




T2K-ND280 Computing Model

Thomas Lindner (TRIUMF) for ND280 Computing Group

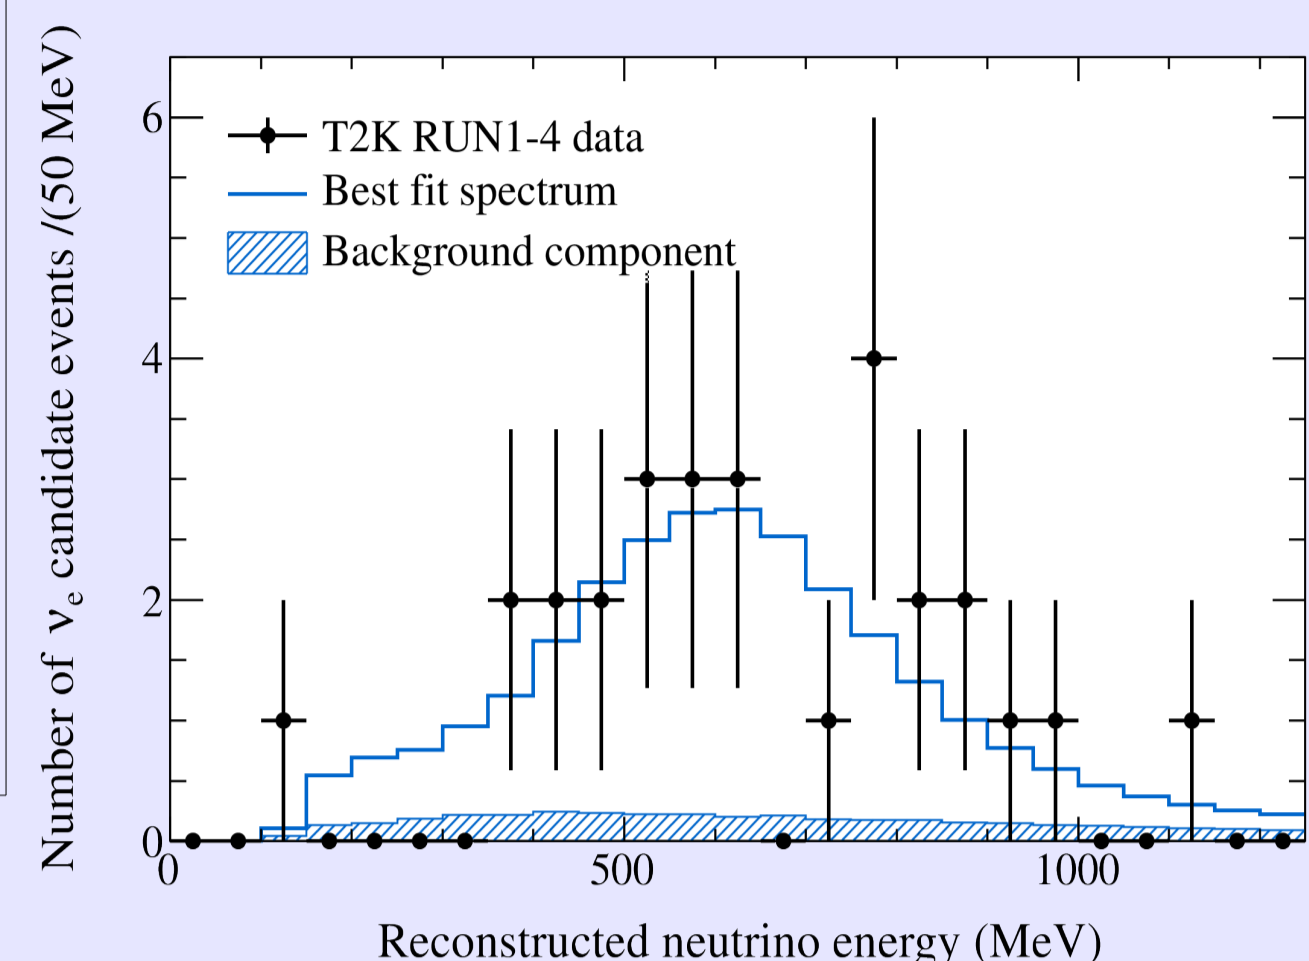


T2K experiment

- Long baseline off-axis neutrino oscillation experiment.
- High purity beam of ν_μ produced at JPARC proton accelerator in Tokai, Japan.
- Measure the rate of ν_μ and ν_e 295 km away at the Super-Kamiokande water Cherenkov detector.

