



# LLview-based Job-Reporting

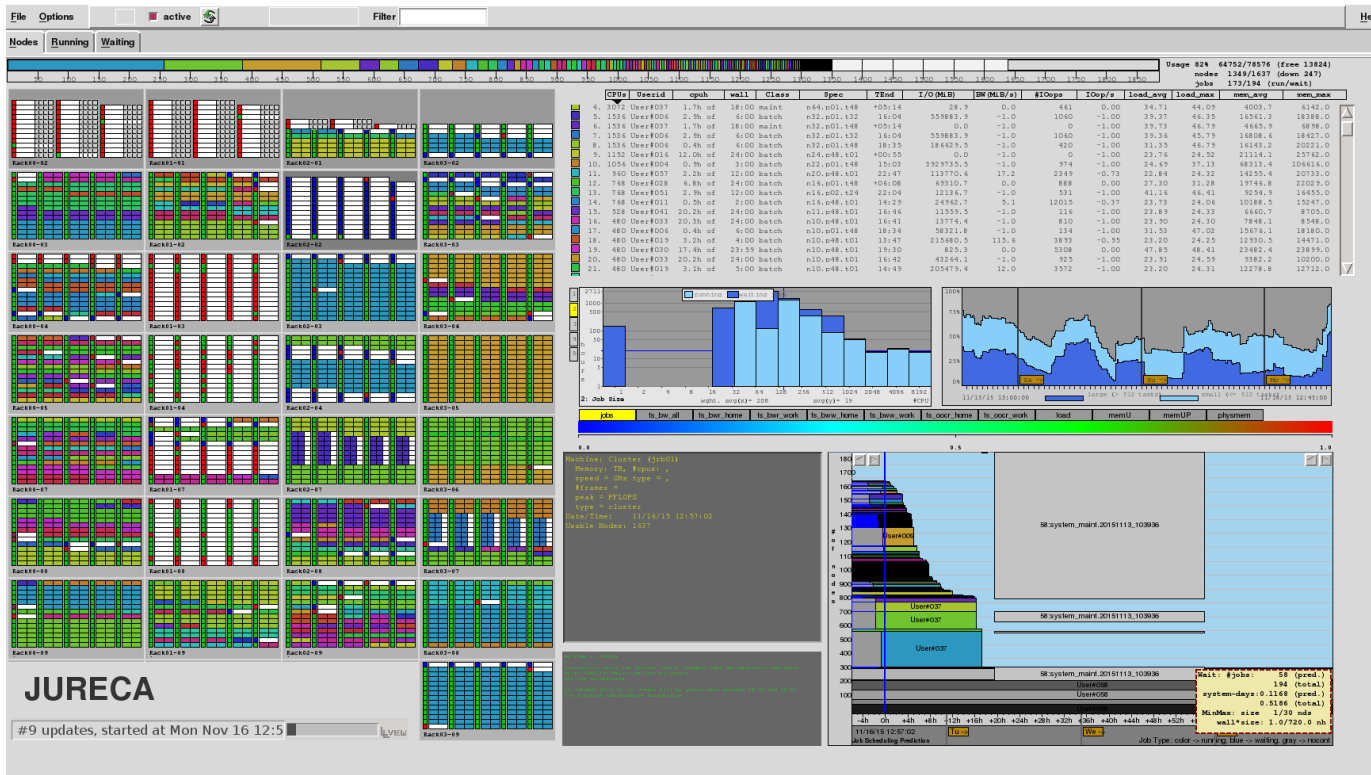
Sebastian Lührs, Jülich Supercomputing Centre (JSC)

24.09.2021



# The history of LLview

- Initially designed to monitor and track individual jobs and full system scheduler status: “HPC visualisation on a single screen”



