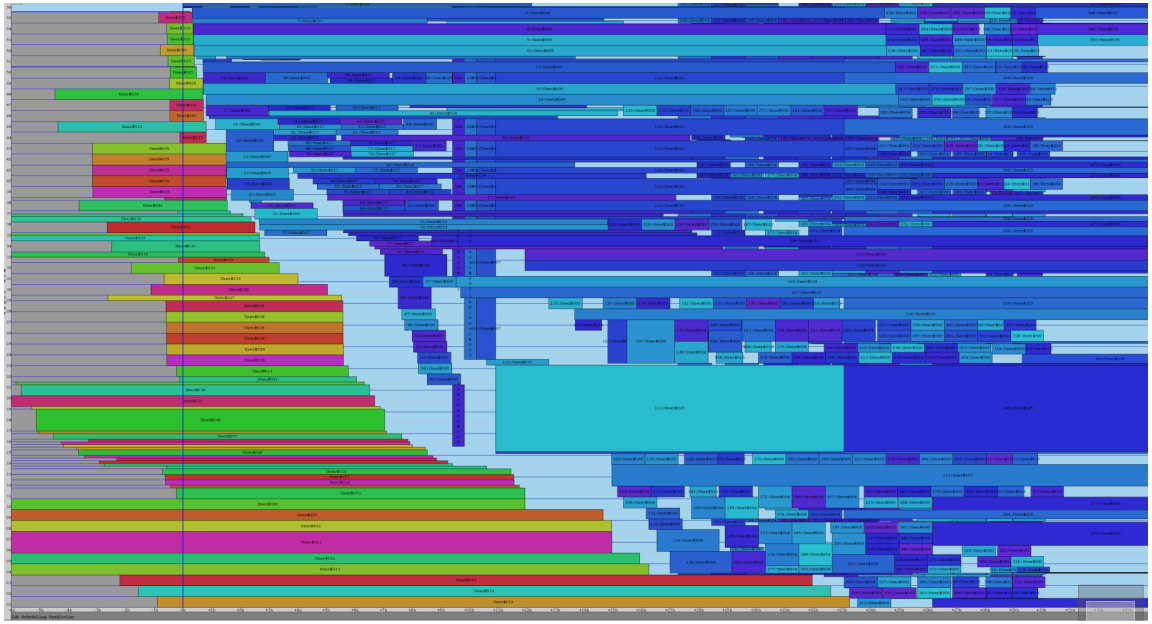

JÜLICH FORSCHUNGSZENTRUM

Statistics

- Statistic overview on system status
- Histograms on job size, wait time, queue load
- Highly configurable, define x-axis/y-axis domain, logarithmic/linear scale
- Overlayed diagrams for waiting/running jobs

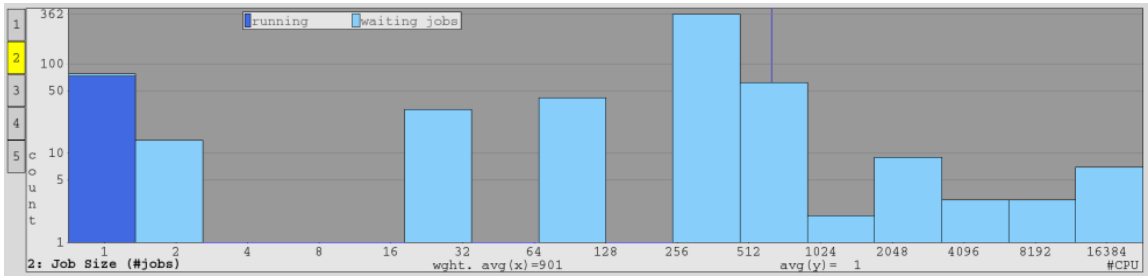

JÜLICH FORSCHUNGSZENTRUM

Infobox

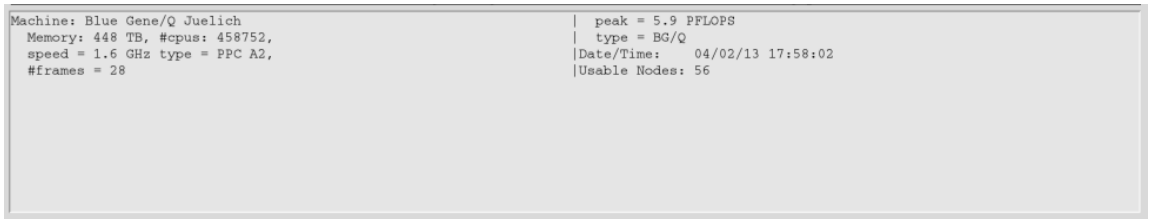

- Shows details on the currently focused object
- Mouse-Over triggers to display detailed data on the focused job/node/system/diagram