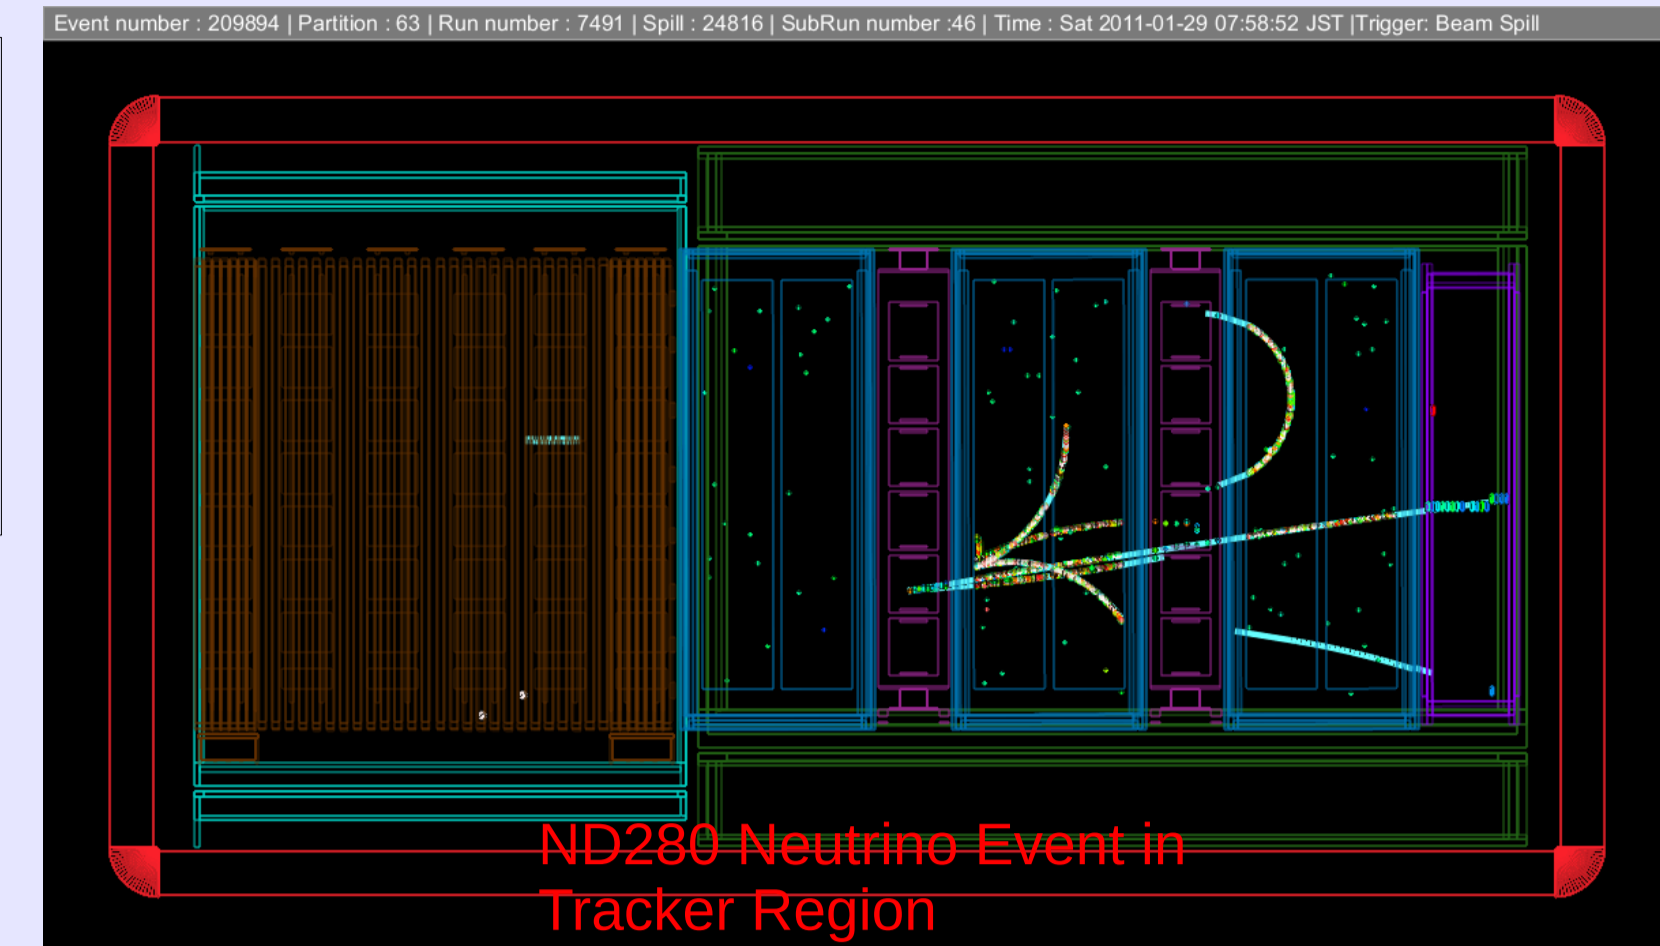

• Allows extraction of neutrino oscillation parameters:

- $\theta_{13}, \theta_{23}, \Delta m^2_{32}, \delta_{CP}$
- Recently presented first observation of ν_e appearance in ν_μ beam (7.5σ effect).



ND280 Detector At JPARC

- Magnetized neutrino detector, 280m from ν target.
- Critical for characterizing neutrino beam flux and composition before oscillations.
- Technology: uMEGAS TPC + scintillator with MPPC.



ND280 Software

- Standard HEP software suite using Geant4 for simulation and ROOT for analysis and I/O.
- NEUT and GENIE for neutrino interaction simulation.
- Custom reconstruction, with RecPack toolkit.
- CVS, CMT and bugzilla for code management and development.

Data Processing and MC Production

Data processing and MC production occurs around world:

Europe:

- 7 sites in UK (with 35% at RAL-LCG2)
- Standard EMI/LCG tools: WMS job submission to CREAM-CEs
- output to local SE, FTS replication back to RAL-SE.



North America:

- Bugaboo and Scinet clusters in Canada (batch queues)
- Colorado CMS Tier-3 Cluster (batch queue)
- Files transferred back to TRIUMF-SE with lcg-cp and globus-url-copy.
- Last major production (for Summer 2013 results) used ~2500 core-years (HEPSC06 normalized).
- Typically have 1 production per year.

Data Distribution & Management

Principal Data Storage

- A pair of large data storage elements are core of ND280 data management:
- RAL-LCG2: ~500TB allocation at RAL Tier-1 (GridPP-supported)
- CA-TRIUMF-T2K: ~500TB SE, provided/maintained solely by T2K-Canada group.
- These SE provide interface between LCG-Europe and non-LCG-America.



Raw Data Distribution

- ~100TB (including replicas)
- Data copied from JPARC to primary storage element at KEK Computing Center.
- Replicated with FTS to RAL and TRIUMF storage elements.