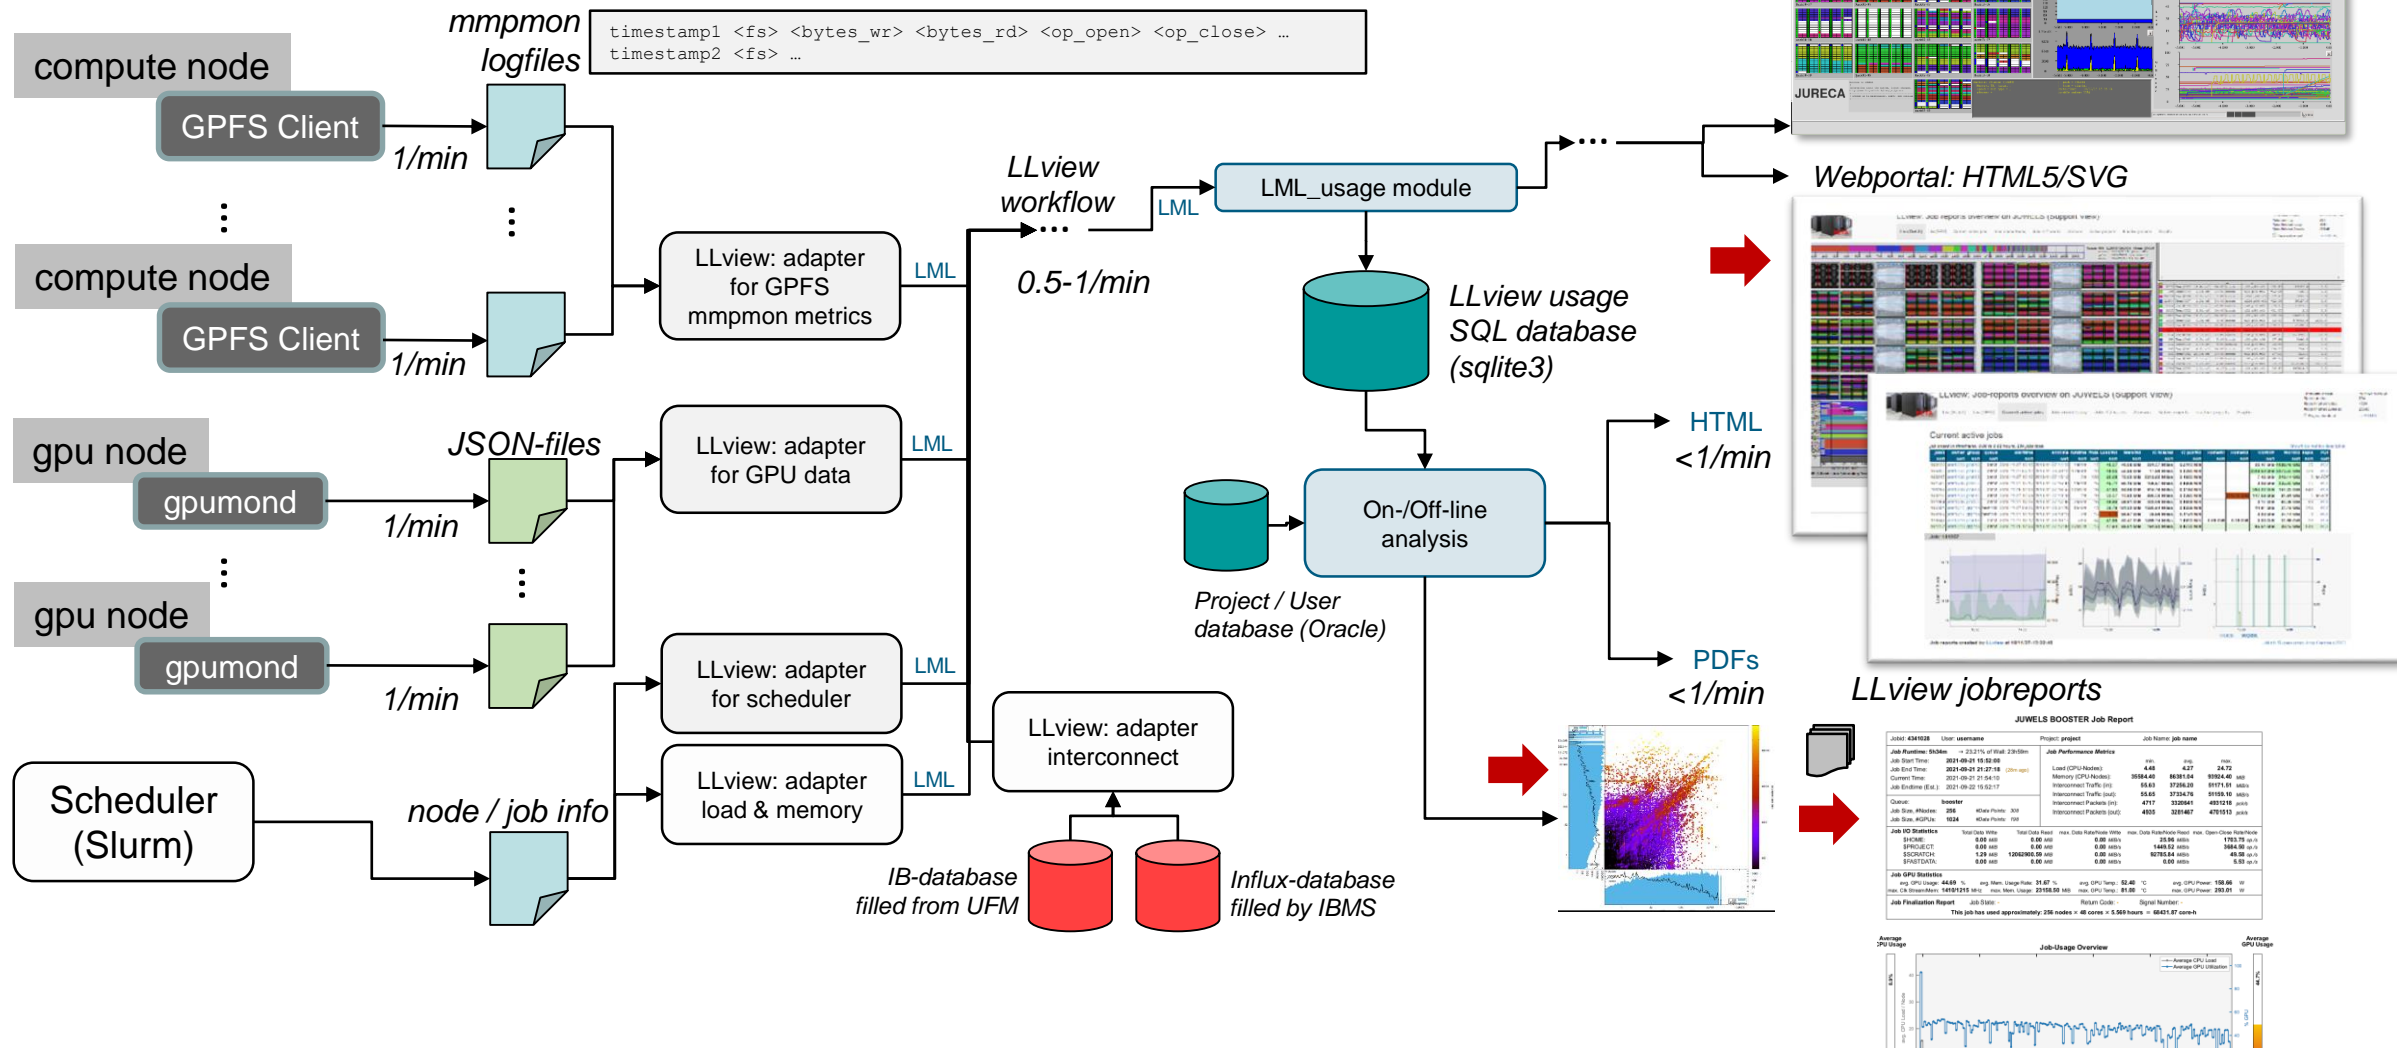
- Configurable, compact and interactive display of all usage data in a single window
- Easy access to system status data
- Open source
- Perl-based client
- Available for all main JSC systems
- <http://www.fz-juelich.de/jsc/llview>