Interaction

- **Mouse-Over** jobs highlights job rectangles for the selected job in all components and shows details on the job in infobox

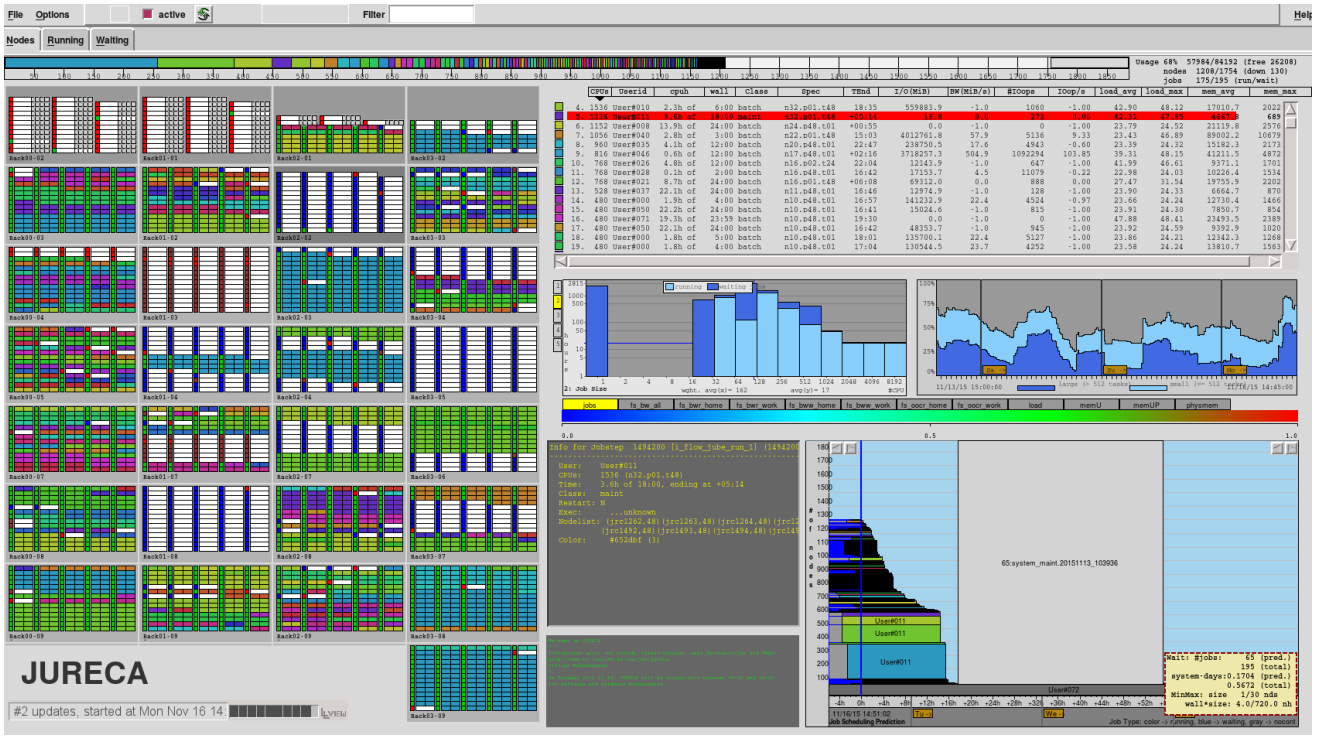

Interaction

- **Mouse-Over** jobs highlights job rectangles for the selected job in all components and shows details on the job in infobox
- **Mouse-Down** (Hold) removes color for all other jobs, only the selected job is colored

