ATLAS-D Meeting 2013 - Würzburg

# GRID Introduction

M. Böhler

September 24th 2013

Albert-Ludwigs-Universität Freiburg

UNI FREIBURG

# Outline

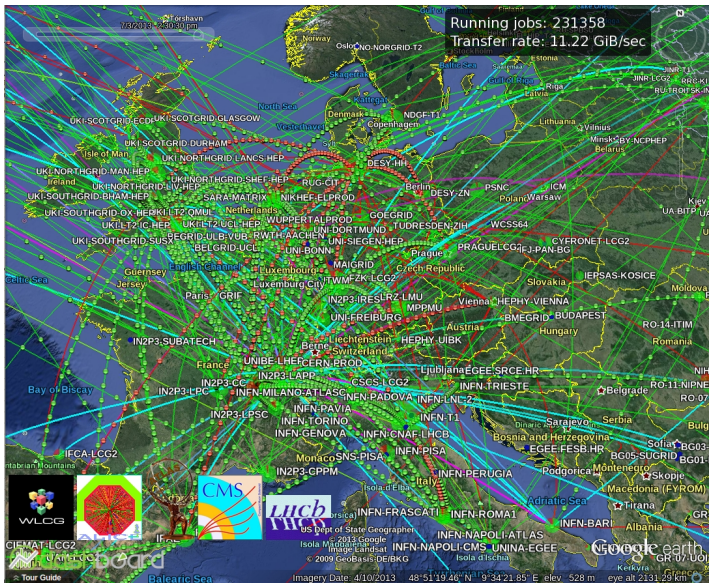


- ATLAS Computing Grid
  - A brief overview of the whole thing
- Distributed Data Management
  - How's data distributed in the Grid?
- Distributed Analysis Interfaces
  - how are analyses scheduled in the Grid?
- GRID and my Analysis?
  - Which tools do I need in order to perform a complete ATLAS analysis?

# The ATLAS Computing GRID

- In fact ATLAS uses three grids:
  - EGI in Europe, Asia and Canada
  - OSG in USA
  - NorduGrid in Nordic countries
- All badged as Worldwide LHC Computing Grid compatible
- However, this is a fact which we mostly try and hide from you