# Targets of LLView based Job-Reporting

- Web-based access to data
  - Role-based data access for users, PIs, mentors and support
- Near real-time monitoring
- Non-intrusive monitoring
- Scalability
- Utilization of available sources for performance data: Load and memory on node, traffic over interconnect, I/O activity, GPU info, ...
- History of job usage and possibility of deeper analysis

# Data sources and data flow



# Job monitoring



## llview Job-reports on JUWELS - User view

Scheduler overview | Live | Queue | Current active jobs | Jobs ended today | Jobs < 3 weeks | Job overviews

ended	jobid	owner	project	queue	starttime	est_endtime	runtime	#nds	Load/Nd	Mem/Nd	IC MiB/Nd	IC Pck/Nd	HomeWr	HomeRd	PrjWr	PrjRd	ScrWr	ScrRd	FDataWr	FDataRd	state	err	#sps
1h23m				batch			15h59m				405.00	0.10									T_OUT		94 PDF
11h10m				batch			56m				468.22	0.12		0.00			2.81	21037.65			CMPL		57 PDF
12h13m				batch			55m				536.85	0.14		0.00				776.73			CMPL		57 PDF
13h15m				batch			54m				1215.10	0.46		0.00				974.89			CMPL		56 PDF
14h29m				batch			4m	1	24.38	12.27	0.21	0.00					0.69	601.03			CMPL		6 PDF
14h50m				batch			5m	1	29.24	32.87	0.21	0.00					0.56	651.47			CMPL		5 PDF
15h19m				batch			4m	1	6.33	12.97	0.21	0.00		0.00			0.56	548.95			CMPL		5 PDF
15h32m				batch			55m	2	31.62	16.37	0.21	0.00		0.00			2.82	21119.81			CMPL		5 PDF
15h34m				batch			5m	1	24.89	13.25	0.00	0.00					0.56	661.40			CMPL		6 PDF
17h57m				batch			21h31m	3	24.30	48.23	1519.94	0.42		0.01			47.56	290843.60			CMPL		1292 PDF


Avg. load & max. memory consumption

Total I/O traffic

Total network traffic

Error states

Job report



Load & memory consumption


Network traffic

I/O traffic



# GPU monitoring

GPUuseAvg	GPUuseMax	GPUmemAvg	GPUmemMax	GPUmemRateAvg	GPUpowAvg	GPUpowMax	GPUtempAvg	GPUtempMax
filter	filter	filter	filter	filter	filter	filter	filter	filter
89.19	92.00	38.98	38.98	96.24	276.69	299.88	57.58	61.00
56.00	100.00	11.23	16.85	22.58	178.90	262.03	53.92	59.00
51.17	87.00	5.90	8.85	18.67	151.99	239.75	51.83	58.00
63.83	100.00	11.23	16.85	26.00	187.41	266.87	53.75	61.00
59.25	100.00	11.90	17.85	23.58	176.91	256.23	54.33	60.00
52.50	90.00	5.90	8.85	20.00	160.75	246.47	52.92	59.00
55.83	95.00	5.90	8.85	21.42	149.07	250.05	52.92	58.00
55.17	99.00	11.90	17.85	22.17	157.27	244.81	53.75	61.00
56.75	100.00	11.23	16.85	23.17	178.83	273.06	53.42	61.00