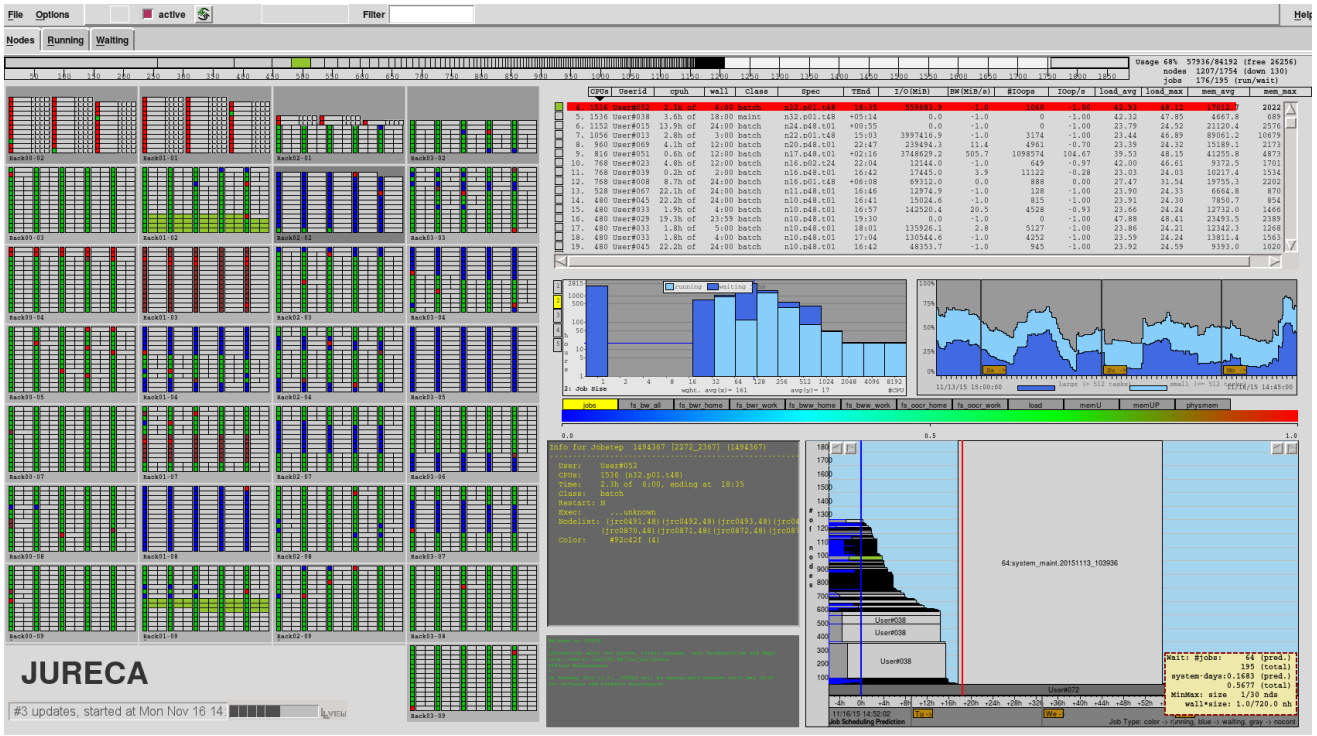

JÜLICH FORSCHUNGSZENTRUM

Interaction

- **Mouse-Over** jobs highlights job rectangles for the selected job in all components and shows details on the job in infobox
- **Mouse-Down** (Hold) removes color for all other jobs, only the selected job is colored
- **Double-Click** show job detail dialog
- Mouse interaction helps to link the information shown in all components

