# Core – IT / News

- Mailbox move to Exchange2010 finished
- FAIR mail domain fair-center.eu in production
- Home license programm (MS Campus contract)
- New windows domain concept starting with administration
- Plans to build a GSM antenna at GSI to increase mobile reception (especially at KBW) => Vodafone is here today!
- All other projects continue…

# Upgrade Tape Drives March 2012

Tape library in RZ (IBM 3584-L23) :

- 8 new tape drives (IBM 3592-E07)
- 250 MB/s read/write (before 160 MB/s)
- 4 TB/medium uncompressed (before 1 TB/med.)

=> overall max:

- ~160 TB/day read/write bandwidth ATL
- 8.8 PB data capacity currently (before 2.2 PB)
- expandable to ~50 PB

Tape storages costs (for primary) will be recalculated, backup remains at 125 k€/PB

# gStore 2012

offline
Clients

lustre
3 PB

Tape
Roboter

5 GB/s

2 GB/s

0.5 GB/s

capacity 8.8 PB
expandable: 50 PB

gStore Data Movers
200 TB Buffer Storage

online DAQ
Clients

# Core IT – WAN

Dark Fiber-Connection to Frankfurt:

- 12x10 Gb LOEWE CSC <-> GSI Lustre
- 10 Gb Uni FFM <-> GSI Lustre
- Optional Uni Mainz <-> GSI (10-40 Gb)
- Expandable to Terabit
- Using one dark fiber and dense wavelength-division multiplexing technology (DWDM)
- Open for offers now, order May, delivery June/July, **production -> August 2012**

# Core IT – Network Plans for Green IT-Cube

- 2 additional small backbone switches
- all 4 connected by metro ring protocol (failover)
- „old“ modules can be used

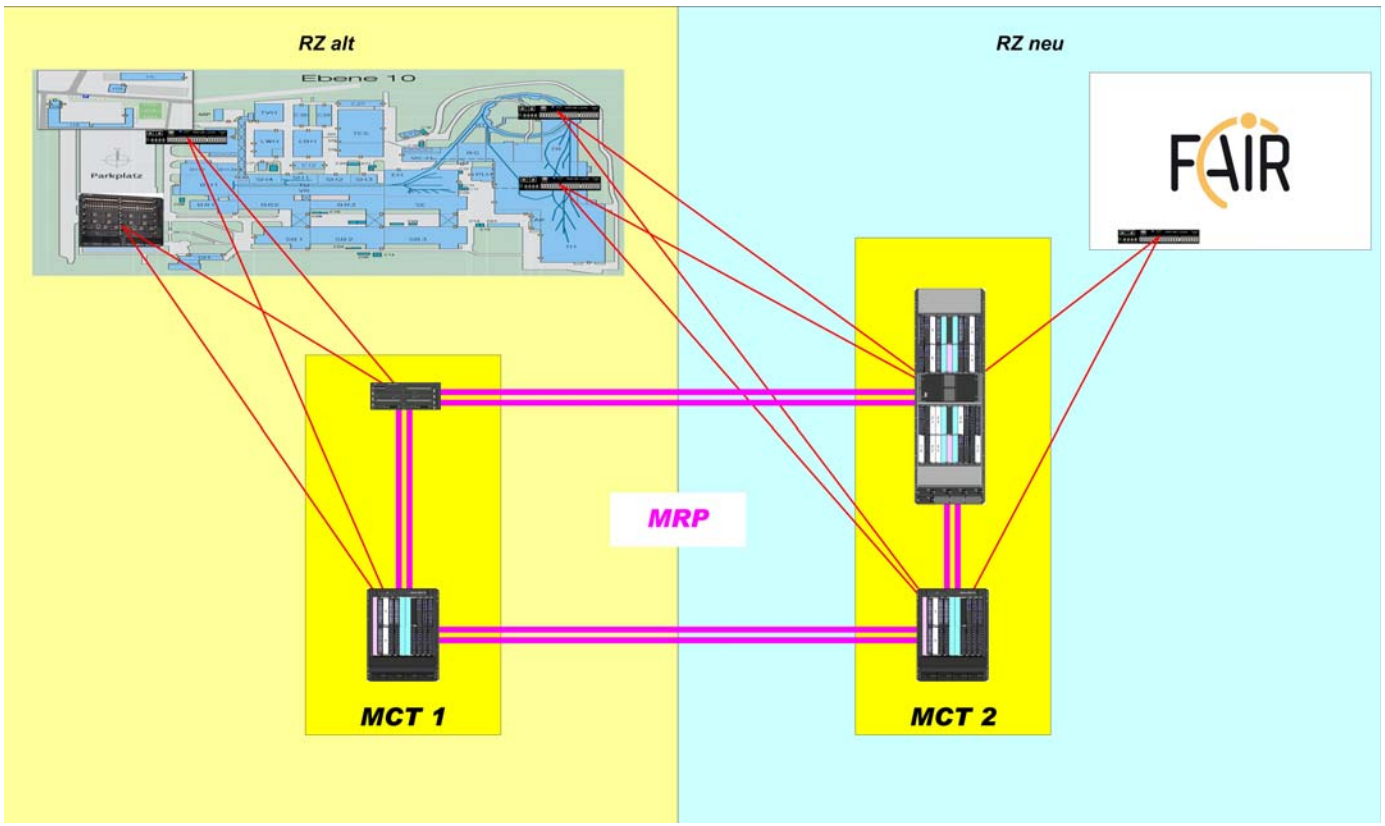

# Administrative IT

- Content management typo3 and new web site – production delayed due to content
- Project tools FAIR/GSI – SAP PS <-> MS Project (PS customzing in progress)
- SAP-ebiss start June EE and Core-IT/July for everybody
- New sub-domain for administration (-> Core-IT)

# Document Management System

- GDS (OnTEAM) insecure, no features…
- EDMS: in future, use will require contribution
- Invenio – optimized for library service
- Evaluation of DMS solutions such as ELO, DOXiS, OpenText, Alfresco…
- Goal – one system for GSI (administration, accelerator and experiments)
  - Feature wish list?
  - Requirements (how many users, file system space…)

Presentation of possible solutions and discussion (May 2012) – who is interested?

# Prometheus / Hera
# - IB cluster in the Minicube

Executive summary: works! Success! :-)

- **Prometheus: 10,000 core cluster, based on Infiniband**
- **Hera: 2,300 disk parallel filesystem (lustre)**
- **IB interconnect + connection to IP world (IP router)**
- **4 LNET router: connecting 6,000 disk lustre system**
- **Deployment of system, configuration, GE, Monitoring of Prometheus and Hera in a couple of days: Thanks to our new „large scale fabric management“ (including boot over IB :-) )**

# Prometheus + Hera – first results

- 200,000 physics jobs already finished,
- 2.8 PB data throughput ….
- ~.5 Tbit/s I/O troughput already reached (in the Hera part)

=> Already (successfully) used by Alice, Beschleuniger-Physik, HPC

# Prometheus + Hera – first results

- IB works reliable and much much faster as Ethernet ...
- Speed measured with Prometheus
  - RDMA per node: 3.1 Gbyte/s (expected: Mellanox: ~3GB)
  - Parallel access to lustre per node: Read ~2 Gbyte/s
- Fileserver performance now limited by the switched backplane...
- core perfomance: constant (about the same as with the older Icarus cluster (16 cores per node)

=> however we have now 24 cores per node …

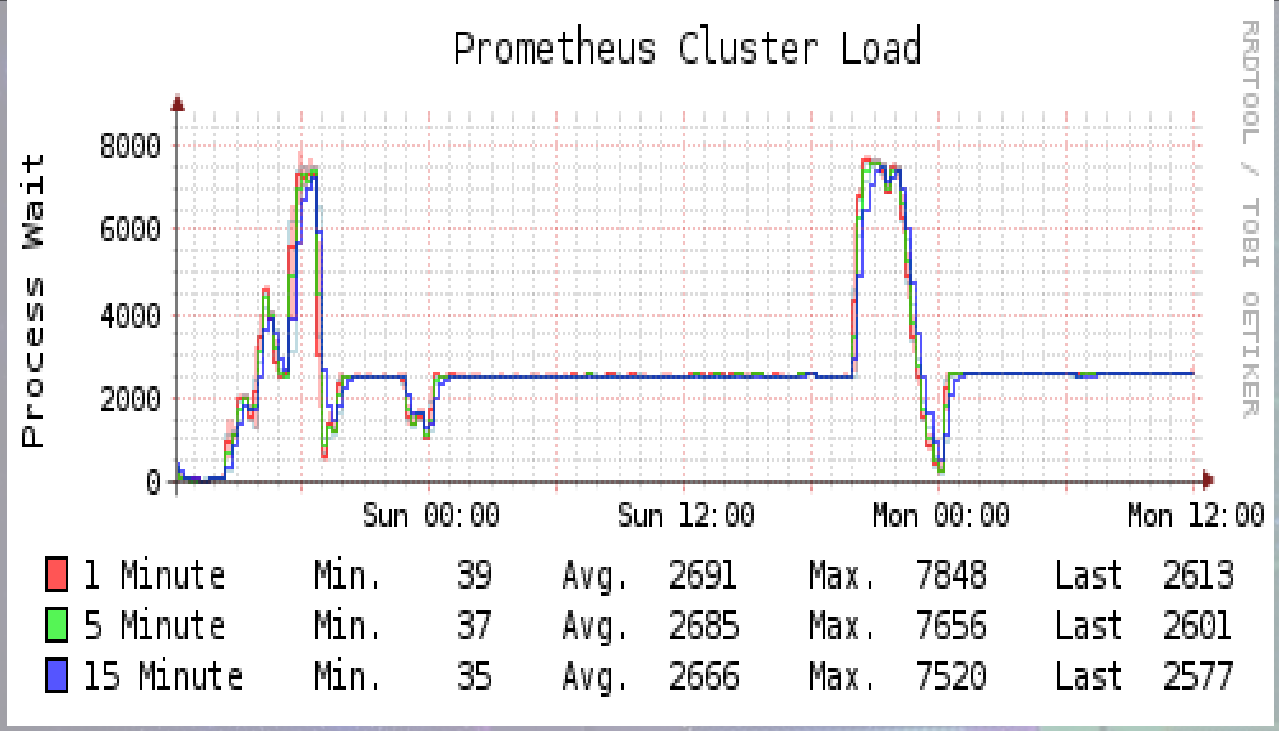


HPC Test
450 Gbit/s
(close to the backplane limit)