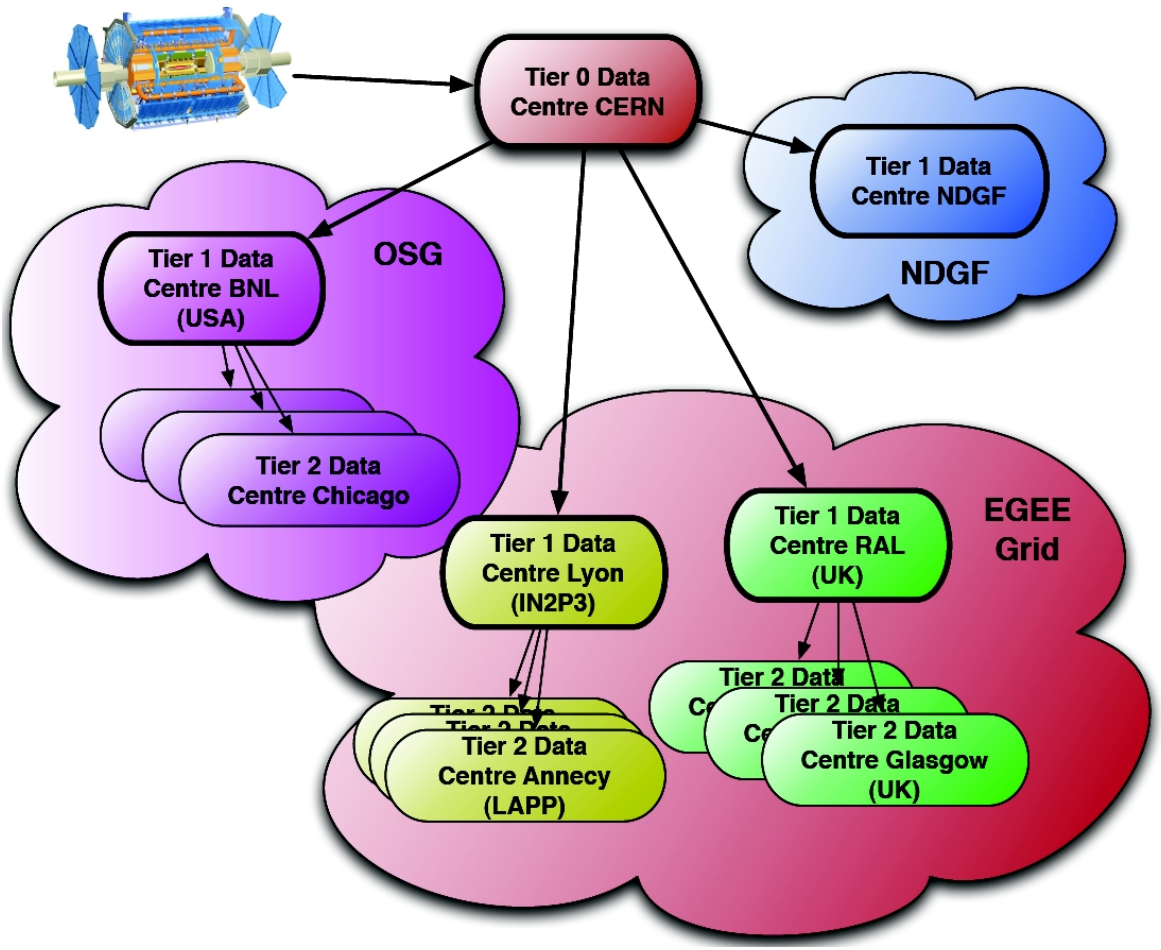

# The ATLAS Computing GRID

- In fact ATLAS uses three grids:
  - EGI in Europe, Asia and Canada
  - OSG in USA
  - NorduGrid in Nordic countries
- All badged as Worldwide LHC Computing Grid compatible
- However, this is a fact which we mostly try and hide from you

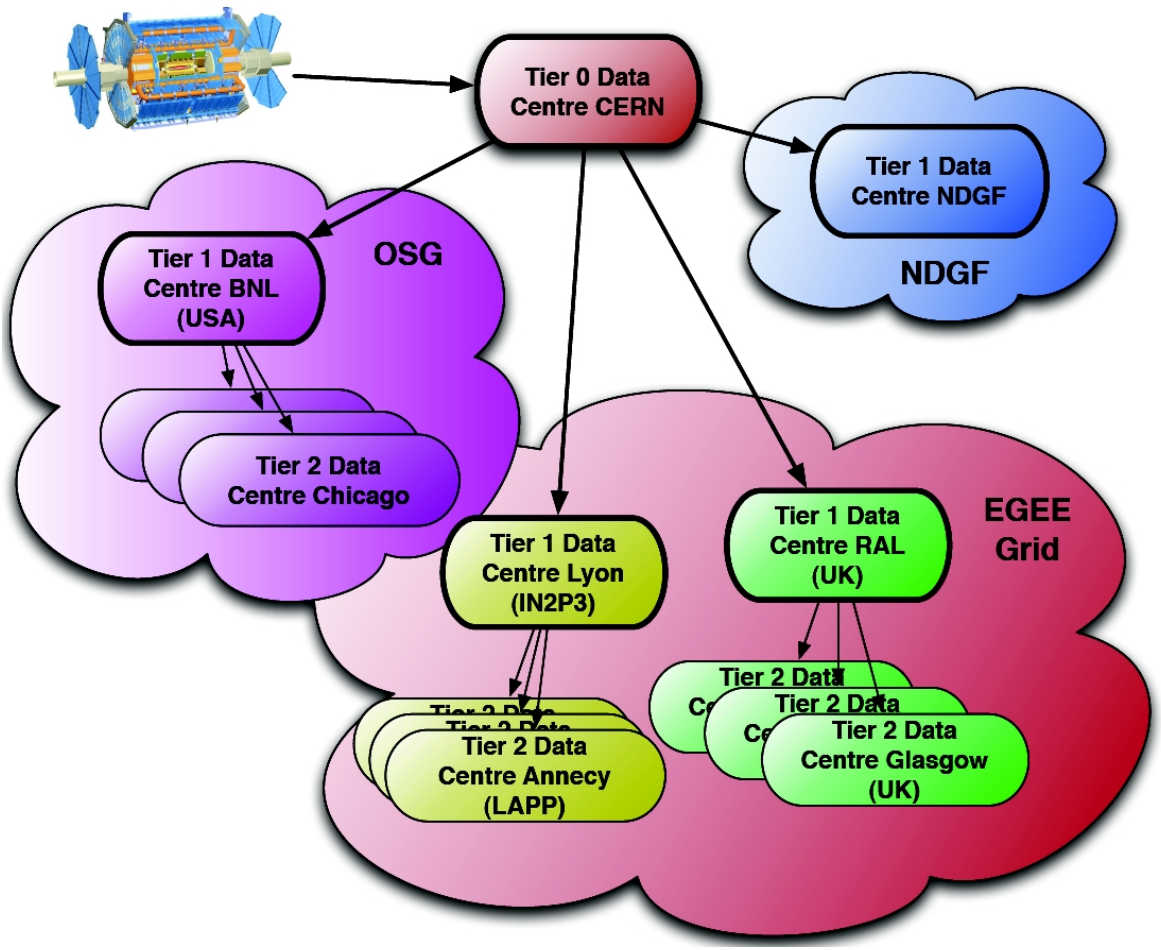


- 1 **Tier-0**: CERN
- 10 **Tier-1**: National Computing Centres (BNL, RAL, IN2P3, ..., FZK)
- $\sim$ 40 **Tier-2**: Regional Computing Centres (CSCS, CYF, DESY-HH,...)
- Composed of multiple individual sites for local User Anlaysis (NAF2.0,...)
- $\sim$100 Analysis queues in PanDA