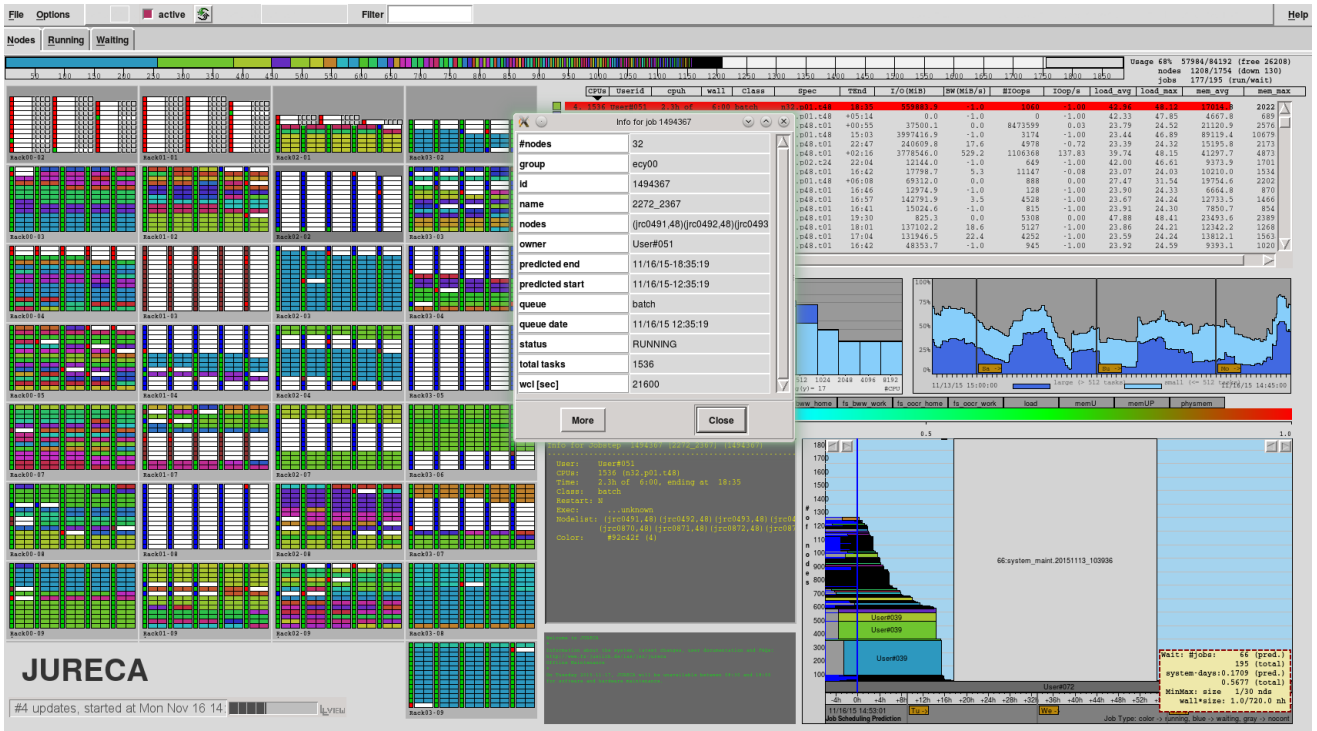

Part IV: Appendix – LLview Customization

November 24, 2017 | Carsten Karbach

JÜLICH FORSCHUNGSZENTRUM

LLview

- LLview is **highly customizable** due to numerous options
- Settings specific to HPC system type
- Start through option menu of main window or `Ctrl+o`
- Most options have **immediate effect** in the main window
- Some will become active at next start of LLview (e.g. *Data source* change)
- LLview layouts use **absolute** positioning
- You can use arrow keys to add/subtract one on numeric values
- Use Page up/Page down keys to add/subtract a bigger step on numeric values

JÜLICH FORSCHUNGSZENTRUM

LLview configuration files

Three configuration file locations (highest priority first):

1. anywhere on your file system passed to LLview with the `-rcfile` option
2. local `.llview.rc` configuration file in current directory or in HOME directory of the user
3. `llview.rc` in the installation directory of LLview. This file contains the system-wide settings

- Configuration files contain all LLview options
- You can change them in any text editor or via the LLview Option window