# Job reports



Jobid: <b>4288631</b> User: <b>username</b>		Project: <b>project</b>		Job Name: <b>job name</b>	
<b>Job Runtime: 15h08m</b> → 63.11% of Wall: 23h59m		<b>Job Performance Metrics</b>			
Job Start Time:	<b>2021-09-03 23:45:59</b>		min.	avg.	max.
Job End Time:	<b>2021-09-04 14:55:03</b> (29m ago)	Load (CPU-Nodes):	<b>34.45</b>	<b>48.00</b>	<b>48.80</b>
Current Time:	2021-09-04 15:23:06	Memory (CPU-Nodes):	<b>0.00</b>	<b>76851.58</b>	<b>92264.40</b> MiB
Job Endtime (Est.):	2021-09-04 23:46:05	Interconnect Traffic (in):	<b>2.80</b>	<b>923.52</b>	<b>7736.99</b> MiB/s
Queue:	<b>batch</b>	Interconnect Traffic (out):	<b>16.53</b>	<b>793.50</b>	<b>7106.87</b> MiB/s
Job Size, #Nodes:	<b>6</b> #Data Points: 811	Interconnect Packets (in):	<b>1068</b>	<b>84528</b>	<b>847428</b> pck/s
Job Size, #GPUs:	<b>0</b> #Data Points: 0	Interconnect Packets (out):	<b>1902</b>	<b>83659</b>	<b>846460</b> pck/s
<b>Job I/O Statistics</b>	Total Data Write	Total Data Read	max. Data Rate/Node Write	max. Data Rate/Node Read	max. Open-Close Rate/Node
\$HOME:	<b>0.00</b> MiB	<b>0.00</b> MiB	<b>0.00</b> MiB/s	<b>0.07</b> MiB/s	<b>76.82</b> op./s
\$PROJECT:	<b>9.74</b> MiB	<b>2974.23</b> MiB	<b>0.05</b> MiB/s	<b>30.64</b> MiB/s	<b>579.60</b> op./s
\$SCRATCH:	<b>0.00</b> MiB	<b>0.00</b> MiB	<b>0.00</b> MiB/s	<b>0.00</b> MiB/s	<b>0.13</b> op./s
\$FASTDATA:	<b>0.00</b> MiB	<b>0.00</b> MiB	<b>0.00</b> MiB/s	<b>0.00</b> MiB/s	<b>0.13</b> op./s
<b>Job Finalization Report</b>	Job State: <b>FAILED</b>	Return Code: <b>15</b>	Signal Number: <b>0</b>		
<b>Node System Error Report</b>	# Msgs: <b>1</b>	# Nodes: <b>1</b> (Out-of-memory)	→ detailed list of error messages at end of report		
<b>This job has used approximately: 6 nodes × 48 cores × 15.134 hours = 4358.59 core-h</b>					