# ATLAS GRID Data Flow

- prompt reconstruction and several data formats are directly processed at CERN
- reprocessing campaigns (e.g. new SW releases, new calibrations) are processed on the GRID (mainly Tier1s)
- MC production and user analysis (Tier1s and Tier2s)


Johannes Elmsheuser (LMU München)

# Data Formats

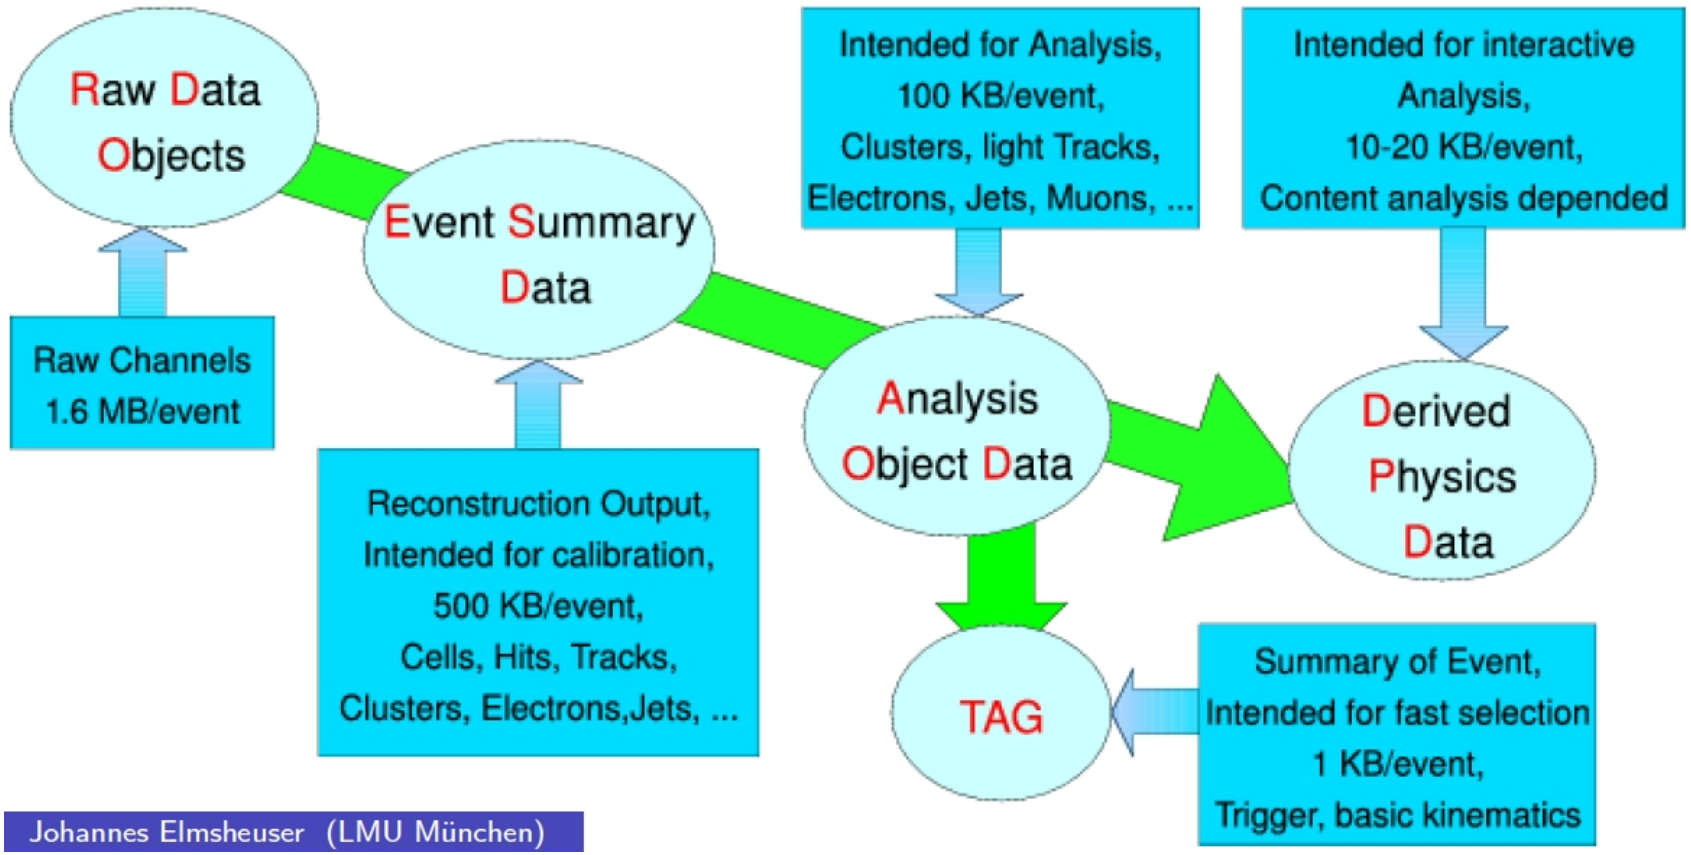


Missing here: dESD and dAOD (Distribution similar to DPD/AOD)

- from most detailed data format RDO to final Derived Physics Data
  - data is combined into higher level objects
  - data is slimmed, skimmed, and thinned

# ATLAS GRID Architecture

- **Production System (prodsys + Panda)** :
  centralized MC simulation and Data reconstruction
  - Useful Tool: ATLAS Metadata Interface (AMI)
    - AMI knows about all datasets, their dependencies and their metadata
      http://atlas-ami.cern.ch:8080/AMI/servlet/

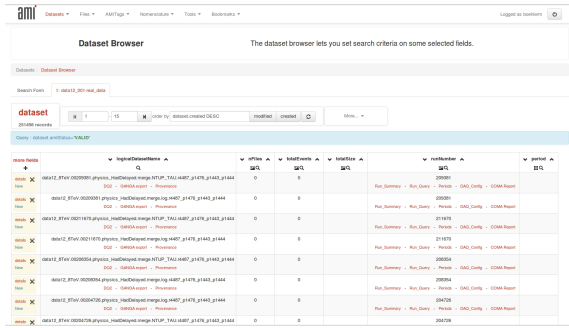


- **Distributed Data Management (DDM/DQ2)** :
  centralized data movement and catalog system

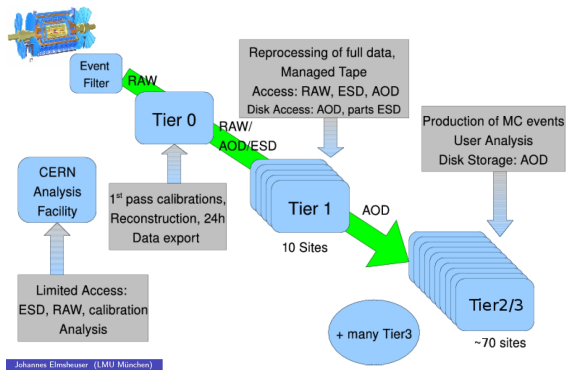