JÜLICH FORSCHUNGSZENTRUM

LLview options

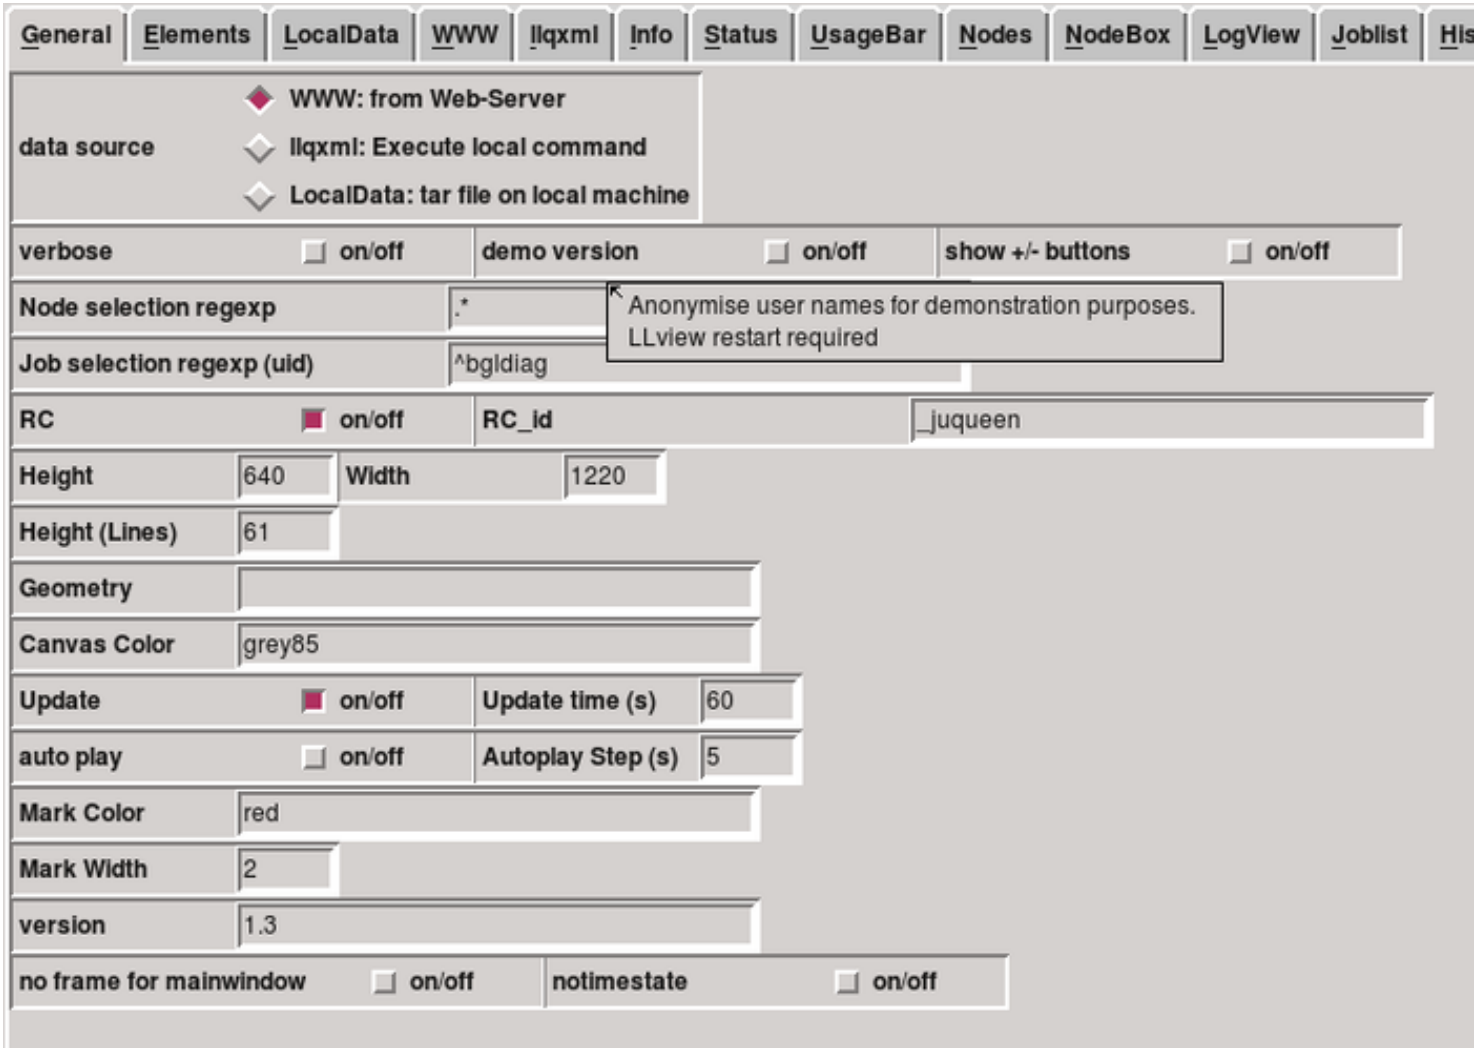

JÜLICH FORSCHUNGSZENTRUM

General Options

- *General* options for the main window
- Choose your preferred data source (*Web-Server*, *LocalData* or *local command*)
- *demo version*: anonymise usernames (for public display)
- *Job selection regexp*: filtering jobs by regular expressions
- Customize *Height* and *Width* of the main window
- *Canvas Color*: background, *Mark Color*: color for marking job in job list
- Choose time until next update
- *auto play* lets LLview mark different jobs automatically (for public display)