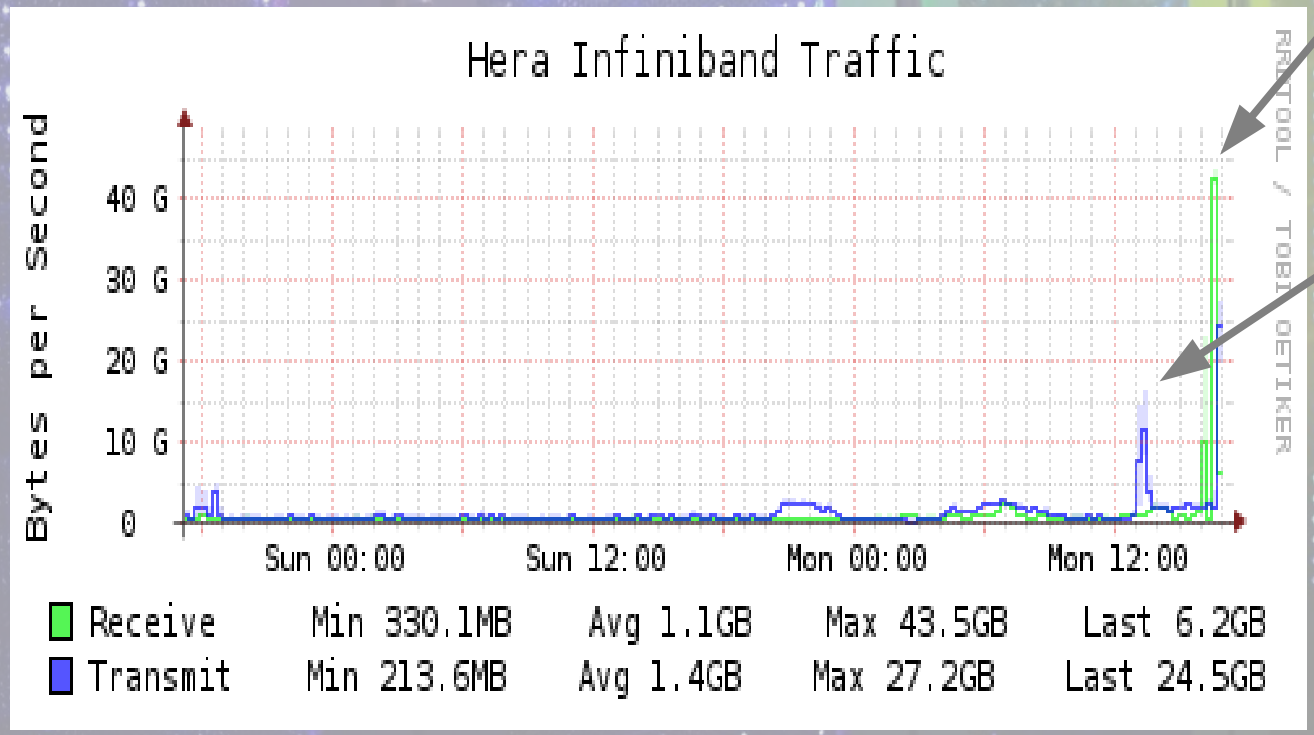


Alice Train,
160 Gbit/s

messaging.gsi.de

prometheus.odp - OpenOffice.org Impress

Report - Iceweasel

File Edit View History Bookmarks Tools Help

http://lxrm02.gsi.de:5000/report?display=gridengine%2Faccounting

Most Visited Getting Started Latest Headlines

Report | Graph | Graph | vpenso/collectd-inter... | Graph

Graph Data Report Help

## GridEngine Accounting Summary of Last Week

Last updated: Mon Apr 16 14:00:52 +0200 2012

Clear

| User | Jobs | CPU Time | Run Time | Memory Used Over Time | I/O | Jobs Max Resident Memory <1GB | <2GB | <4GB | >4GB | Jobs Run Time <1h | <4h | <12h | >12h | Exits not 0 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| cpreuss | 26 | 5m0s | 1h11m | 0.0B/s | 490.72KB | 26 | 0 | 0 | 0 | 26 | 0 | 0 | 0 | 23 |
| vpenso | 5 | 0s | 1m6s | 0.0B/s | 4.15MB | 5 | 0 | 0 | 0 | 5 | 0 | 0 | 0 | 0 |
| root | 11 | 0s | 8m12s | 0.0B/s | 11.98MB | 11 | 0 | 0 | 0 | 11 | 0 | 0 | 0 | 0 |
| sma | 34739 | 8kh | 28kh | 32.02EB/s | 318.2TB | 9853 | 24728 | 64 | 94 | 28241 | 6493 | 5 | 0 | 1 |
| jthaeder | 142336 | 238kh | 250kh | 1.27EB/s | 2.5PB | 128004 | 175 | 14157 | 0 | 83177 | 44686 | 11912 | 2561 | 12827 |

/report/gridengine/accounting?format=html

142,000 jobs

2.5 PB

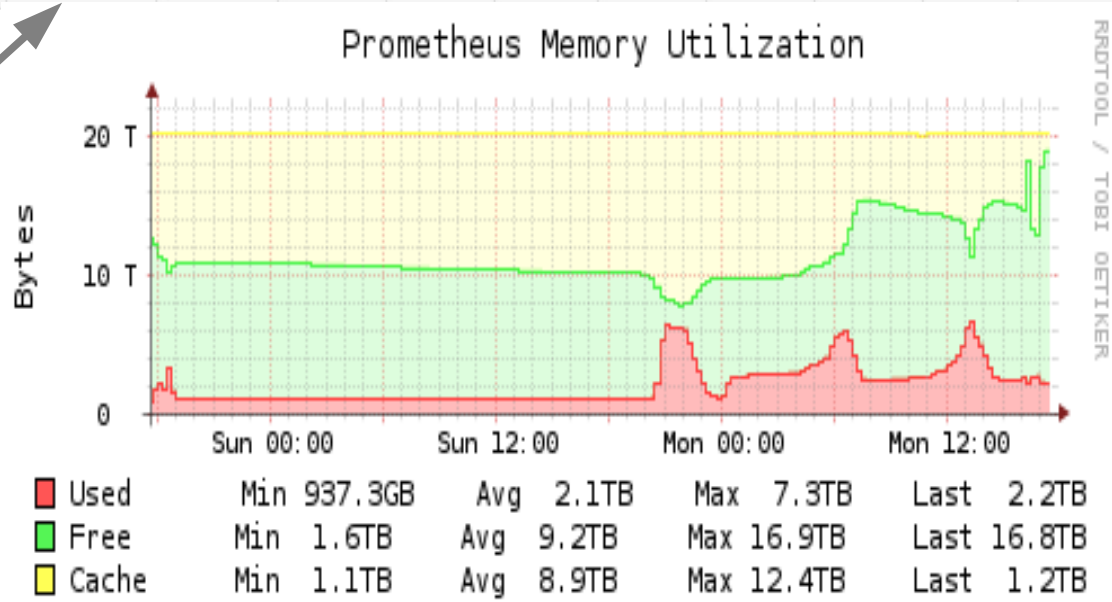

To do:

- One rack is missing (PDU issue => Infrastructure)
- LNET router instability at the Ethernet side
  => HPC, under investigation
- ALICE: migrate data from /lustre => /hera
- HPC: migrate “Icarus” cluster to Prometheus
  - Decommissioning of old, aged, unreliable lustre file servers
  - Buy, install new file servers at minicube
  - Building + testing a LNET router cluster for the WAN lustre Tera Scale project