# Job reports

Average CPU Usage

Job-Usage Overview

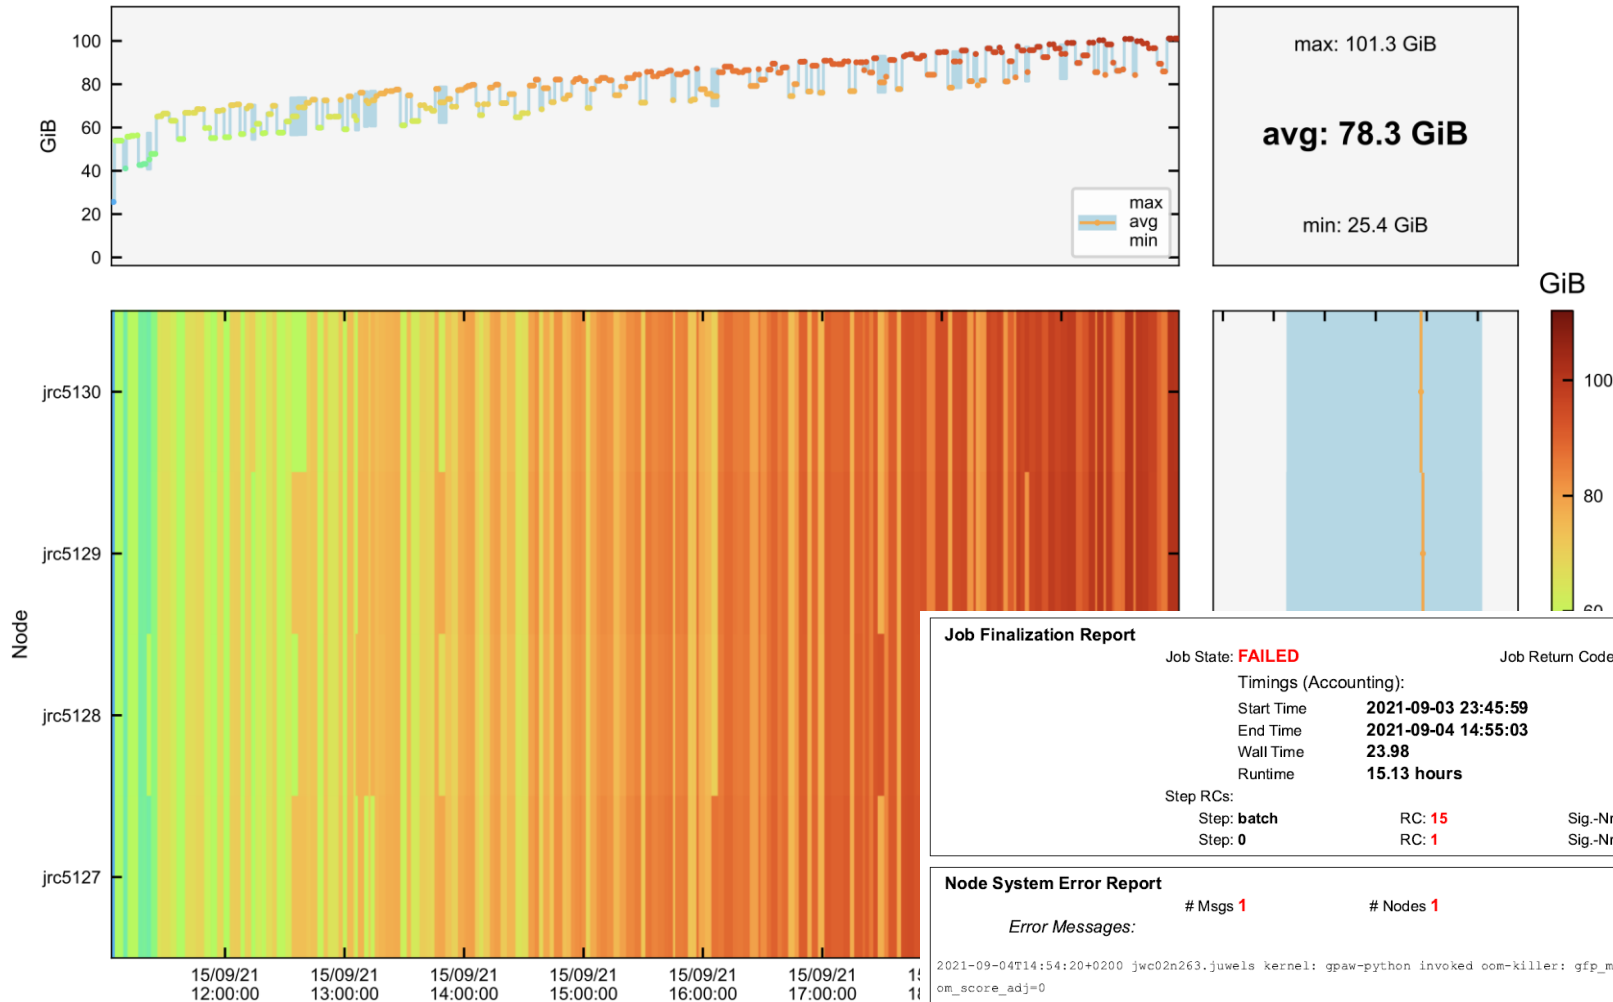
Average GPU Usage





# Job reports

CPU: Memory Usage



**Job Finalization Report**

Job State: **FAILED**      Job Return Code: **15**      Job Signal Number: **0**

Timings (Accounting):

Start Time	<b>2021-09-03 23:45:59</b>
End Time	<b>2021-09-04 14:55:03</b>
Wall Time	<b>23.98</b>
Runtime	<b>15.13 hours</b>

Step RCs:

Step: <b>batch</b>	RC: <b>15</b>	Sig.-Nr: <b>0</b>
Step: <b>0</b>	RC: <b>1</b>	Sig.-Nr: <b>0</b>