LLview Element options

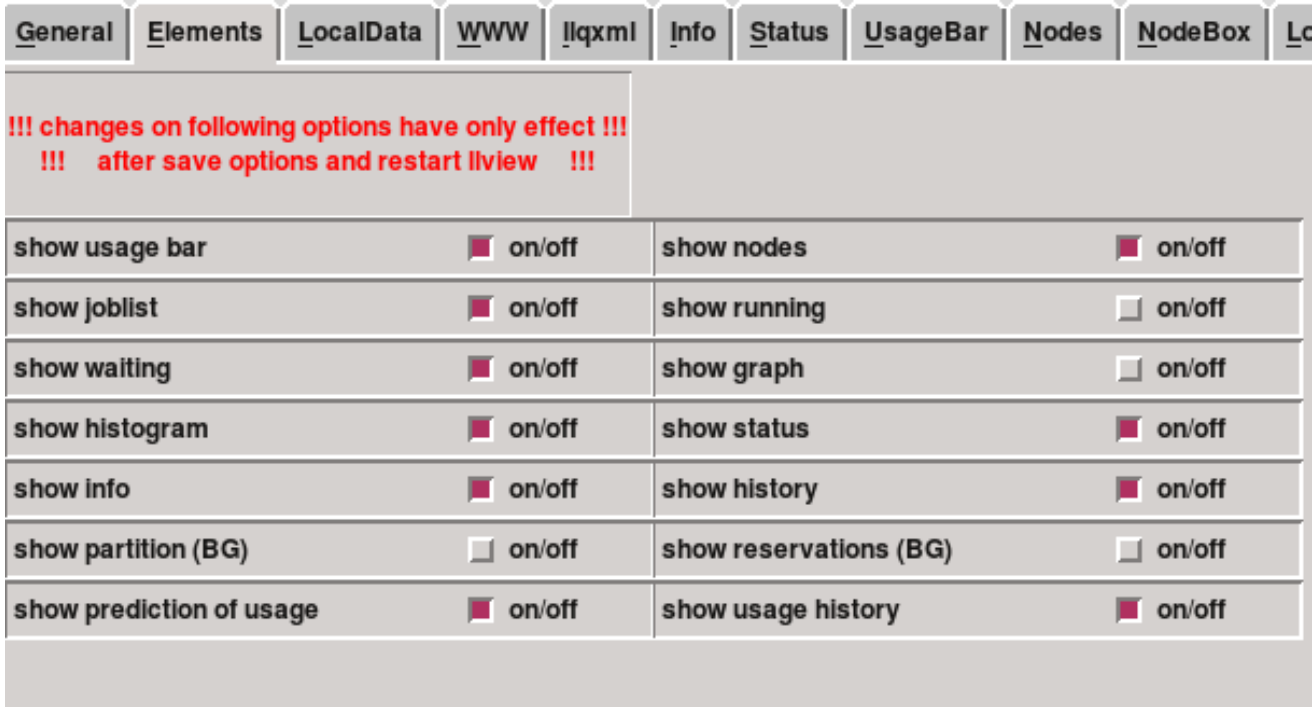

- Choose, which Elements to show
- For end users: components like *joblist*, *info*, *nodes* and *prediction* etc.
- Changes take effect after **restarting** LLview

JÜLICH FORSCHUNGSZENTRUM

LLview Node options

JÜLICH FORSCHUNGSZENTRUM

Node Options

- Customize *Height*, *Width* and *Margins* of the node display
- Logical node view available for BG systems e.g. JUQUEEN $\Rightarrow$ *Twin View* places adjacent midplanes next to each other in torus network
- Options for the *Twin View* are available in a new option subfolder $\Rightarrow$ *Log View*
- *Node attr.*: show scalar data e.g. temperature or power usage
- Show logo or site name