# Focus of Scientific Computing:

Subgroups: Experiment Simulation and Analysis, e-Infrastructure, DBs

- Development, maintenance and support of the common software framework of the FAIR experiments (FairRoot)
- Participation in national and international research - and development activities for software of the experimental and theoretical particle - and nuclear physics; Community software support at GSI / FAIR
- Support of the software environment of the ALICE experiment at the German Tier-1 center GridKa at KIT and the Tier-2 center at GSI, LOEWE CSC, ...
- Strategic planning and design of the e - Infrastructure for the FAIR experiments, development of models for the resource integration of regional universities, national and international FAIR partners with methods of the Grid and Cloud Computing

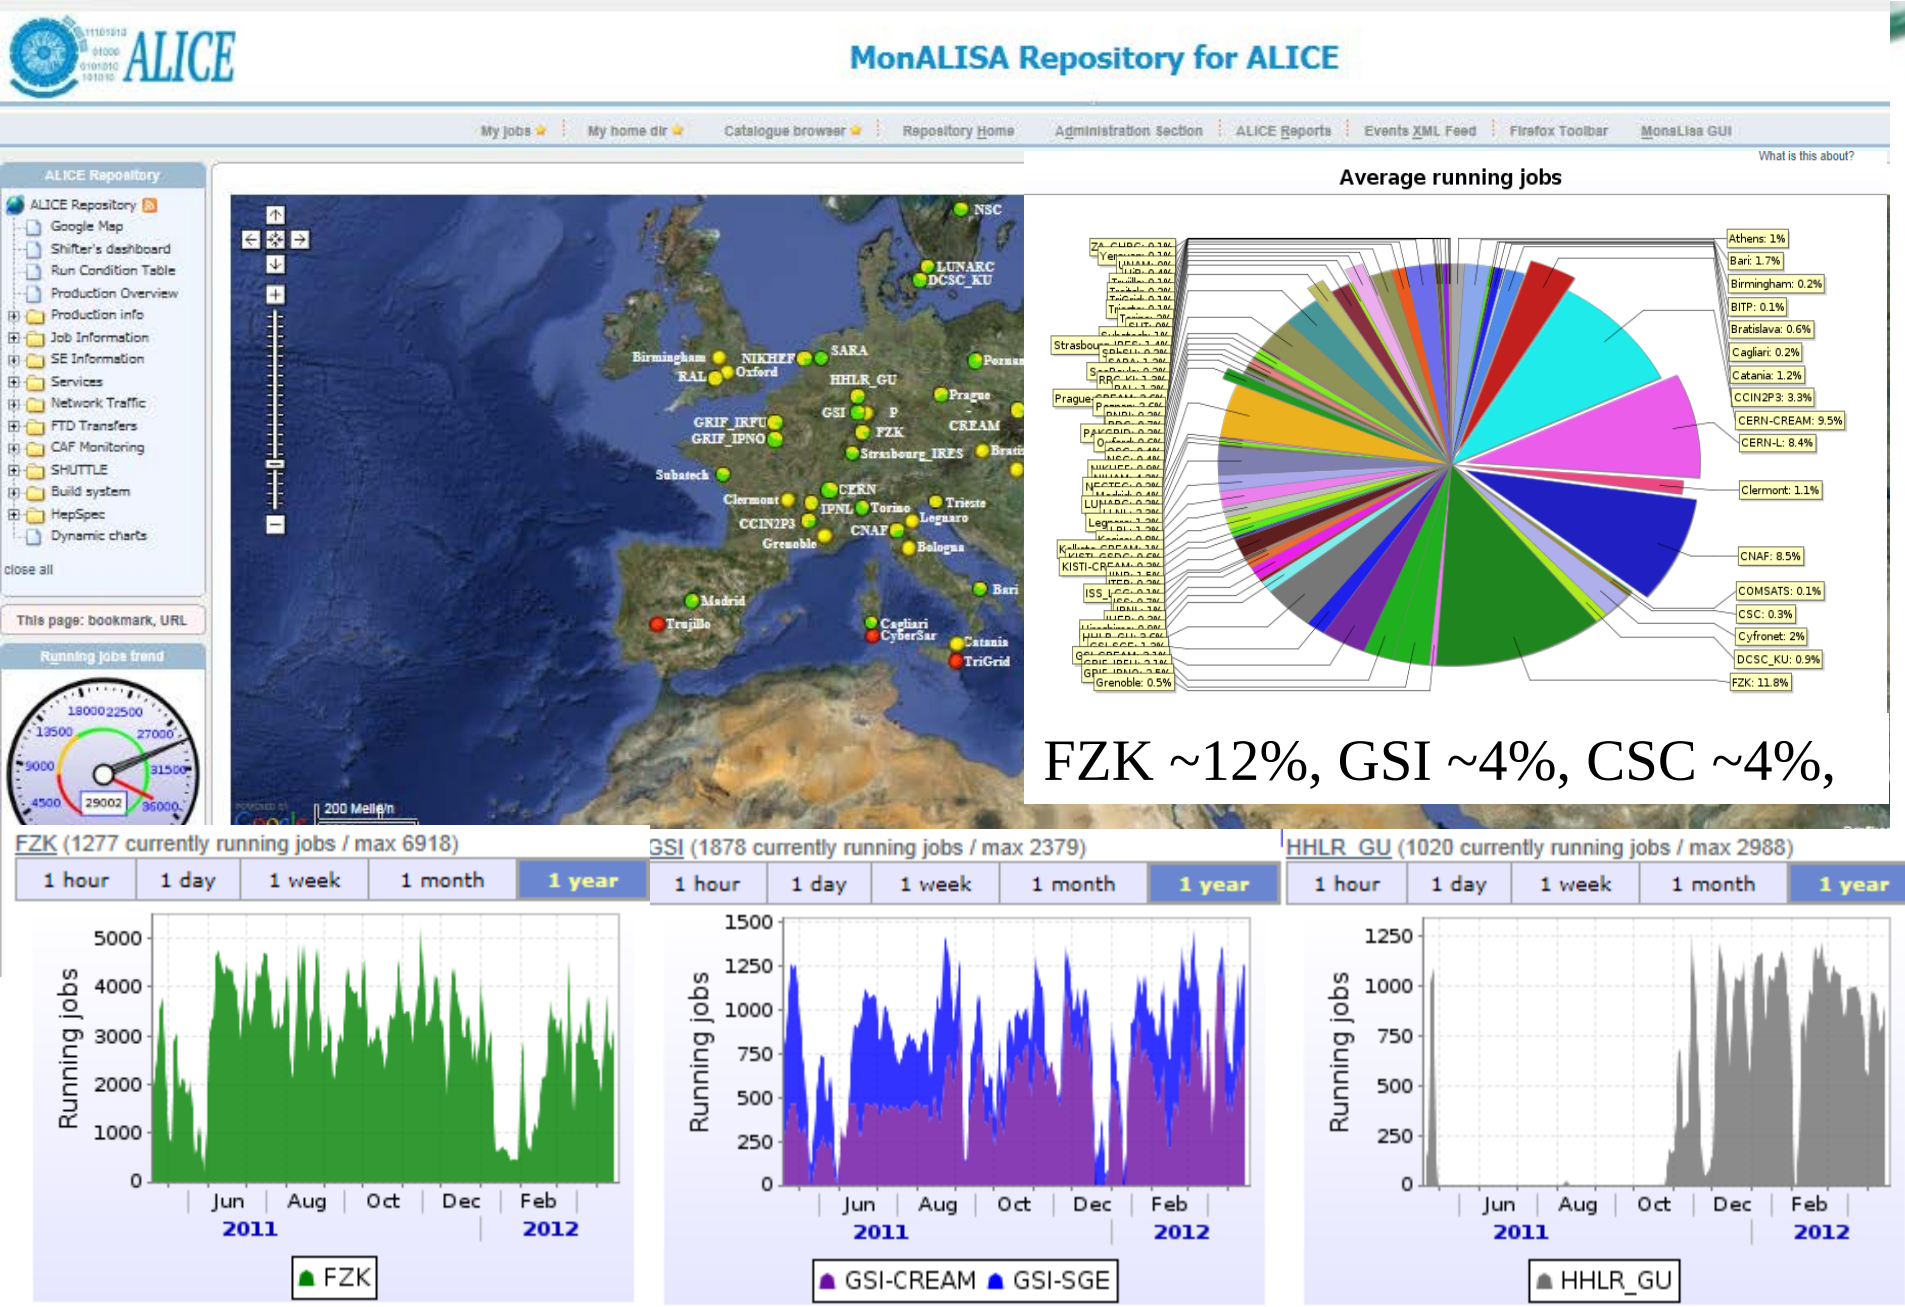
ALICE
MonALISA Repository for ALICE
My jobs
My home dir
Catalogue browser
Repository Home
Administration Section
ALICE Reports
Events XML Feed
Firefox Toolbar
MonsLisa GUI
ALICE Repository
ALICE Repository
Google Map
Shifter's dashboard
Run Condition Table
Production Overview
Production info
Job Information
SE Information
Services
Network Traffic
FTD Transfers
CAF Monitoring
SHUTTLE
Build system
HepSpec
Dynamic charts
close all
This page: bookmark, URL
Running jobs trend
What is this about?
Average running jobs
Athens: 1%
Bari: 1.7%
Birmingham: 0.2%
BITP: 0.1%
Bratislava: 0.6%
Cagliari: 0.2%
Catania: 1.2%
CCIN2P3: 3.3%
CERN-CREAM: 9.5%
CERN-L: 8.4%
Clermont: 1.1%
CNAF: 8.5%
COMSATS: 0.1%
CSC: 0.3%
Cyfronet: 2%
DCSC_KU: 0.9%
FZK: 11.8%
Grenoble: 0.5%
FZK ~12%, GSI ~4%, CSC ~4%,
FZK (1277 currently running jobs / max 6918)
1 hour
1 day
1 week
1 month
1 year
Running jobs
FZK
GSI (1878 currently running jobs / max 2379)
GSI-CREAM
GSI-SGE
HHLR_GU (1020 currently running jobs / max 2988)
HHLR_GU
Jun
Aug
Oct
Dec
Feb
2011
2012
Main problem: we do not have enough disc: < 300TB SE, we need 550 TB

# Grid & Cloud for FAIR

PandaGrid, Cloudification of LOEWE CSC, Ffm Cloud, ...