---

**Node System Error Report**

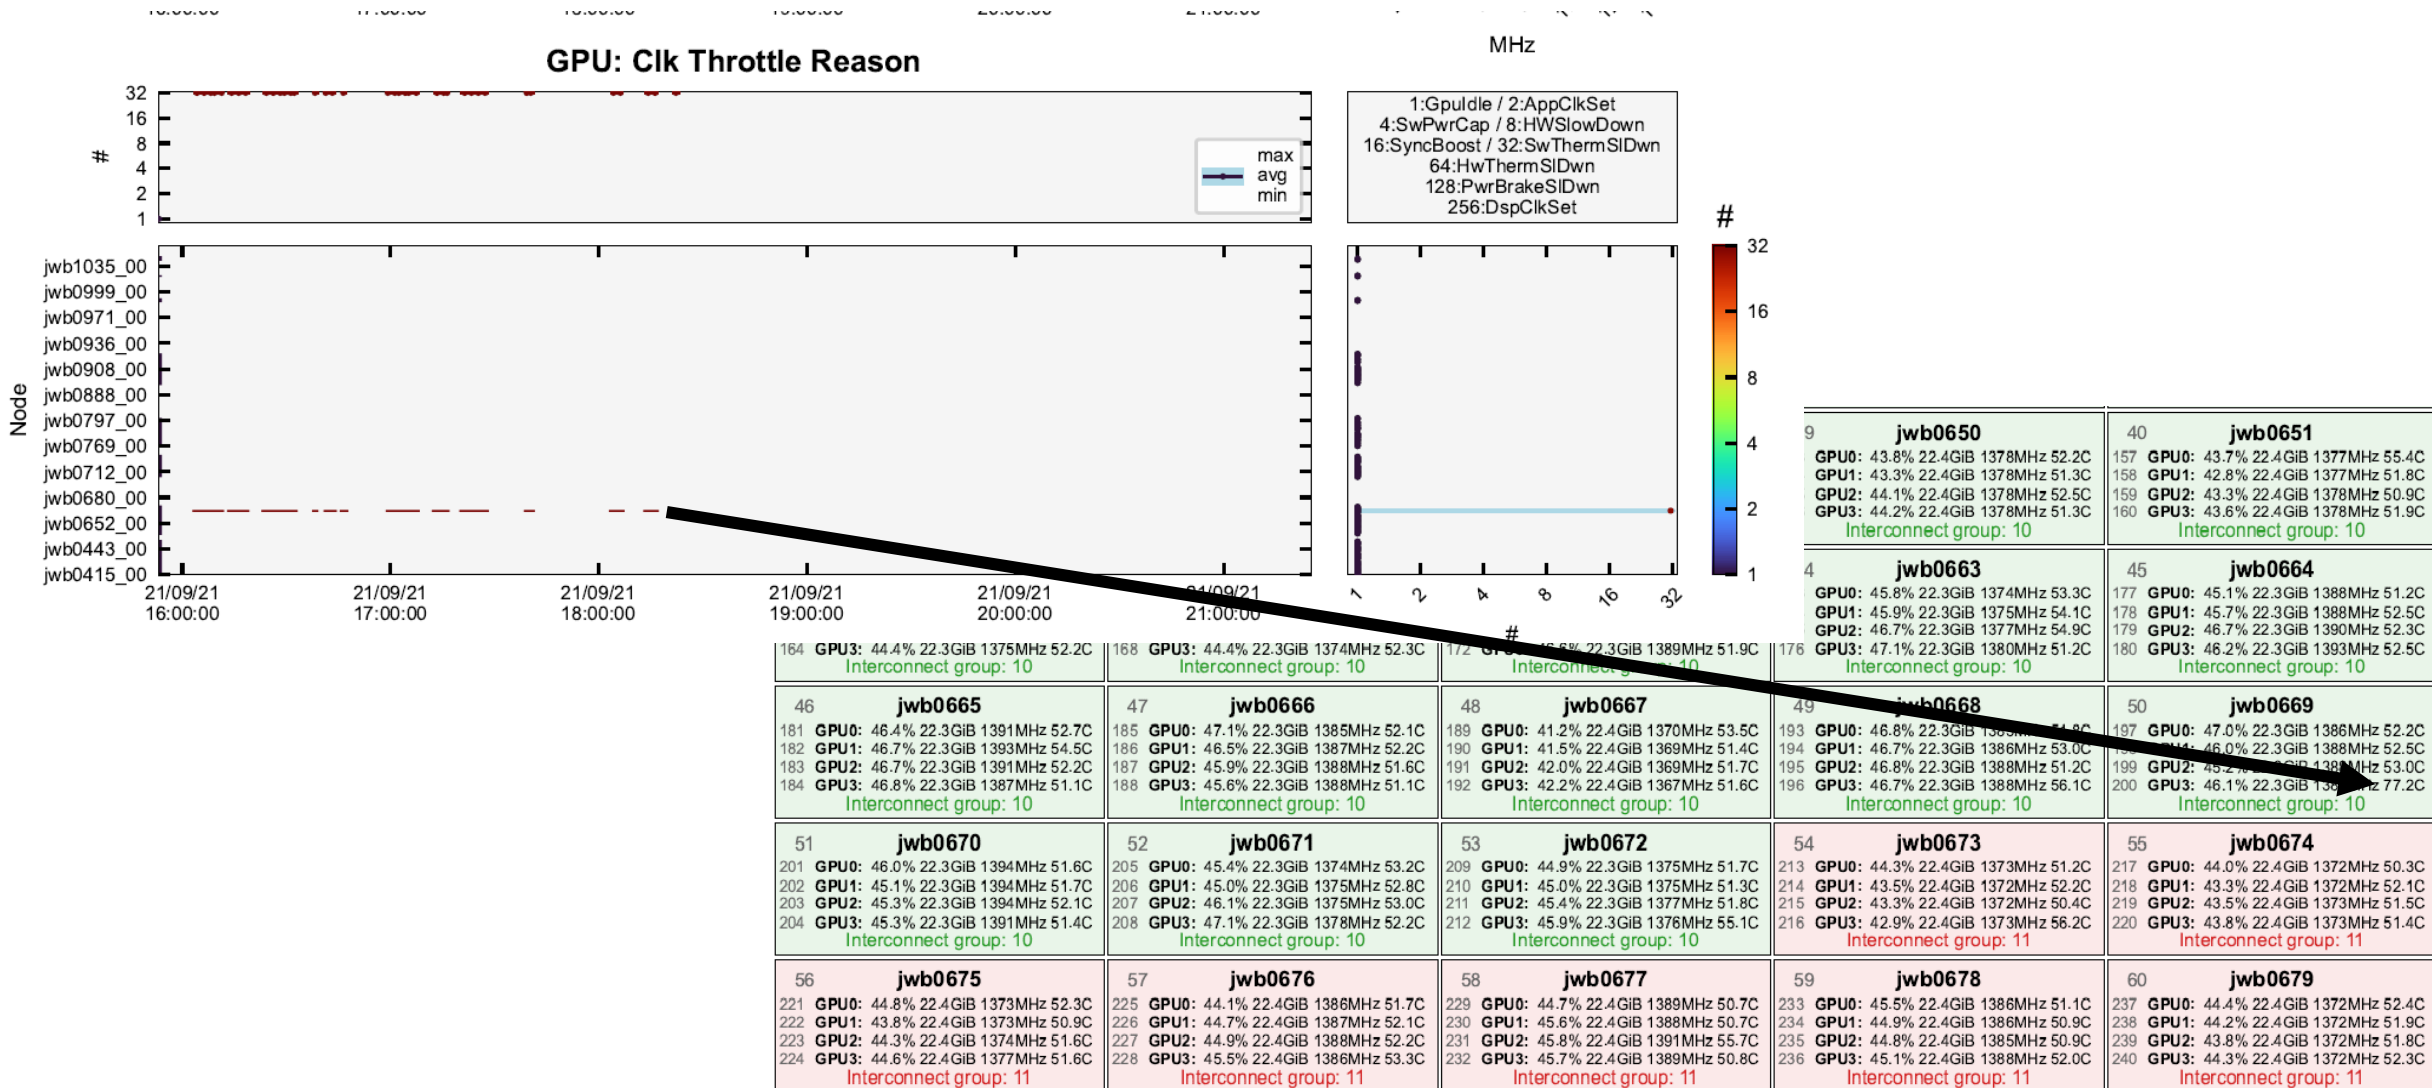
# Msgs **1**      # Nodes **1**

*Error Messages:*

```
2021-09-04T14:54:20+0200 jwc02n263.juwels kernel: gpaw-python invoked oom-killer: gfp_mask=0x6280ca(GFP_HIGHUSER_MOVABLE! __GFP_ZERO), order=0, oom_score_adj=0
```

