

Parallel I/O Monitoring at JSC

Wolfgang Frings

Jülich Supercomputing Centre

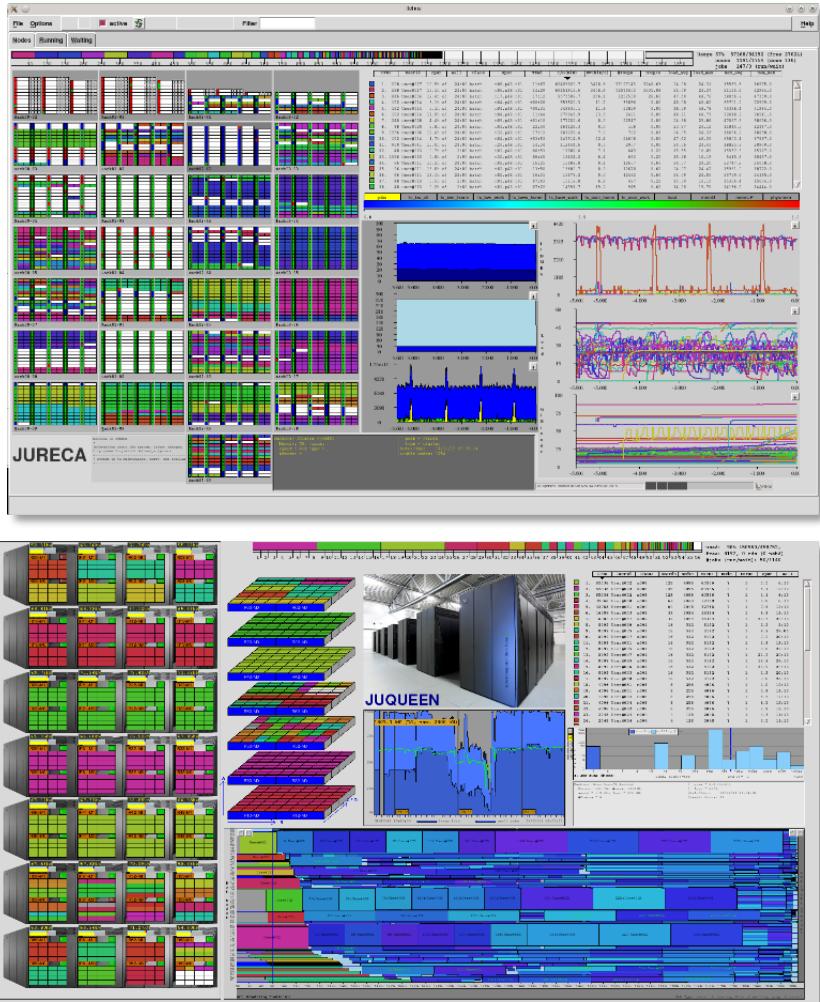
BoF@SC15: Analyzing Parallel I/O

November 2015



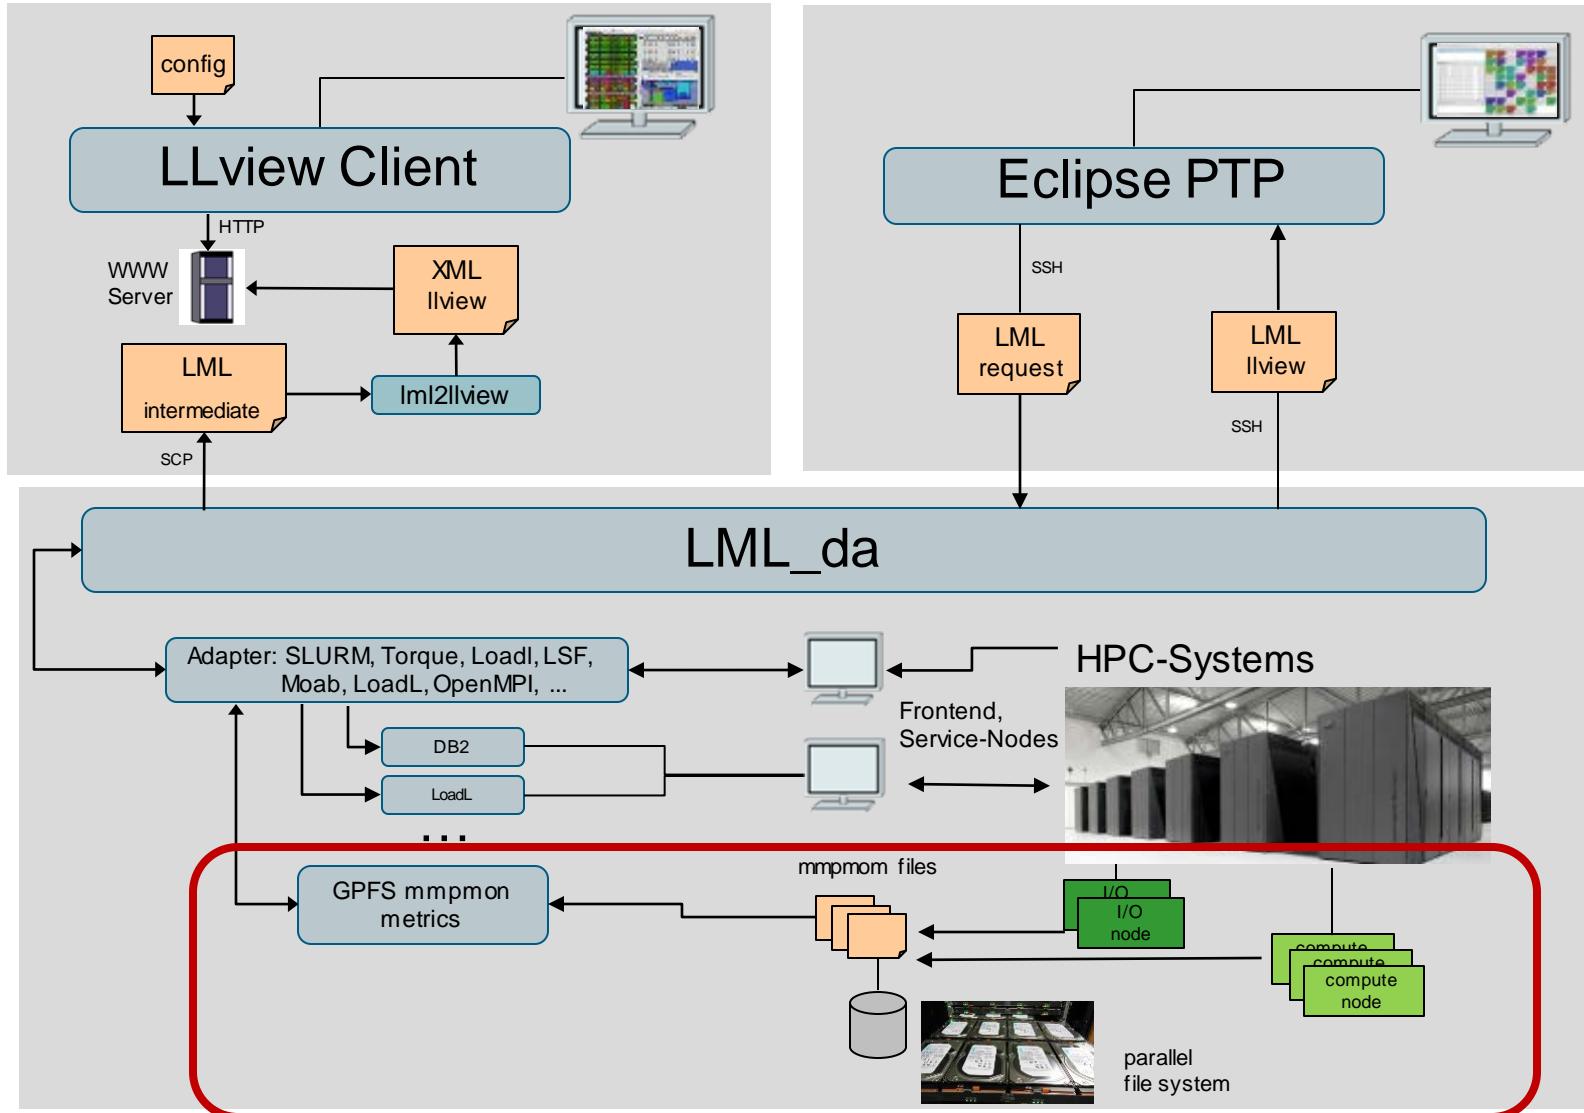
LLview: User-level Monitoring

- Efficient supervision of node usage, running jobs, statistics, history
- Prediction of system usage
- Monitoring of energy consumption, load, memory usage, I/O usage
- Interactive and mouse-sensitive
- Main data source: batch scheduler, runtime system
- No interaction with compute nodes
- Fully customizable, fast and portable client-server application
- Integrated into Eclipse/PTP
- Support for various resource manager, incl. LoadLeveler, IBM Blue Gene, Cray ALPS, PBSpro, Torque, SLURM, Grid Engine and LSF