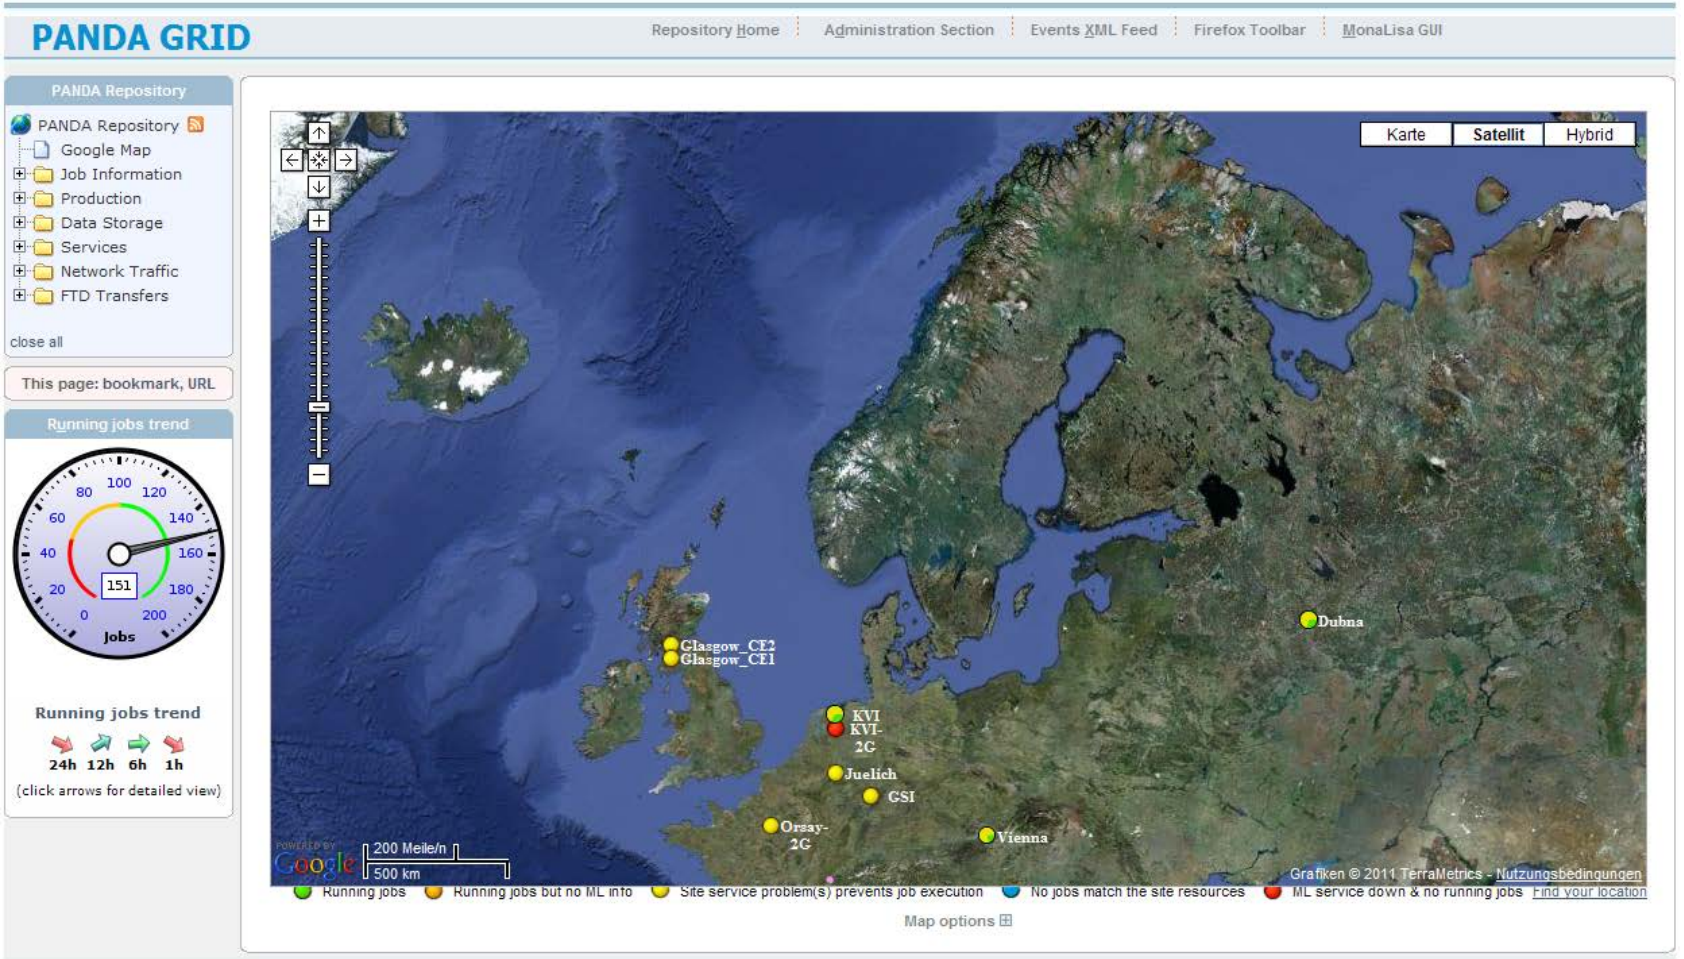

# Grid & Cloud for FAIR

PandaGrid, Cloudification of LOEWE CSC, Ffm Cloud, ...