Production Data Distribution

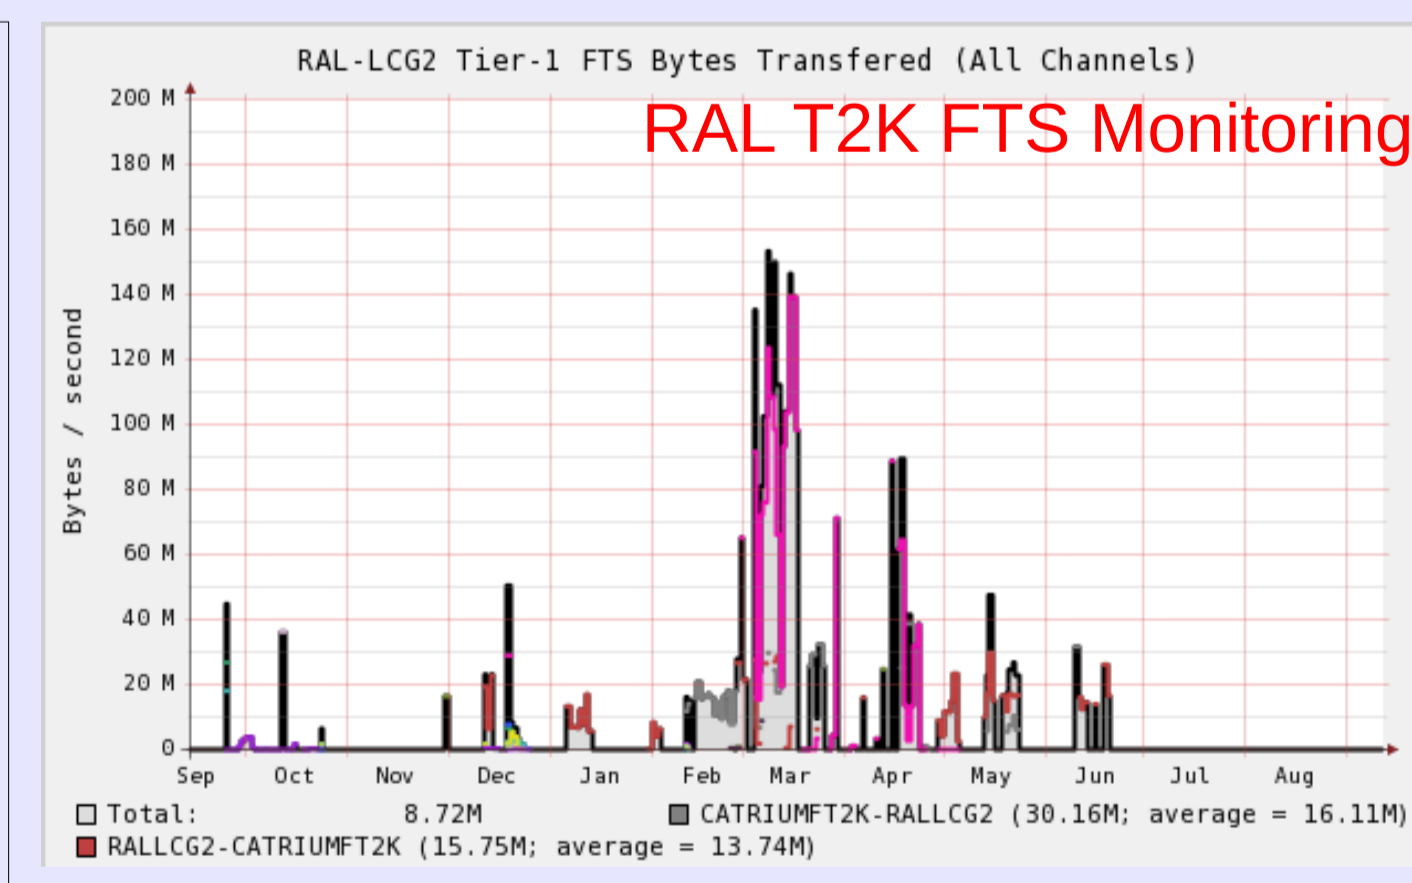
- ~1PB (including replicas)
- Data distribution decoupled from production jobs.
- Data from production sites copied to RAL or TRIUMF
- Then fully synchronized between RAL/TRIUMF with FTS
- Users download the summarized data set (~10TB) for analysis.