- **Distributed User Analysis** :
  de-centralized individual analysis (ganga + pathena)

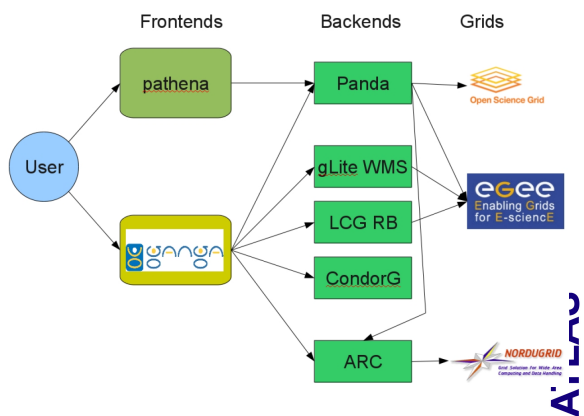

# Distributed Data Management (DDM)

Data management is at the core of all ATLAS GRID activities

- All data organised as datasets
  - Which contain multiple files

```
dq2-ls -f data12_8TeV.00203454.physics_Egamma.merge.NTUP_TAU.r4065_p1278_p1443/

data12_8TeV.00203454.physics_Egamma.merge.NTUP_TAU.r4065_p1278_p1443/
[ ] NTUP_TAU.01224072._000001.root.1 5acc1530-50d2-49ef-8063-5309a886a0b6 ad:575182f9 4826841984
[ ] NTUP_TAU.01224072._000002.root.1 55f20ff3-cbe4-41fa-94d6-9b4a0dffe65a ad:99f3e1cc 3700907305
[ ] NTUP_TAU.01224072._000004.root.1 64017057-887d-4ced-b9a4-aeecc84139ca ad:45bcd11e 3888328769
[ ] NTUP_TAU.01224072._000005.root.1 7bd3f2f7-ad33-4fe3-9dfb-44464bbe75e4 ad:9b1f4745 3809997412
....
[ ] NTUP_TAU.01224072._000015.root.1 f5c35582-6a09-4d26-a8db-9be15473e800 ad:f60637f9 4369755635
total files: 15
local files: 0
total size: 64453708501
date: 2013-03-25 20:09:46
```