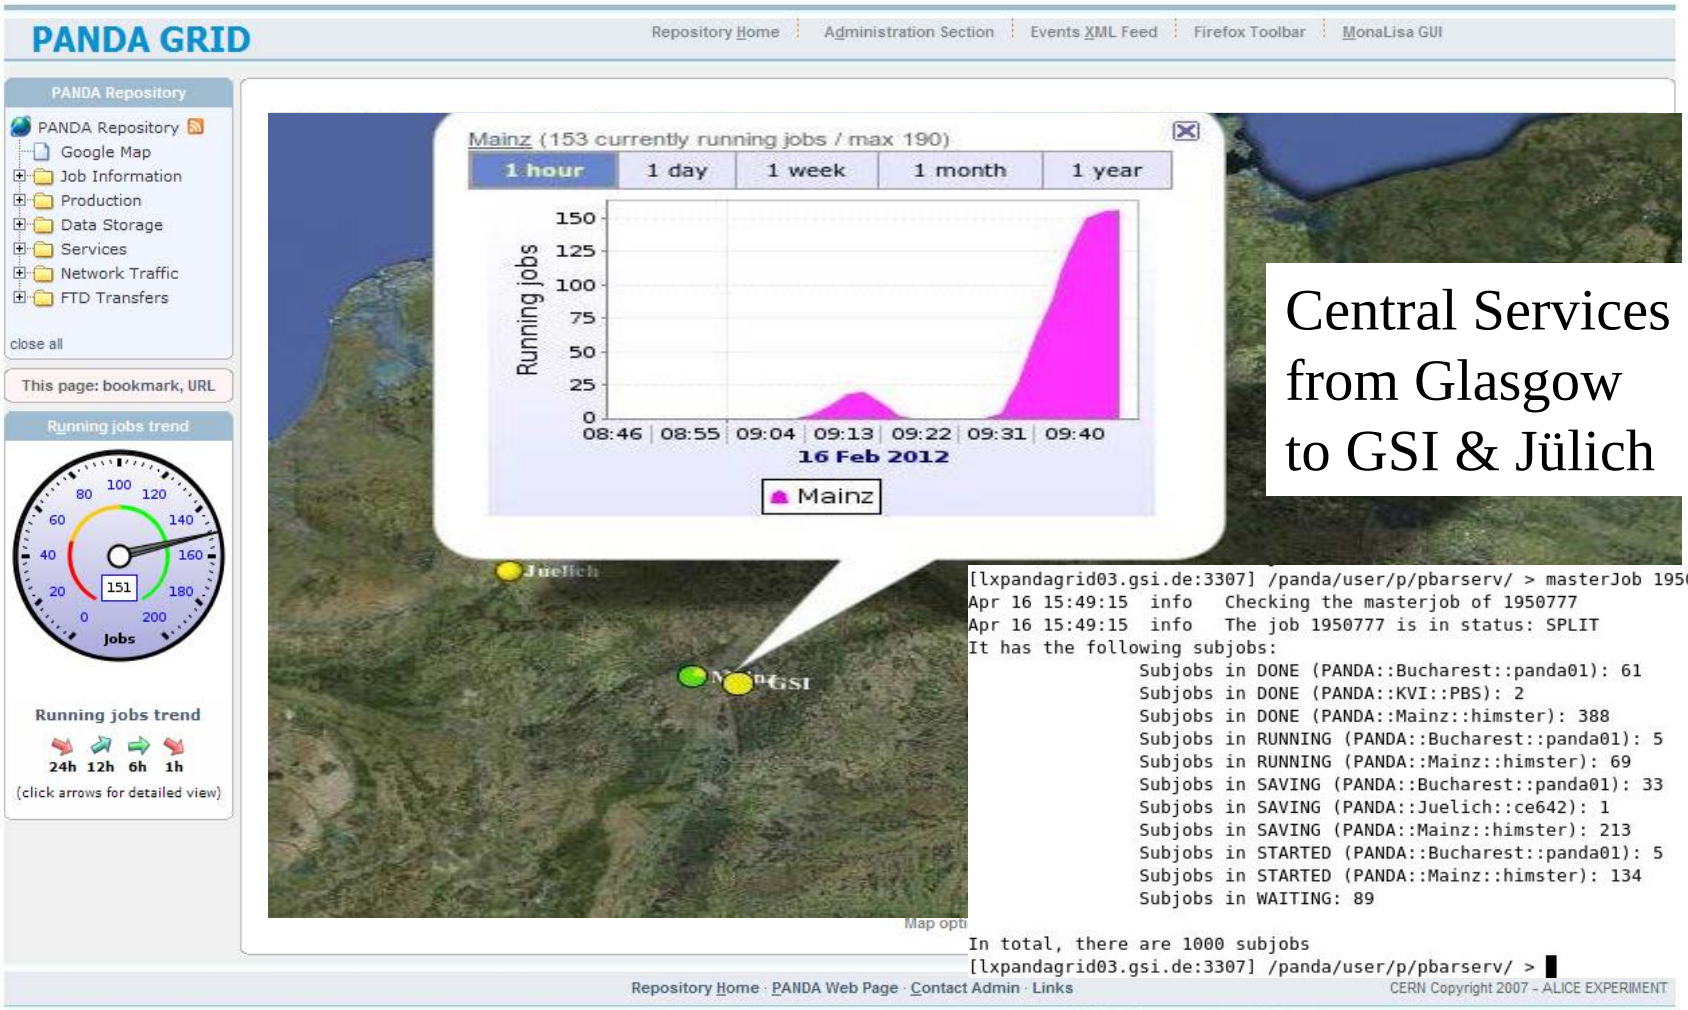


panda GRID AliEn² MonALISA MONitoring Agents using a Large Integrated Services Architecture

http://www.crisp-fp7.eu/


**WP16: Common User Identity System**

**WP19: Distributed Data Infrastructure**

HELMHOLTZ GEMEINSCHAFT

Antrag für das Portfoliothema

***„Large-Scale Data Management and Analysis“ (LSDMA)***

**Beteiligte Helmholtz-Zentren:**

| | | |
|---|---|---|
| KIT | – | Karlsruher Institut für Technologie |
| DESY | – | Deutsches Elektronen-Synchrotron |
| FZJ | – | Forschungszentrum Jülich |
| GSI | – | Helmholtzzentrum für Schwerionenforschung |

**Beteiligte externe Partner:**

| | | |
|---|---|---|
| Uni HD | – | Universität Heidelberg |
| UULM | – | Universität Ulm |
| UHH | – | Universität Hamburg |
| GU | – | Johann Wolfgang Goethe-Universität Frankfurt am Main |
| HTW B | – | Hochschule für Technik und Wirtschaft Berlin |
| DKRZ | – | Deutsches Klimarechenzentrum |
| TUD | – | Technische Universität Dresden |

Koordinierender Sprecher: Prof. Dr. Achim Streit, KIT

# Status of the new data base interface

- Database Interface is in SVN
- Ongoing work on
  - RuntimeDB connection and initialization scheme (GSI)
  - Implementation of Generic Parameter Container (KVI)
  - Preparation of realistic tests ( R3B, Panda ... )
- In Collaboration with KVI and USC (Nustar collaboration) implement and test the online EPICS database interface.

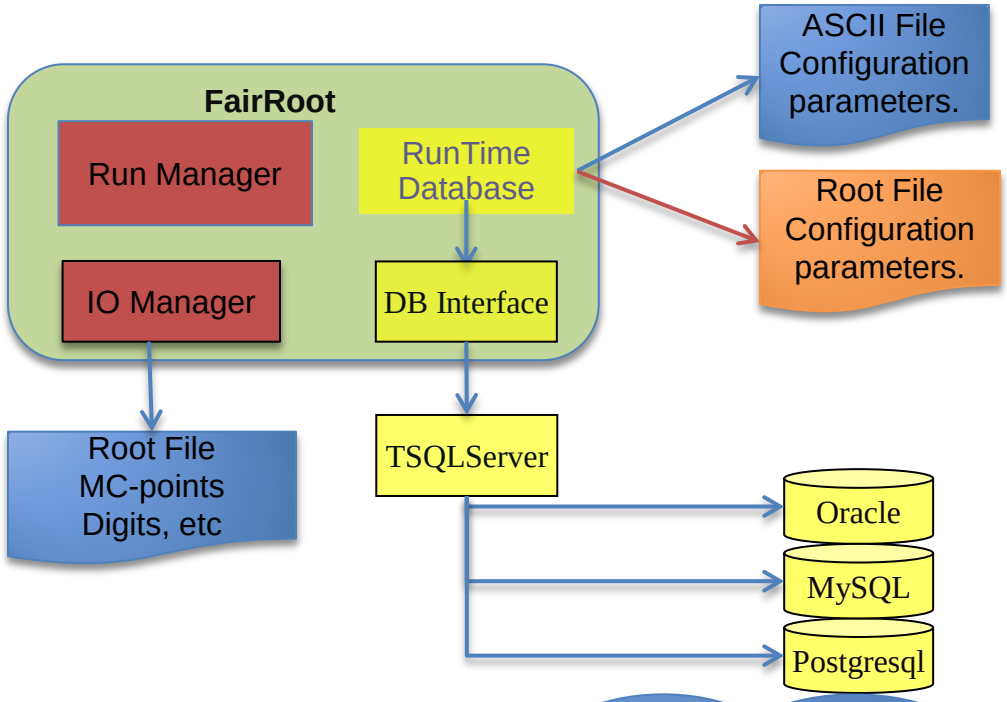


Yago Gonzales
Software engineer from
U Santiago de Compostella
started this month at the
GSI (Two years)

# Proof in FairRoot

In FairRoot change one line in the macro to use it:

FairRunAna *fRun = new FairRunAna();

to

FairRunAna *fRun = new FairRunAna("proof");