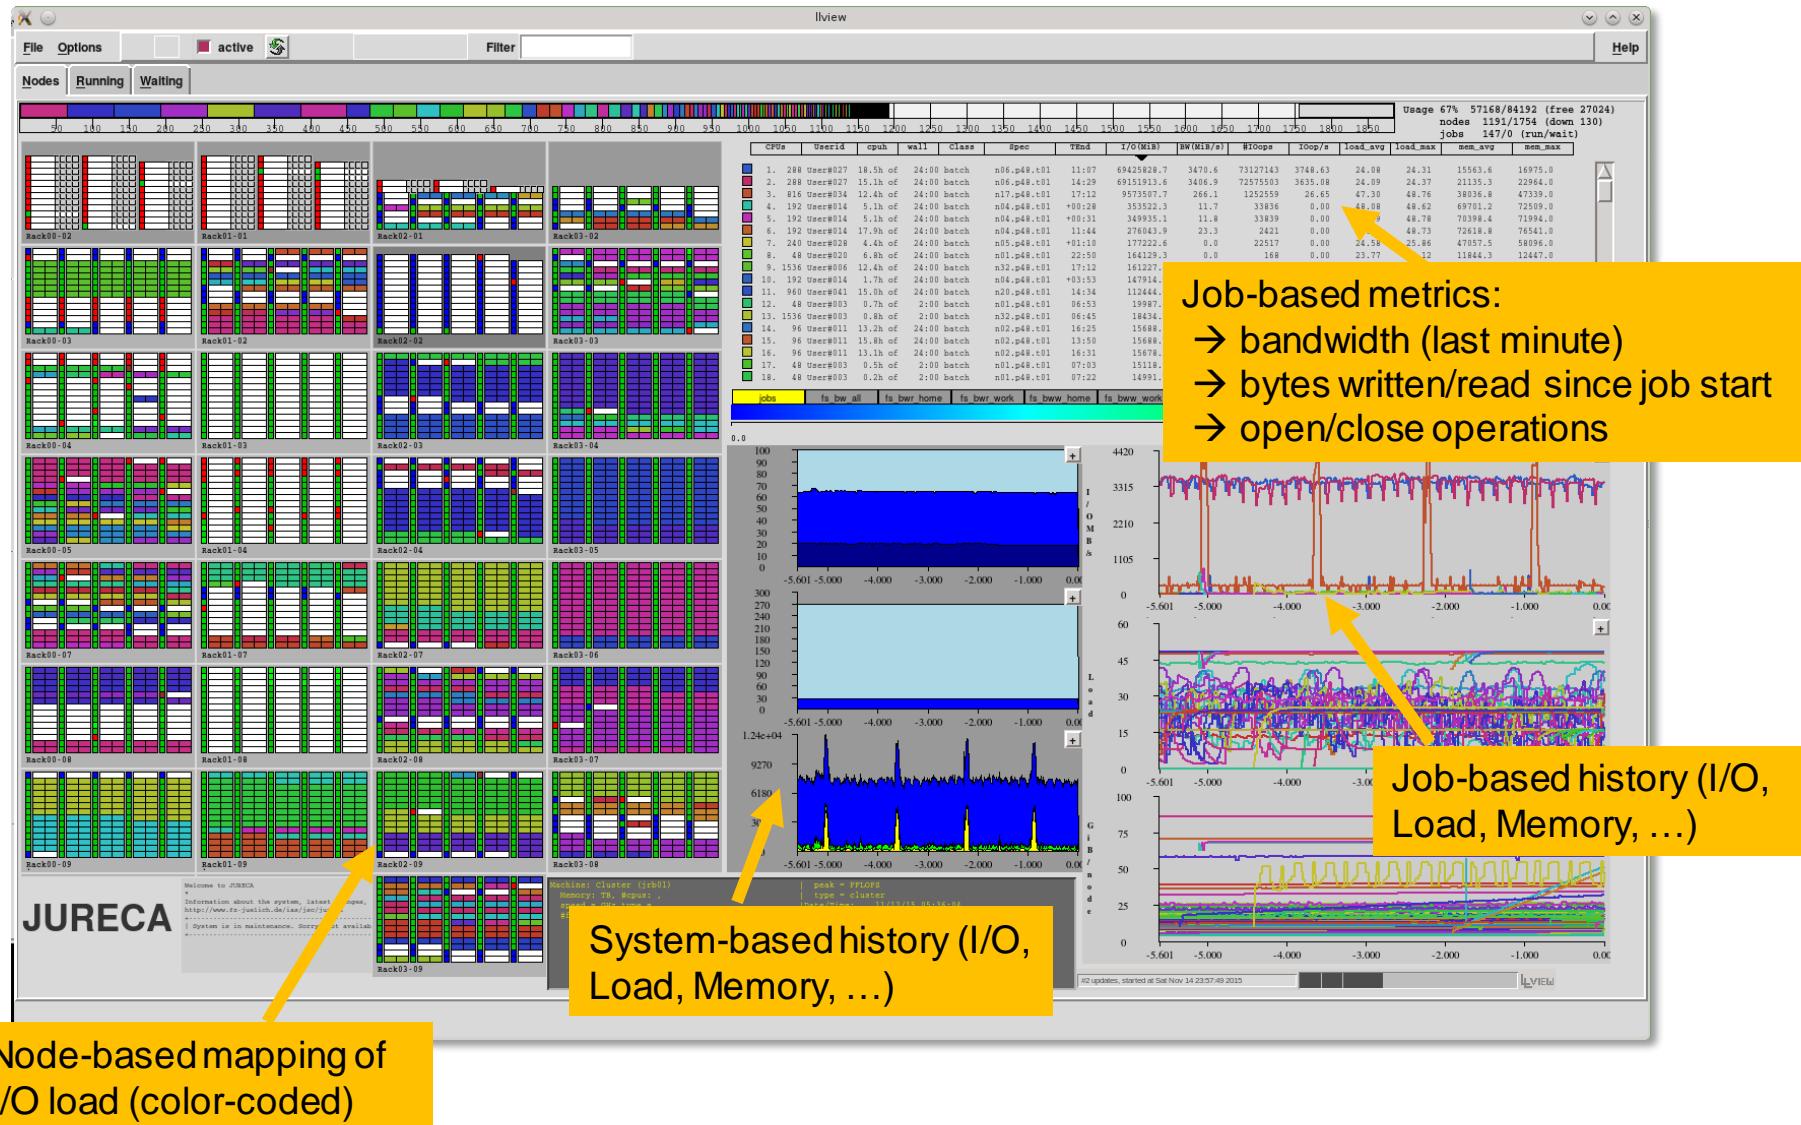


LLview download: (open source)
<http://www.fz-juelich.de/jsc/llview>

LLview Architecture & I/O monitoring

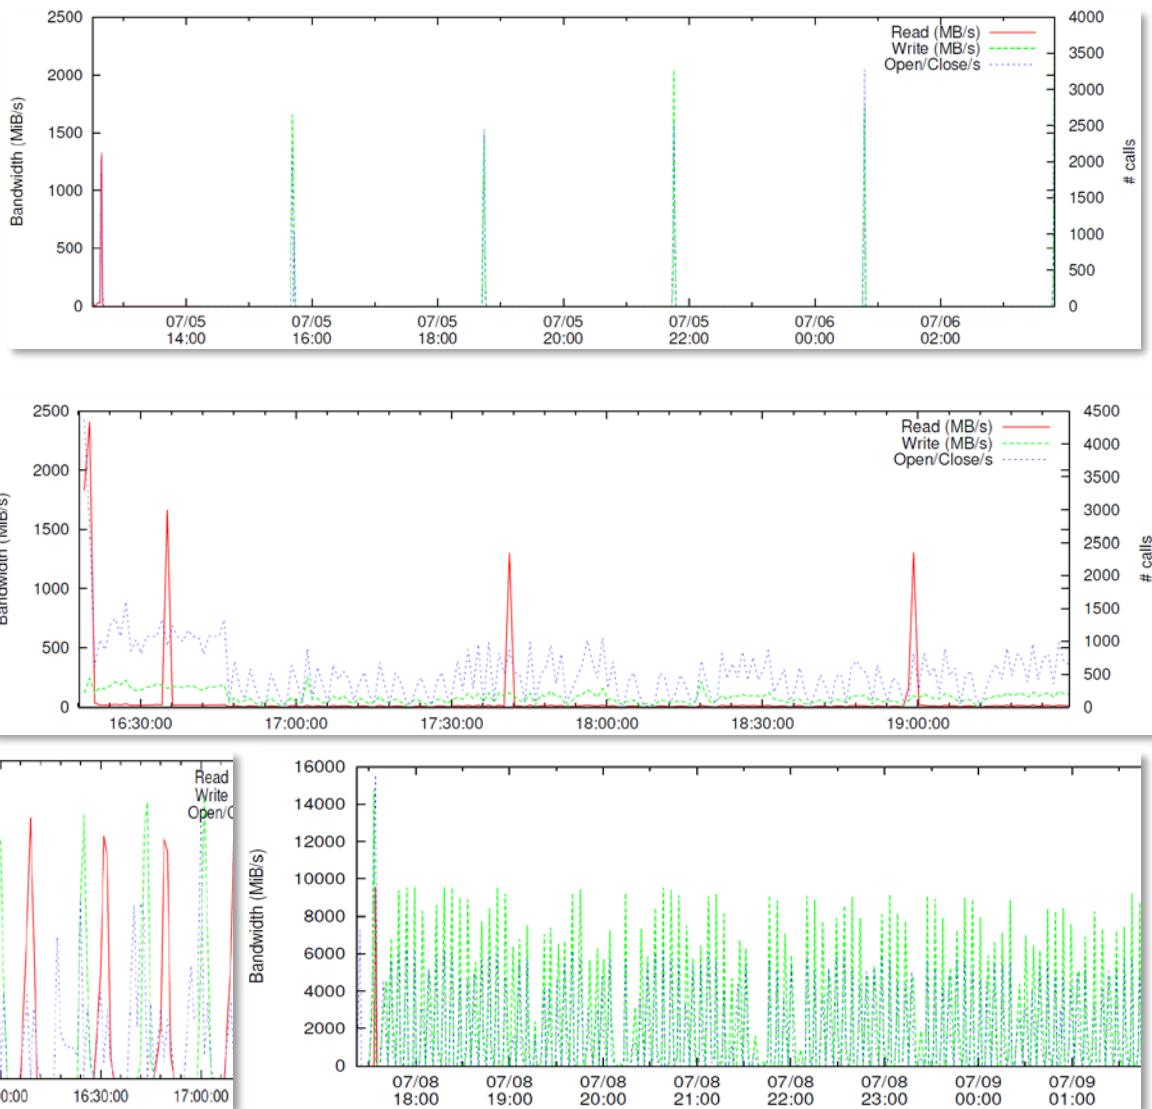


LLview: Node & File System Metrics



I/O Monitoring: GPFS MMPMON

→ Examples:
Job-based
I/O-Usage
on JUQUEEN



Conclusion

- Design of LLview: Batch system monitoring
- New: Combination of job-to-node mapping with additional data sources:
 - Node-based I/O information (e.g. GPFS mmpmon)
 - Node metrics (e.g. load, memory usage from Slurm)
→ Application-based monitoring without instrumenting or modifying application
- Next:
 - Provide users job-based information (LLview, PDF-Report)