- Datasets are the units of replication on the GRID
  - So the DDM central catalog records the location of complete datasets on the GRID, as well as the content of each dataset
  - Datasets can be grouped in containers
    https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/PhysicsContainers
  - Datasets are also the basic input to an analysis
- DDM also moves data between sites on the GRID

# Datasets

- Dataset can have different states
  - Open: new files can be attached to the dataset
  - Closed: no new files can be attached to the current dataset version.
    But a new version can be created and new files added to it
  - Frozen: no new files can be added to the datasets
    not possible to create a new version
- The datasets/containers can be subscribed (copied) from one site to another site.
  - Official datasets (e.g. AOD, ESD, DESD) are placed centrally on different sites and have many replicas on some official area
  - Datasets produced by groups (e.g. D2PD) are managed by group managers and placed on some group area
  - User can request some of the official datasets, group or user datasets to be moved to some user areas
- The different areas are identified by Space Tokens

# Space tokens

- ATLASDATATAPE/ATLASMCTAPE :
  - The TAPE Space Tokens are only available in Tier1s
  - can be read only by people with production role
- ATLASDATADISK/ATLASPRODDISK :
  - official space token only writable by people with special role (Production Role)
  - space tokens are readable by every ATLAS users

DE-Tier1

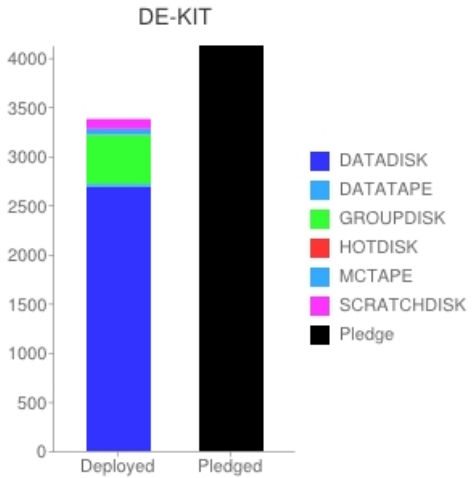


Space tokens dedicated to users:

- ATLASSCRATCHDISK : (DESY-HH_SCRATCHDISK)
  - *transient space*, with a lifetime of less than 30 days
  - default space where Distributed Analysis tools write
  - data needs to be stored permanently somewhere else
    1. download them using DQ2 tools
    2. move them on a permanent storage e.g. using DATRI
- ATLASLOCALGROUPDISK : (DESY-HH_LOCALGROUPDISK)
  - this is permanent storage
  - available on all German sites included in DQ2
  - Only people with group /atlas/de in VOMS can request replication to this space
  - no space limitation per user - BUT if you are using too much space O(5 TB) $\rightarrow$ you might be contacted...

typical Tier2:

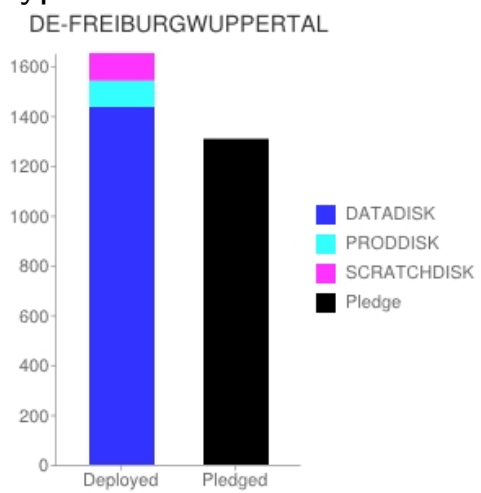

# Distributed Analysis

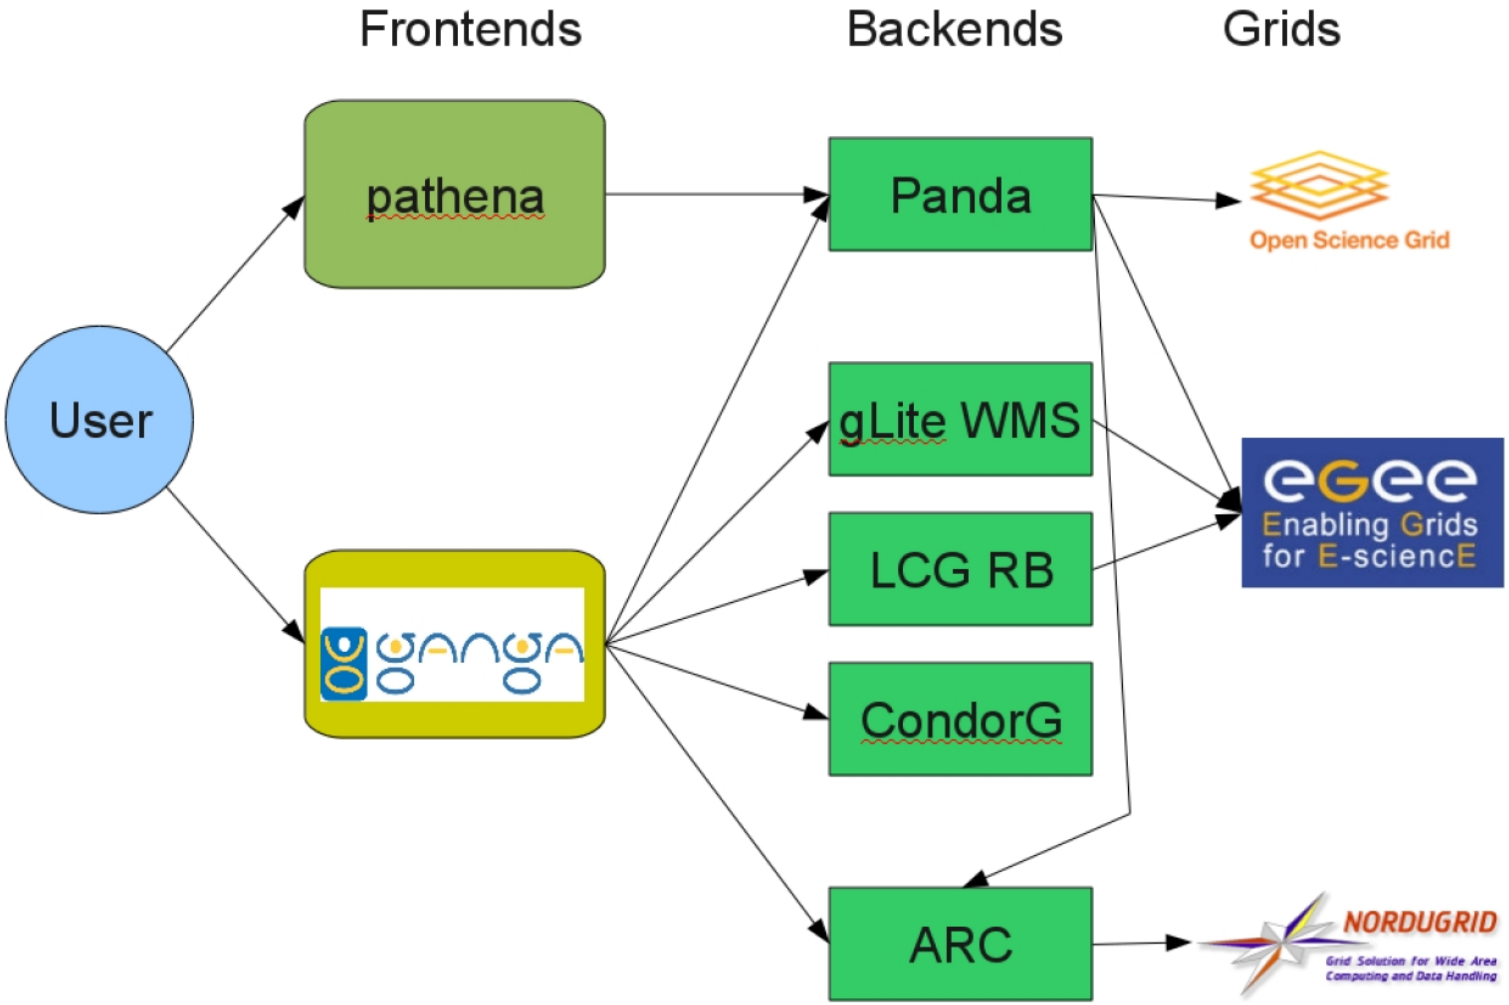


- Data is centrally being distributed by DQ2
- Jobs go to data

## PanDA Client Tools

- PanDA is the main ATLAS Production and Distributed Analysis system
- PanDA Client Tools consists of several tools for job execution on the grid and bookkeeping
  - pathena - for submitting Athena jobs to PanDA
  - prun - for general jobs (e.g. ROOT and Python scripts) to PanDA
  - psequencer allows for a sequence of different tasks to be submitted (e.g. an analysis job followed by the transfer of the output back to the local machine)
  - pbook is a bookkeeping tool for all PanDA analysis jobs