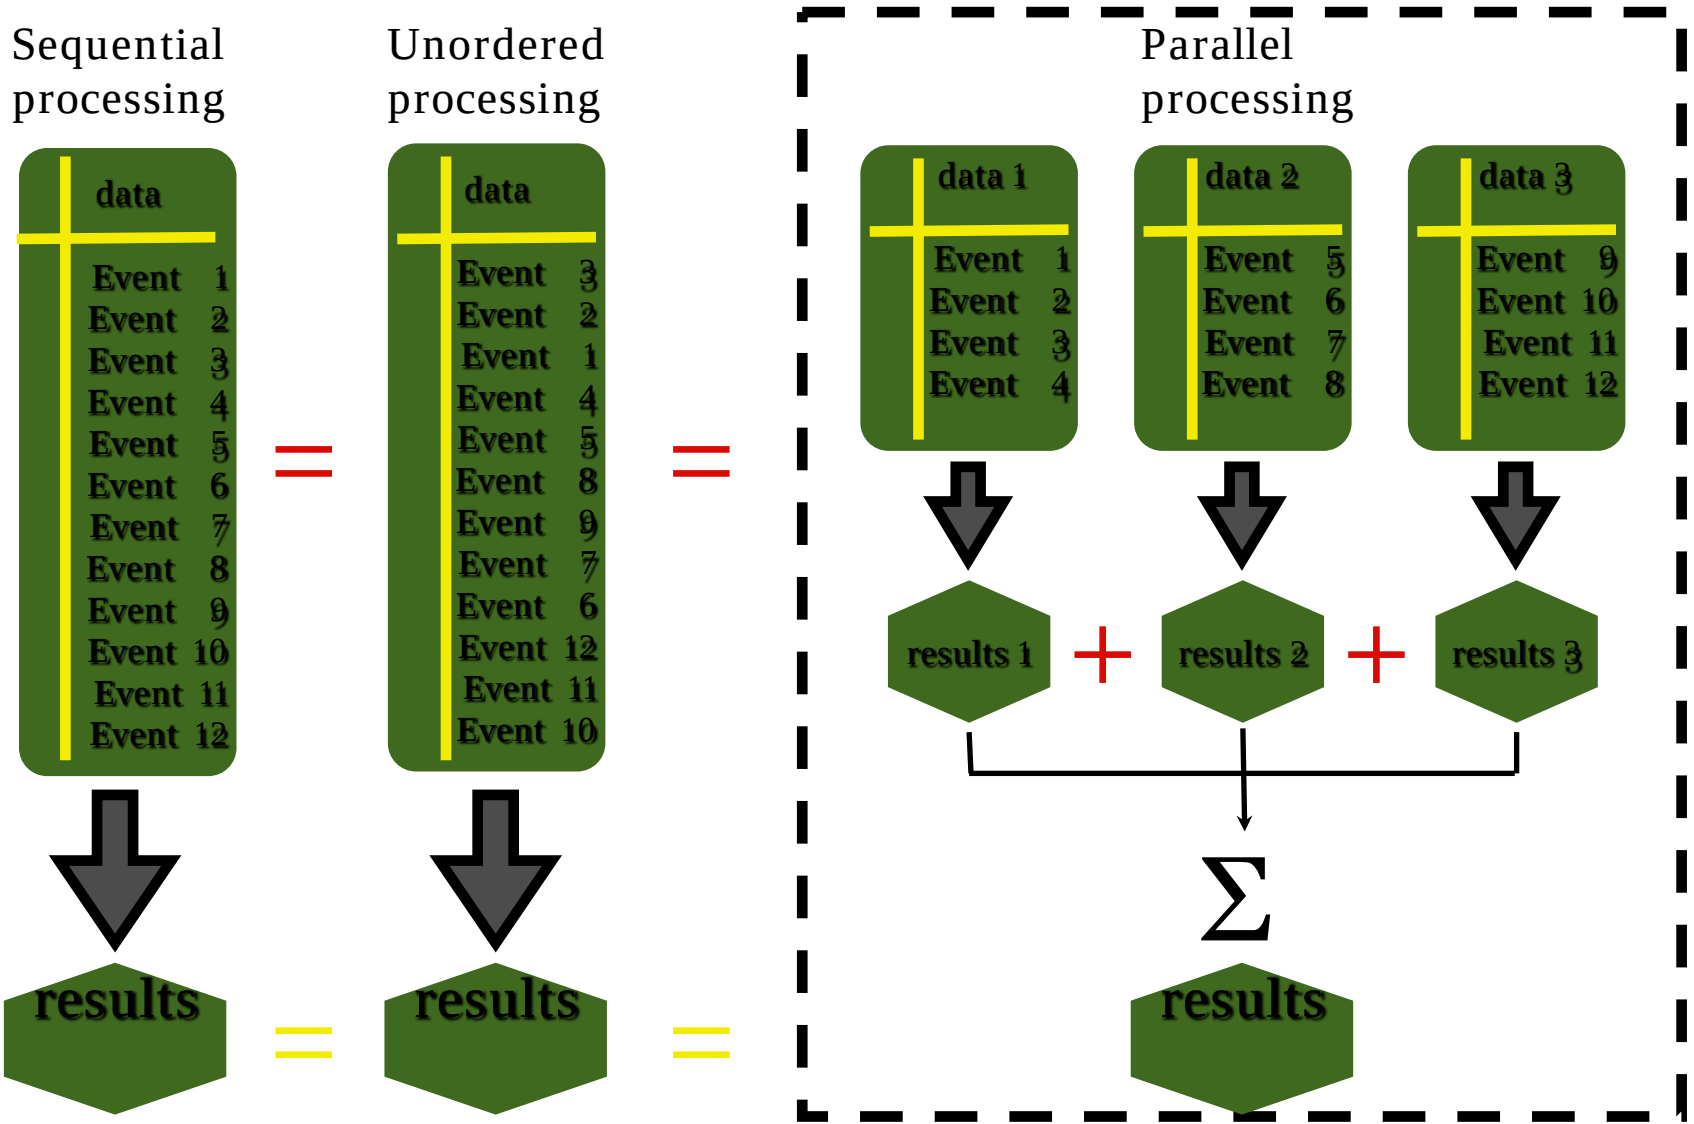

In combination with PoD (PROOF on Demand) : one can use any available resources interactively!

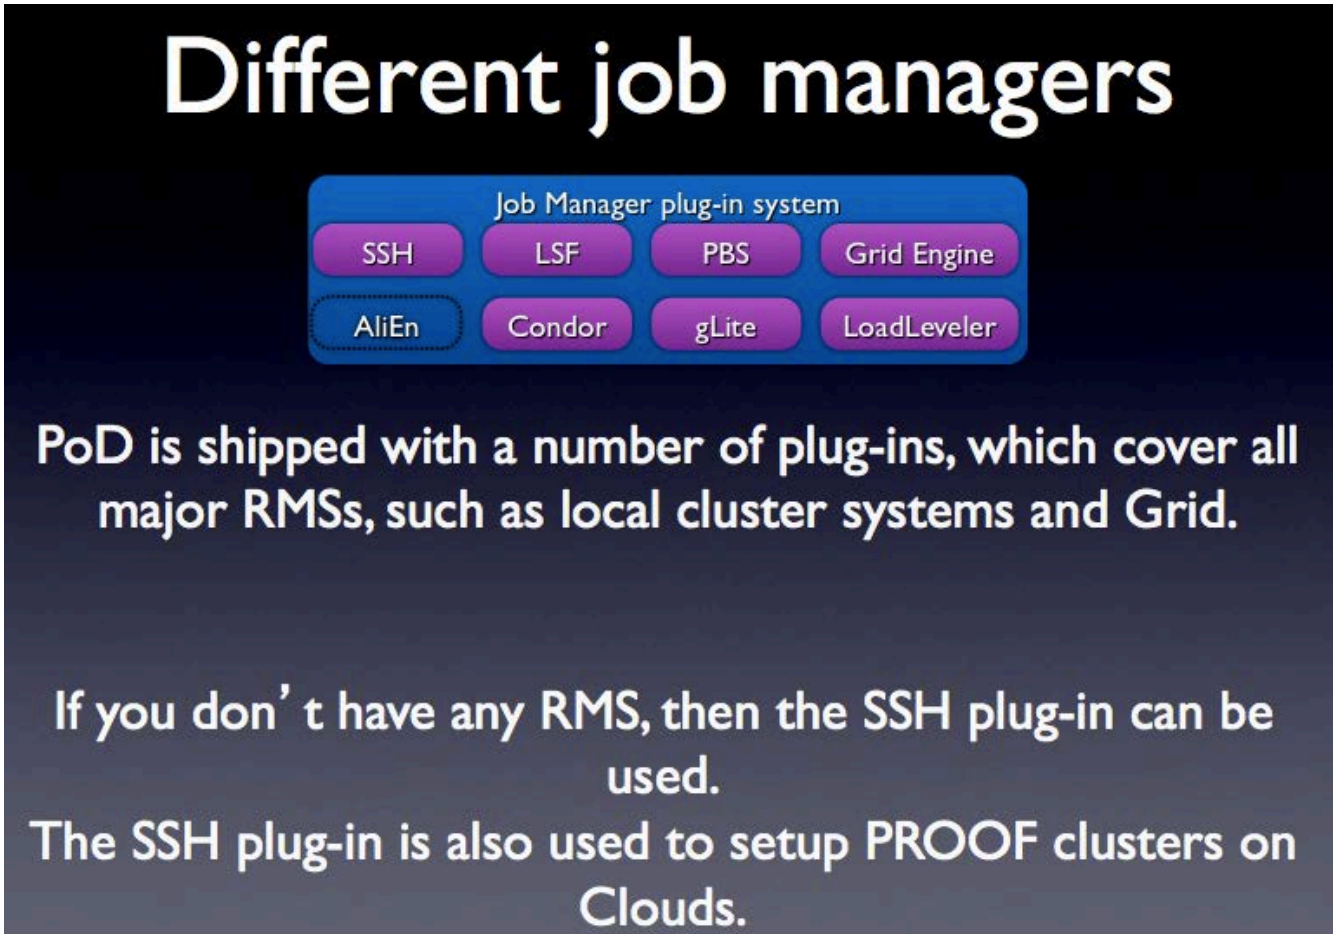


PoD: Anar Manafov

# Event Merging

MC Files:

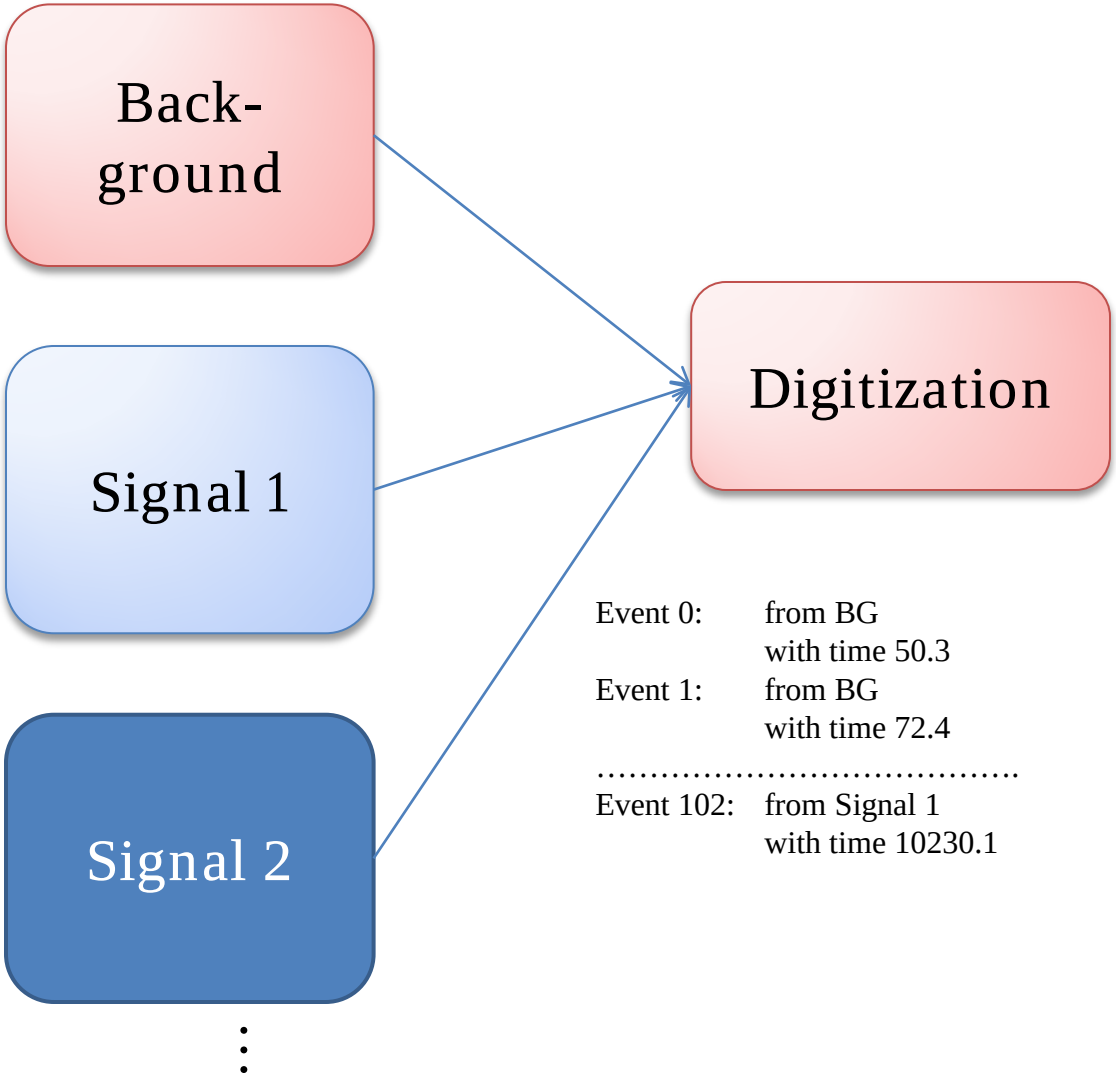


- MC data do not know anything about time structure
- Time structure is calculated when reading the MC for digitization
- Many different signal files can be added to one background file(s)
- Data source (S or B) is stored in EventHeader
- No overlap (pileup) of events at this level

# How do events overlap?

- In Detectors:
  - Sensor elements are still blocked from previous hits
  - Electronic is still busy
  - Hits too close in time cannot be distinguished
  - ...
- Special problem for CBM and PANDA:
  - Continuous beam with Poisson statistics ➔ many events with short time between them
  - No hardware trigger
  - Complex event reconstruction
  - ➔ Necessary to simulate data stream as realistic as possible