Data Management

- Use LFC for global view of data (t2k.org VO)
- Custom MySQL file catalogue for details about files

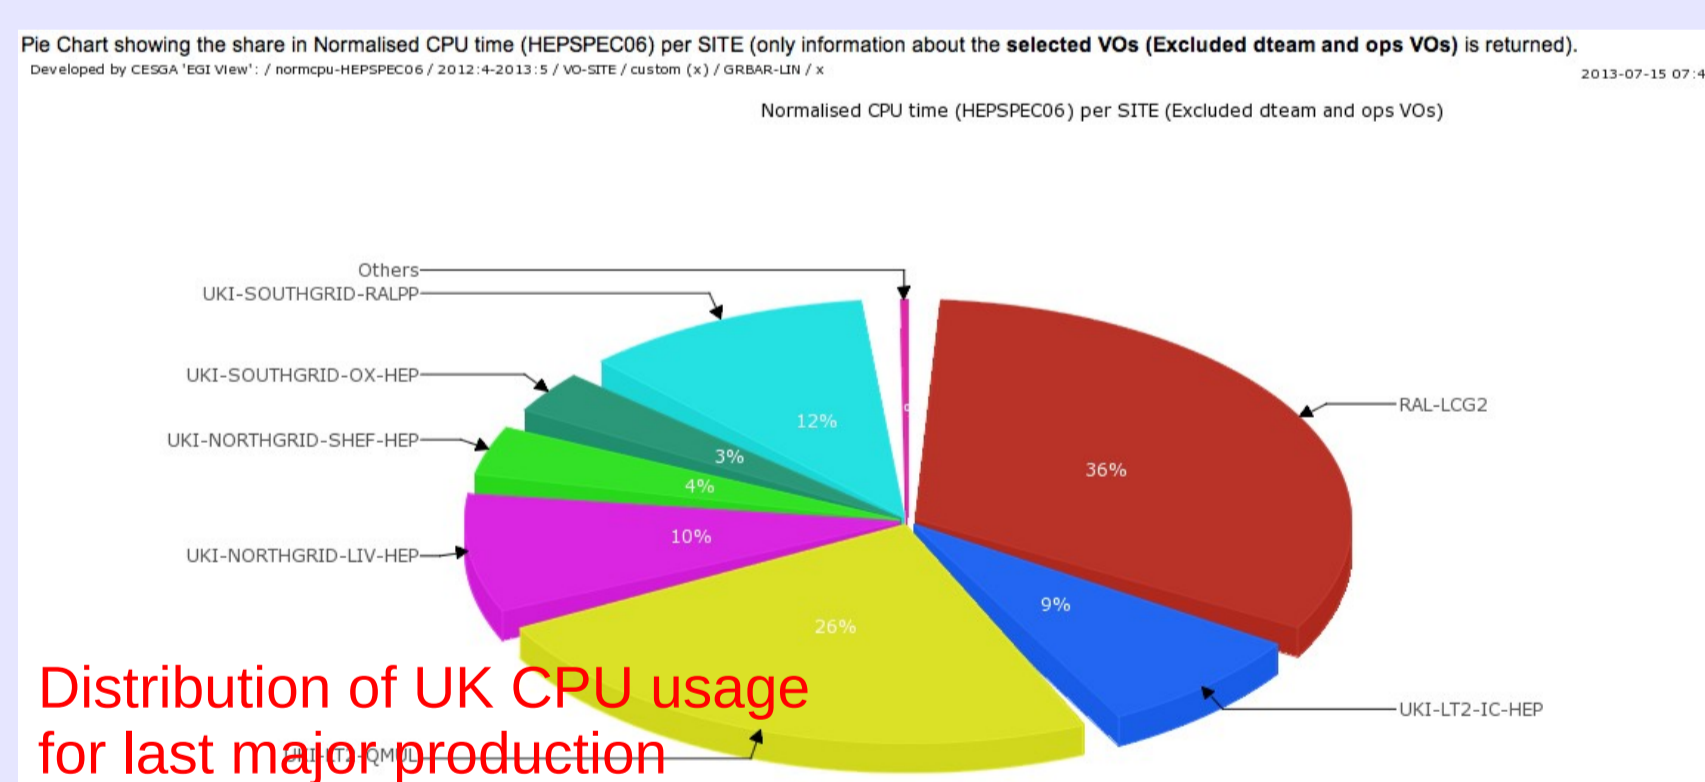
Key Data Challenges:

- Large data set to manage: ~6% of ATLAS data set in Canada.
- FTS transfers don't provide LFC registering → can cause "dark data"
- Summary data set are still very large; difficult for analyzers.