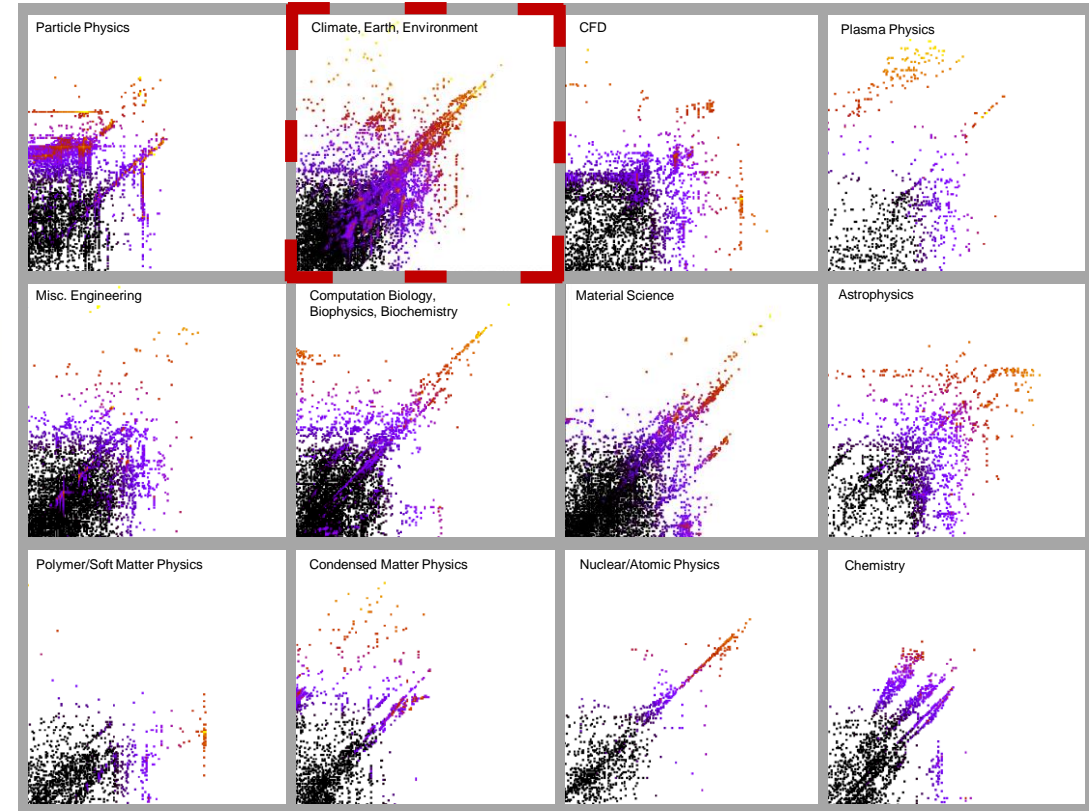



# Job reports

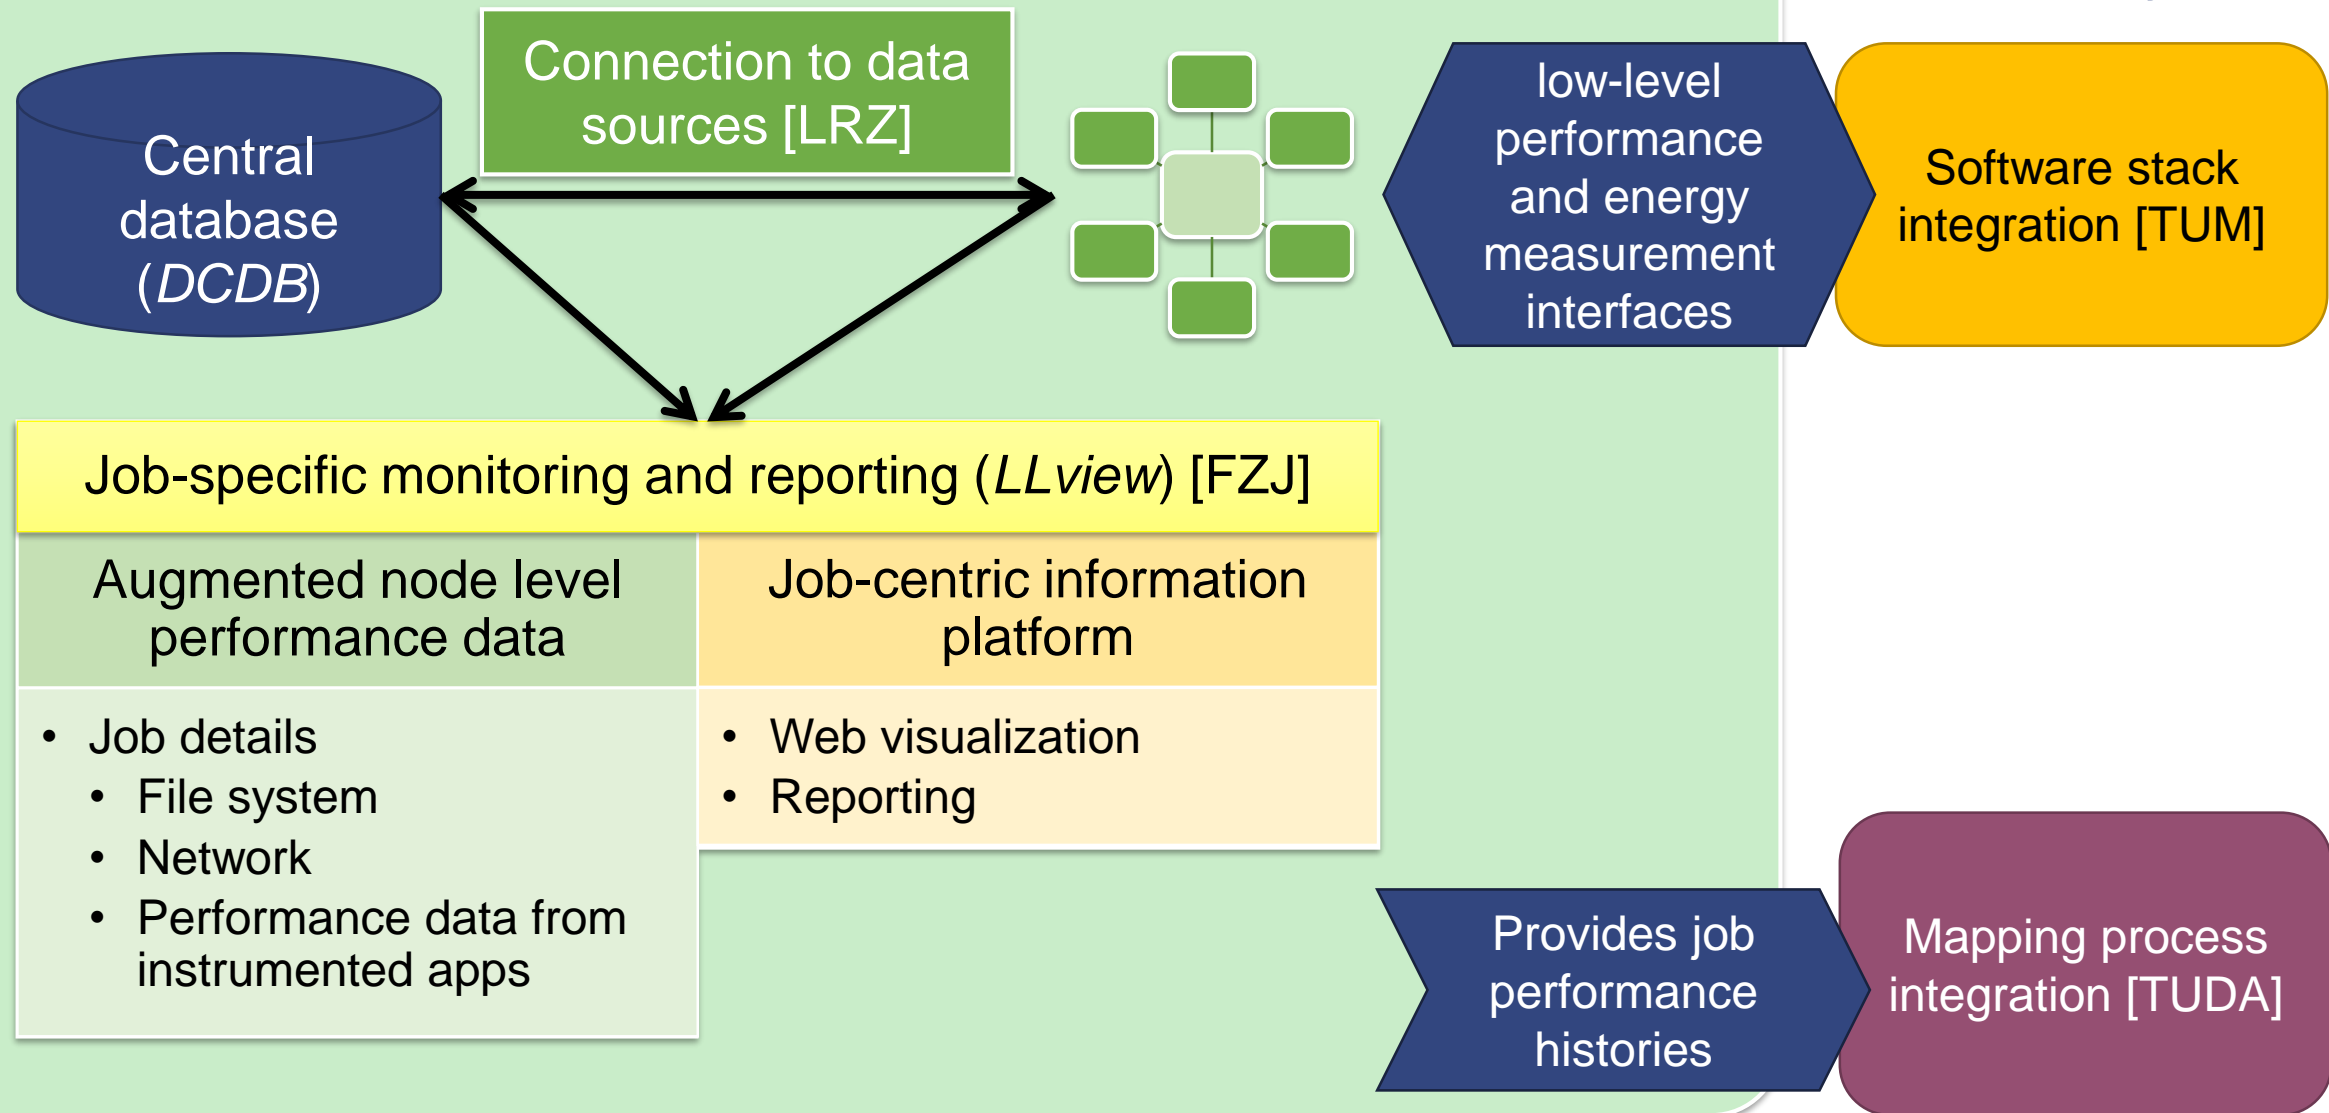


# Offline workload analysis (I/O activity of jobs)

I/O activity by research topic



# DEEP SEA: Application and system monitoring platform





[www.deep-projects.eu](http://www.deep-projects.eu)



@DEEPprojects

# **DEEP** *Projects*



*The DEEP Projects have received funding from the European Commission's FP7, H2020, and EuroHPC Programmes, under Grant Agreements n° 287530, 610476, 754304, and 955606*