# Digi Data Sorted

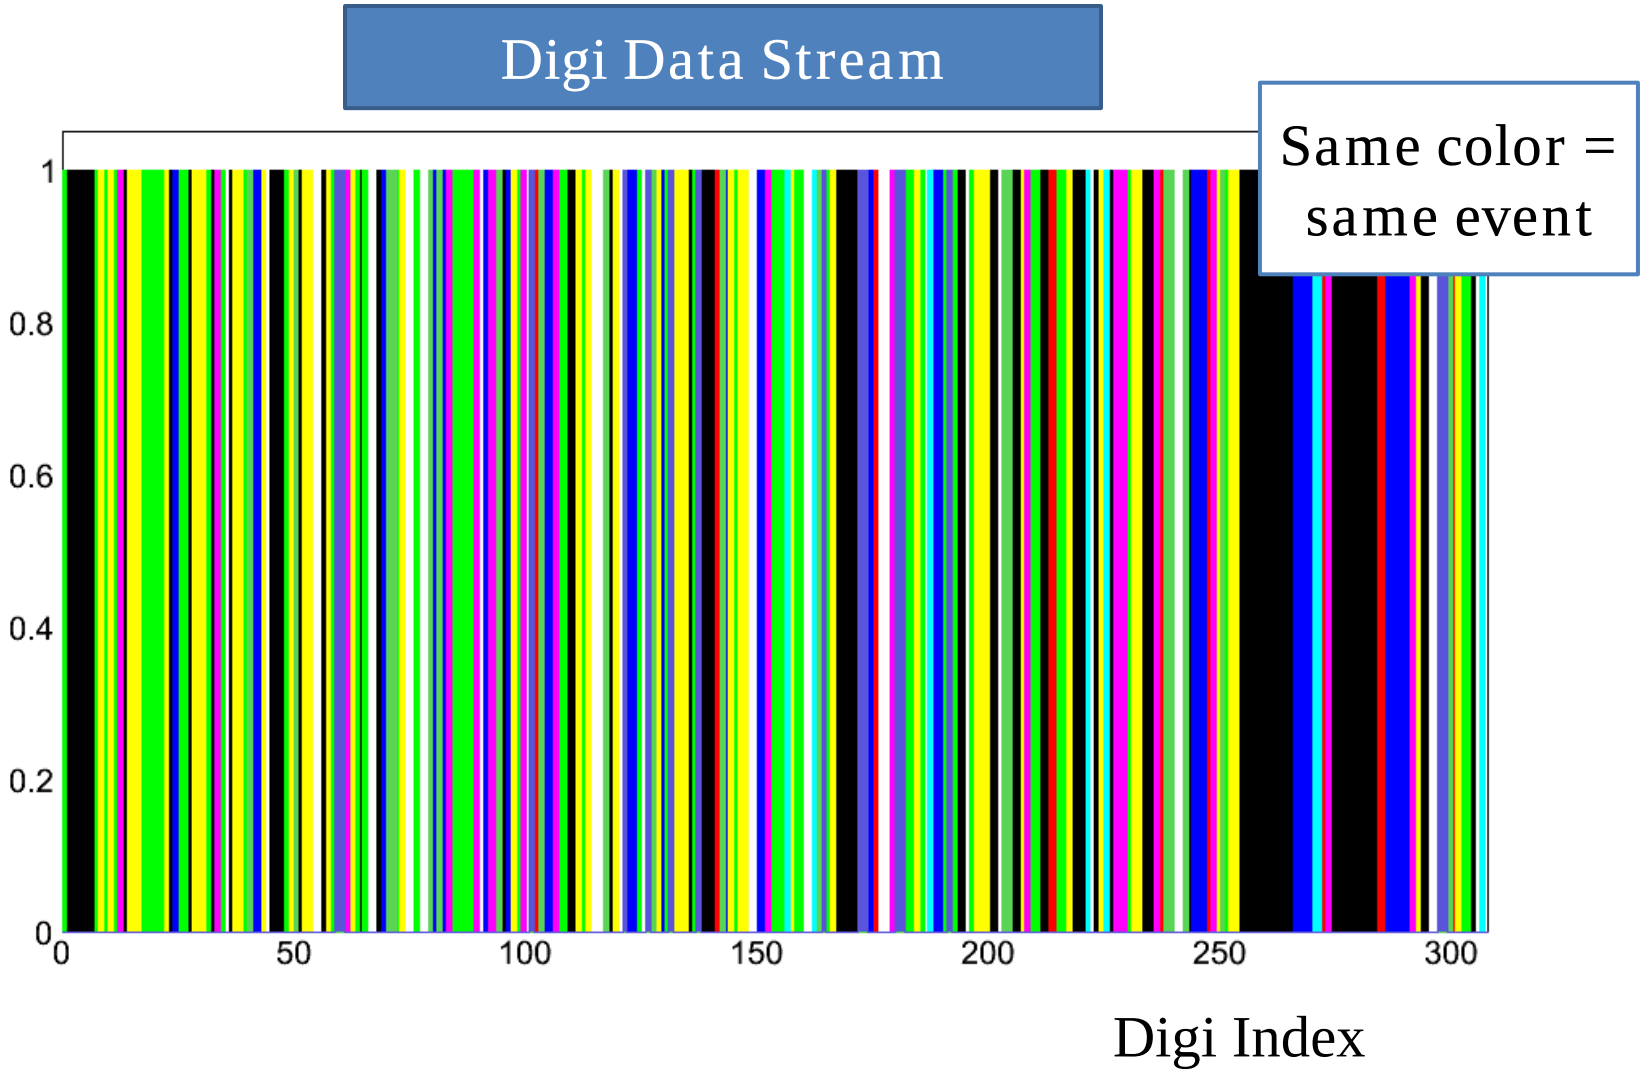

# Digi Data Sorted

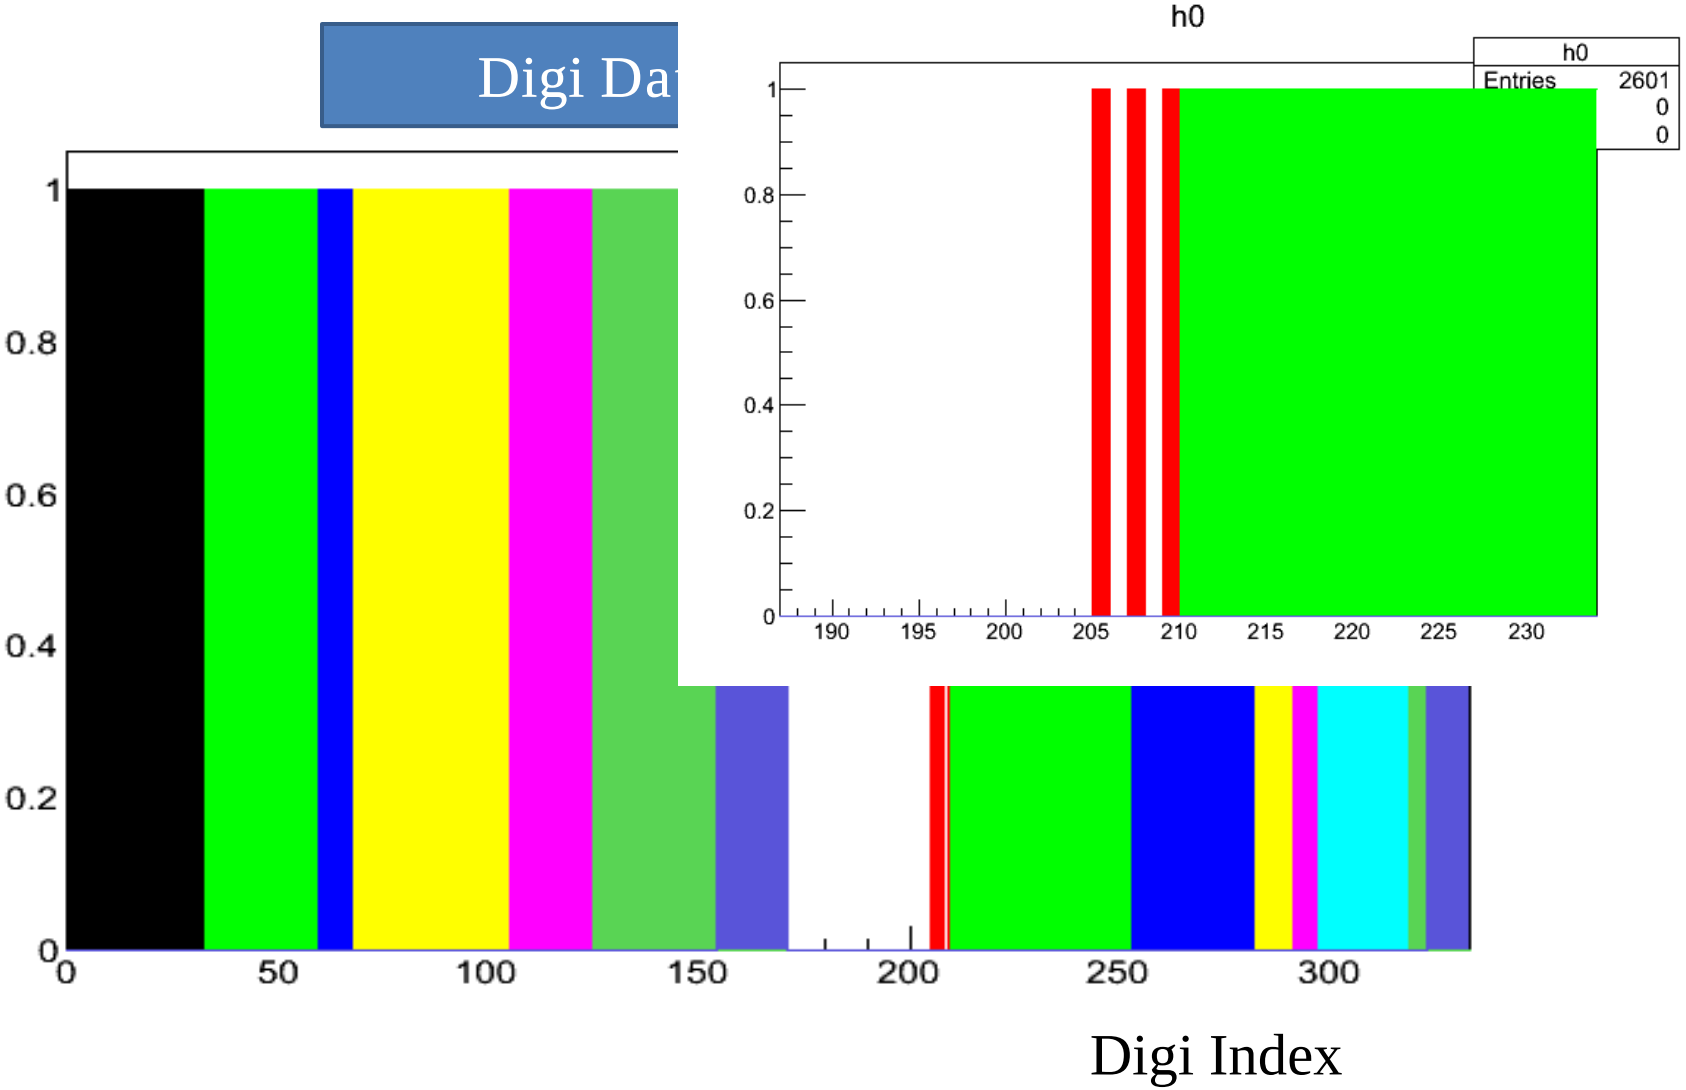

# Reading back data

- FairRootManager has new reading algorithms, which make it possible to use the event wise implemented tasks to run on such data streams
- Different algorithms available to extract data:
  - All data up to a given time
  - All data in a time window
  - All data between time gaps of a certain size
- Other algorithms can be (easily) implemented

# ASY-EOS collaboration (ASYEOSRoot)

- Simulation is done inside FairRoot, ASYEOSRoot branch in SVN
- MBS API in FairRoot is used to read lmd files and analyse experiment data
- Work is ongoing to move all existing data analysis of existing detector parts to FairRoot

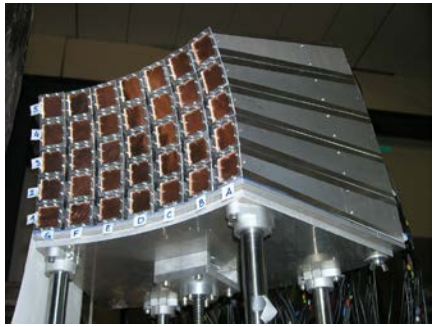

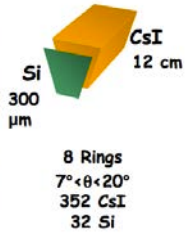


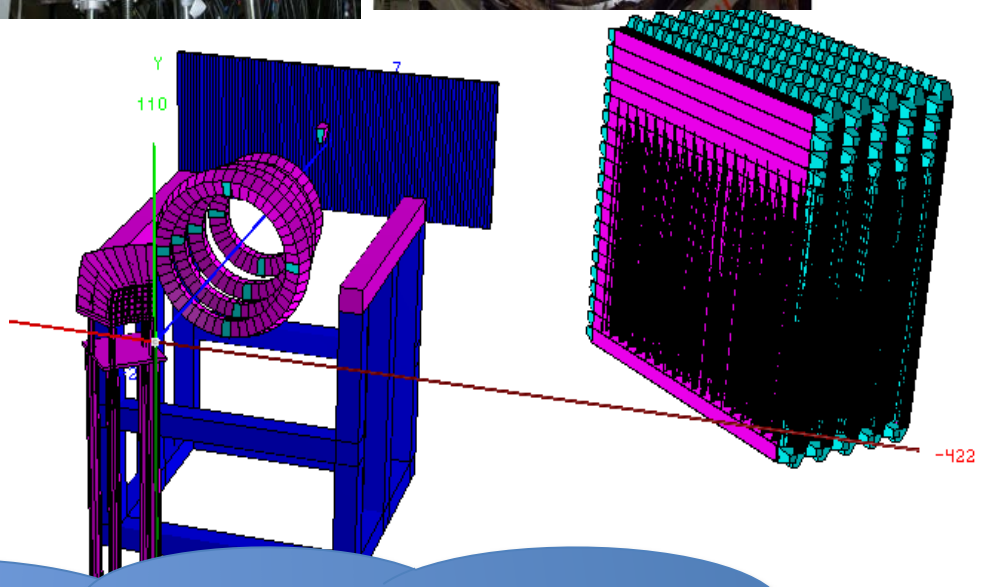


With the ASYEOS and R3B we are developing:

- Client-server based online analysis and monitoring tool
- Online event display (Extend the existing event display to run online)