## PanDA Client Tools

- PanDA is the main ATLAS Production and Distributed Analysis system
- PanDA Client Tools consists of several tools for job execution on the grid and bookkeeping
  - pathena - for submitting Athena jobs to PanDA
  - prun - for general jobs (e.g. ROOT and Python scripts) to PanDA
  - psequencer allows for a sequence of different tasks to be submitted (e.g. an analysis job followed by the transfer of the output back to the local machine)
  - pbook is a bookkeeping tool for all PanDA analysis jobs

- ATLAS analysis has (at least) two stages:
  1. Run Athena on AOD/ESD to produce DPD (pathena)
  2. Run ROOT, Python, or shell scripts to produce final plots (prun)

# PanDA Client Tools

- PanDA is the main ATLAS Production and Distributed Analysis system
- PanDA Client Tools consists of several tools for job execution on the grid and bookkeeping
  - pathena - for submitting Athena jobs to PanDA
  - prun - for general jobs (e.g. ROOT and Python scripts) to PanDA
  - psequencer allows for a sequence of different tasks to be submitted (e.g. an analysis job followed by the transfer of the output back to the local machine)
  - pbook is a bookkeeping tool for all PanDA analysis jobs

- ATLAS analysis has (at least) two stages:
  1. Run Athena on AOD/ESD to produce DPD (pathena)
  2. Run ROOT, Python, or shell scripts to produce final plots (prun)

1. pAthena: Client tool for PanDA to submit user-defined jobs from the command line
   - Works on the Athena runtime environment
   - A consistent user-interface to Athena

When you run Athena locally with *athena jobOptions.py* all you need to do to submit a job to the grid is

```
pathena jobOptions.py [--inDS inputDataset] --outDS outputDataset
```

for more instructions see:

```
pathena -h
```

# Monitoring PanDA jobs (1)

2. prun: submit a simple (e.g. plotting) job:

```
prun --outDS user.boehlerm.pruntest_Hello --exec HelloWorld.py

INFO : gathering files under /pathena
INFO : upload source files
INFO : submit 2 subjobs to ANALY_INFN-PAVIA
===================
JobsetID  : 1124
JobID  : 1125
Status : 0
  > build
    PandaID=1903903212
  > run
    PandaID=1903903213
```

...you can directly follow your job with the PandaID, here: PandaID=1903903212
Go to: http://panda.cern.ch and enter your PanDA id in job field in the left column