JÜLICH FORSCHUNGSZENTRUM

LLview Histogram options

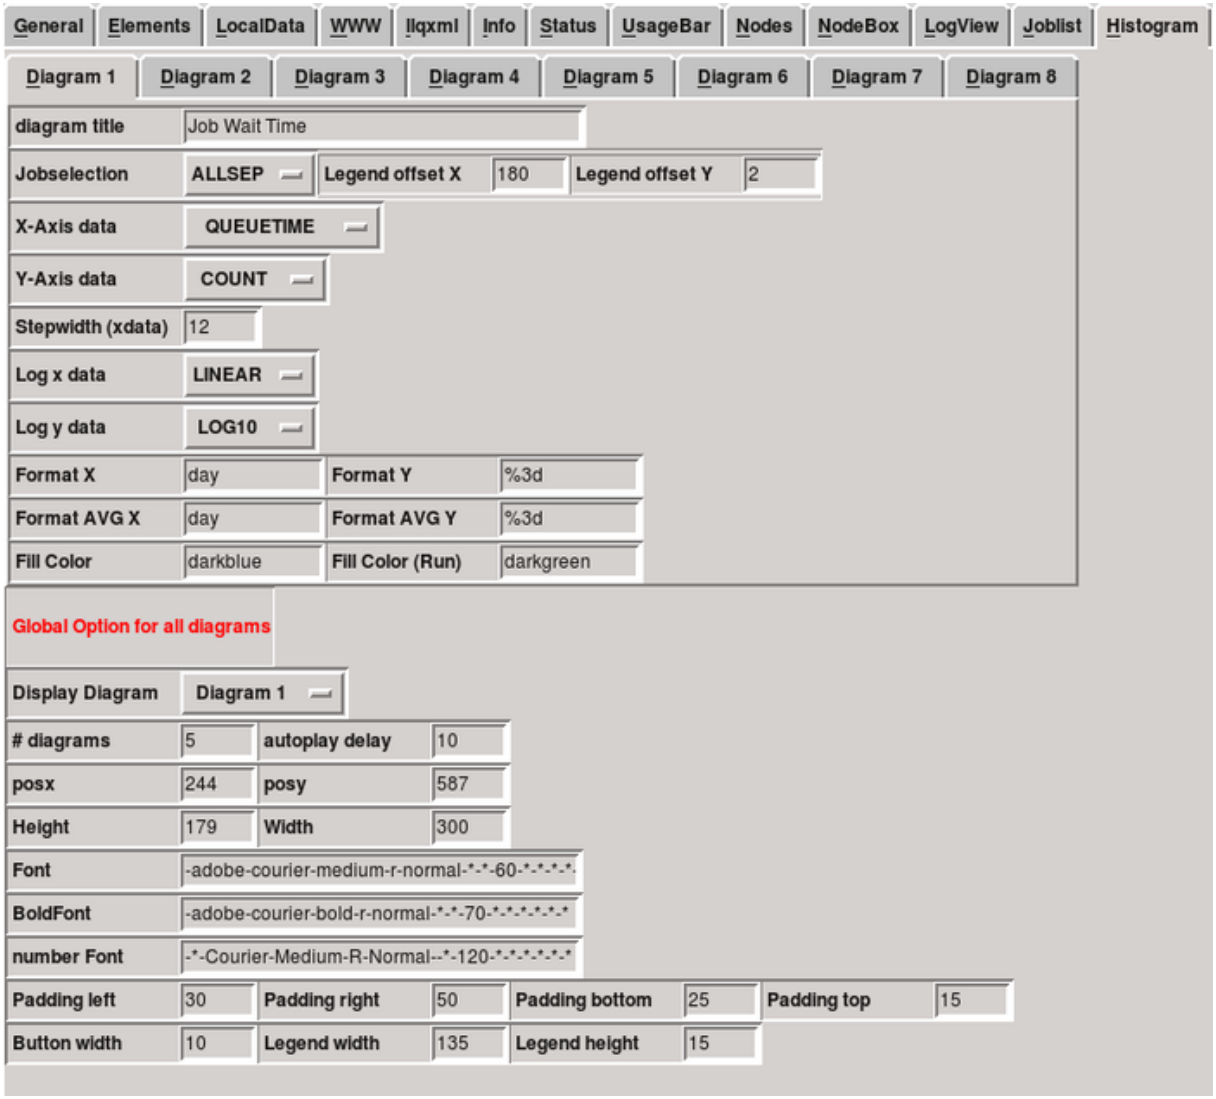

JÜLICH FORSCHUNGSZENTRUM

Histogram Options

- Each histogram shows distribution for a single job attribute, e.g. waiting time
- You can configure up to 8 histograms
- Jobs are grouped into discrete classes, y-axis shows count of jobs in each class
- Y-axis may also show number of CPUs, CPU hours or job duration
- Scaling may be linear or logarithmic
- *Auto play* is available for public display