Key Production Challenges:

- Production work is relatively manpower intensive, because need different skill sets for work at different sites.
- Standard LCG toolset has presented many difficulties, in particular with WMS and Proxy Management.
- North American production is often limited by the rate at which output files can be moved back to TRIUMF SE.
- Limited grid support in many institutions and minimal dedicated computing manpower.
- No unified view of production status.



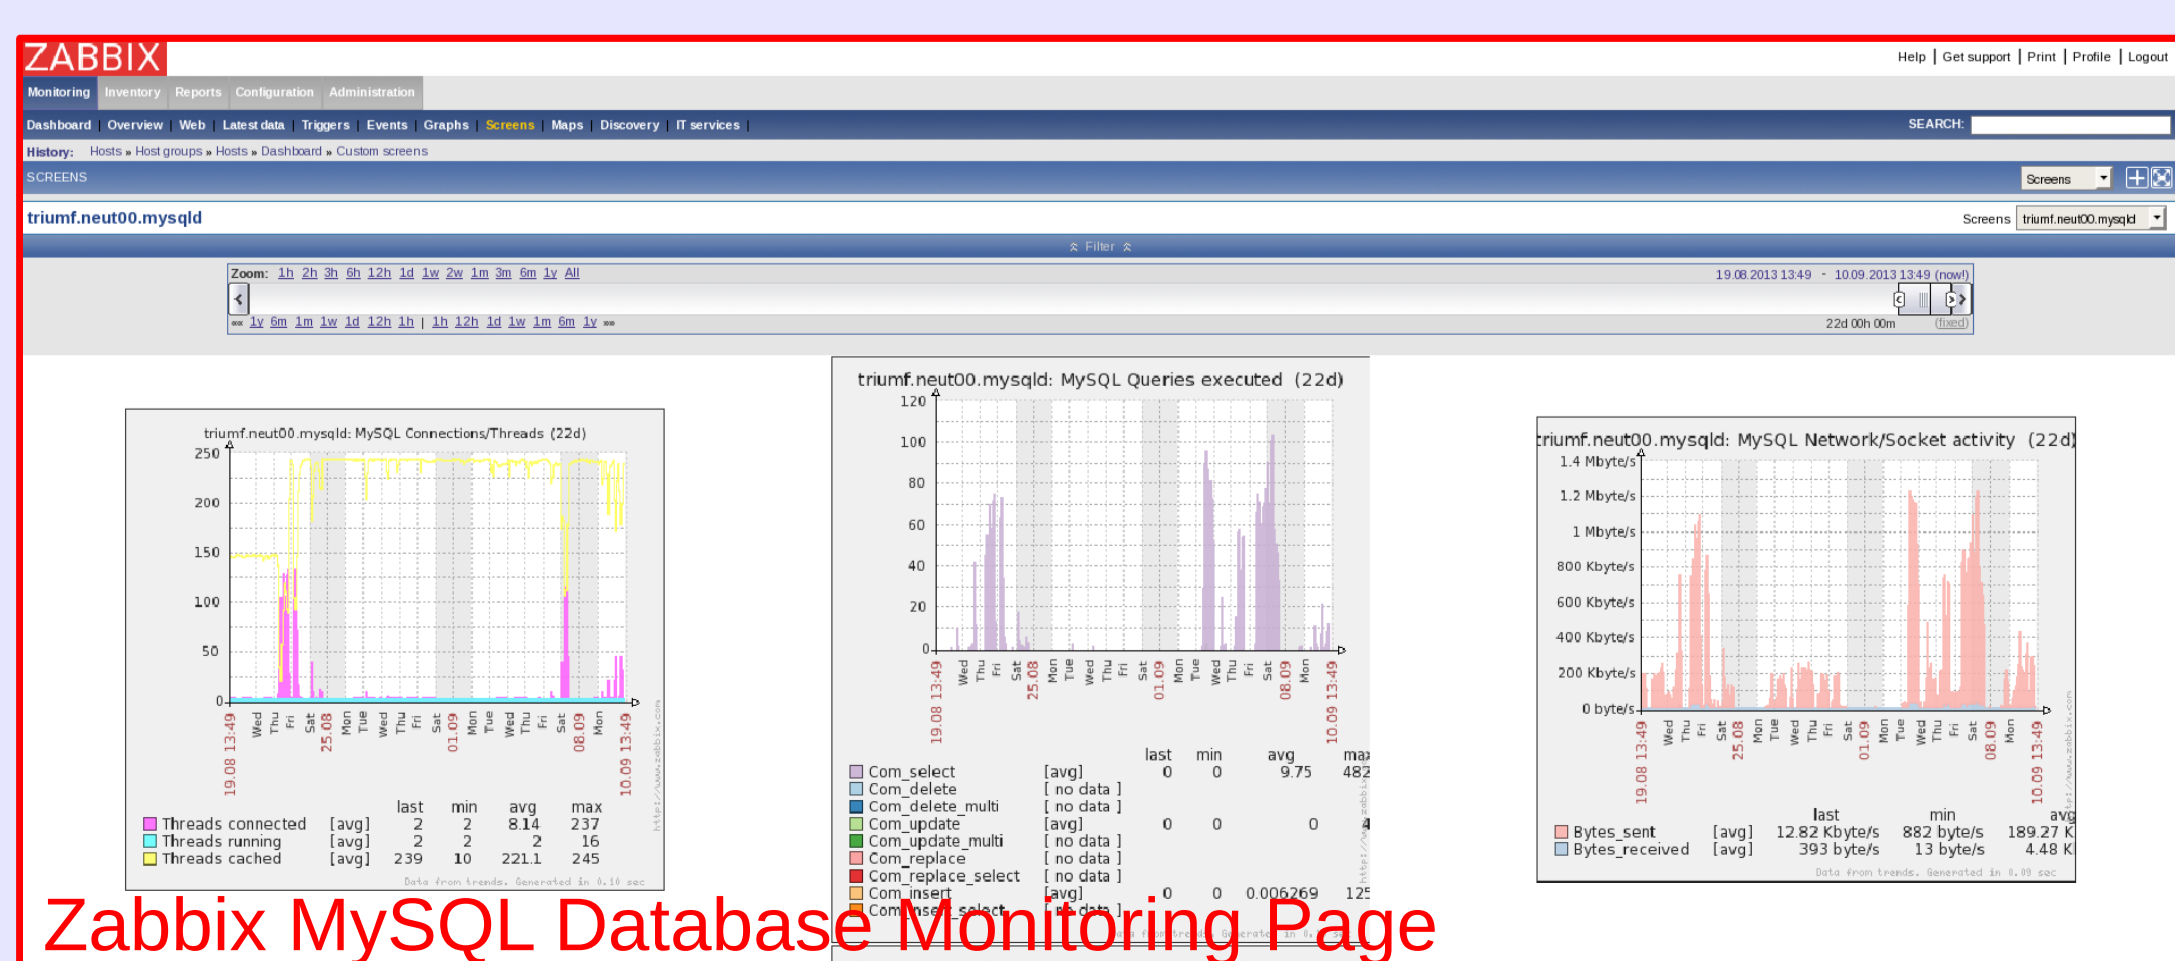
ND280 Databases

ND280 uses a set of MySQL databases for following purposes:

- 1) Slow control information (temperatures, voltages, etc) [70 GB]
- 2) Calibration constants and data quality [5 GB]
- 3) File Metadata Catalogue [3 GB]

Each database has different schema and C++ interface classes.

Master databases are hosted at JPARC and TRIUMF, with some replication to Europe. Databases monitored with Zabbix.



Future Directions

T2K has only collected 8% of approved data-set. Will need continued improvements to ND280 software and computing to be able to handle full T2K data-set. Planned improvements:

- Reduction in file sizes; too much of effort is spent on data management.
- Investigating using Canadian resources through LCG interfaces.
- Investigating DIRAC and GANGA as ways of improving the stability/simplicity of production work, as well as unifying regional efforts.
- Switching to FTS3; perhaps automating file replication and registration.
- RAL storage increasing to 1PB next year; probable to have similar increase in TRIUMF storage.
- But need to also understand whether network links between RAL-TRIUMF are sufficient.