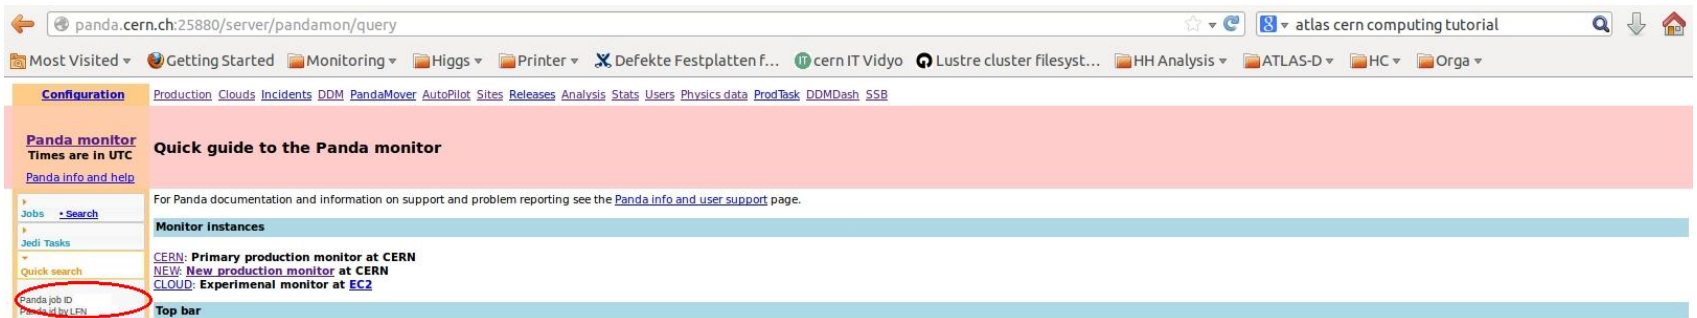

# Monitoring PanDA jobs (2)

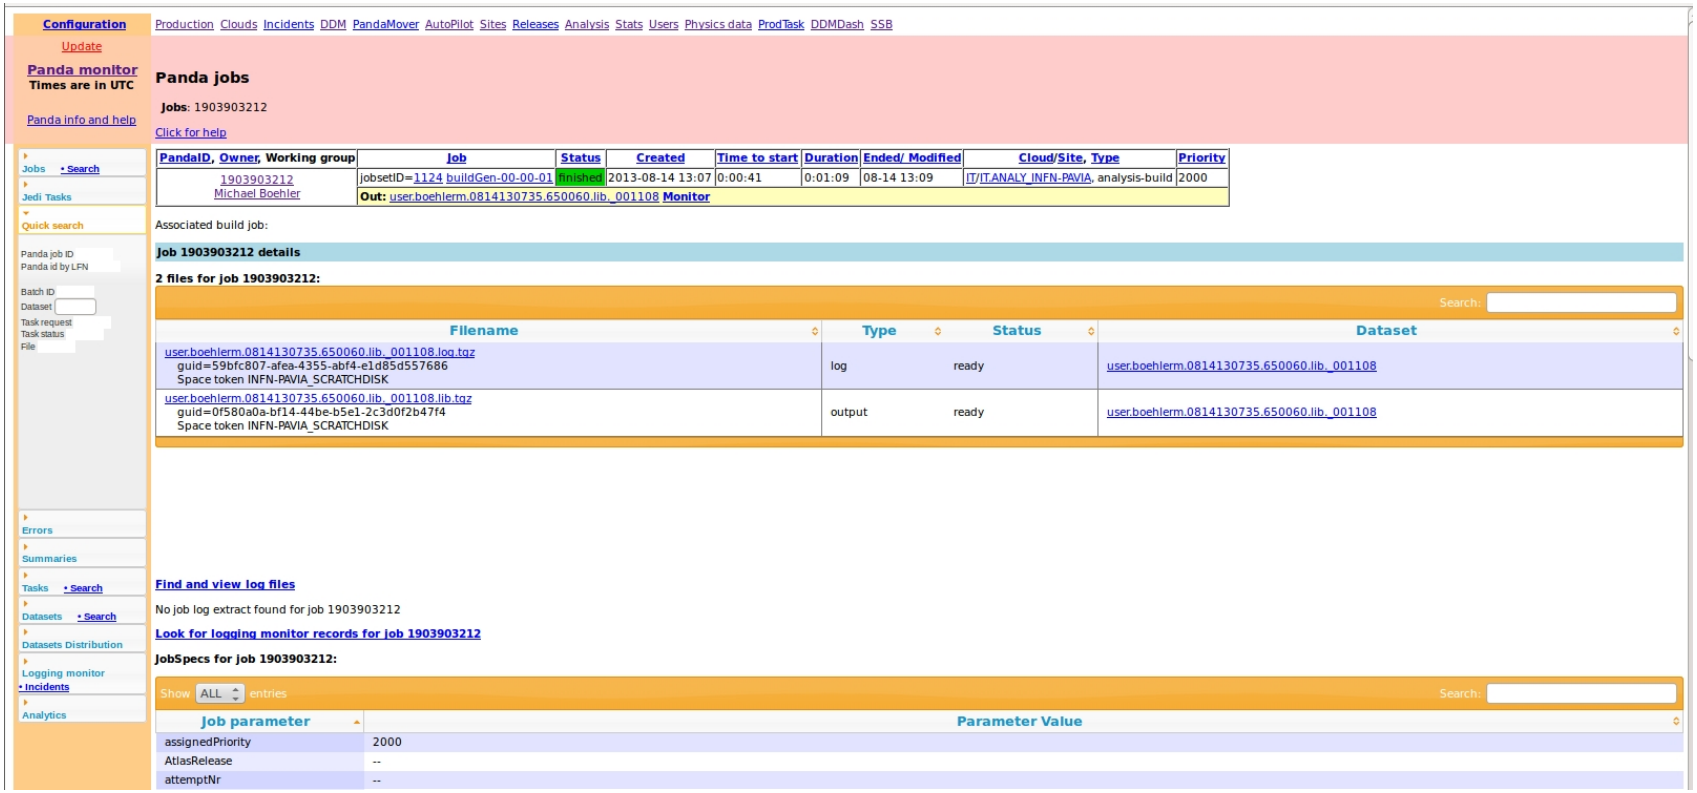


- job finished correctly (green field)
- Most common grid failures: stage-in or stage-out
- For less obvious failures, check log files → **Find and view log files**

# Ganga

- Ganga combines components to provide a front-end client for interacting with Grid infrastructures
- Ganga allows simple switching between testing on a local batch system and large-scale data processing on Grid distributed resources
- Jobs look the same whether they run locally or on the Grid
  $\rightarrow$ Configure once, run anywhere

GANGA offers three ways of user interaction:

- Shell command line (example on this page)
- Interactive IPython shell
- Graphical User Interface

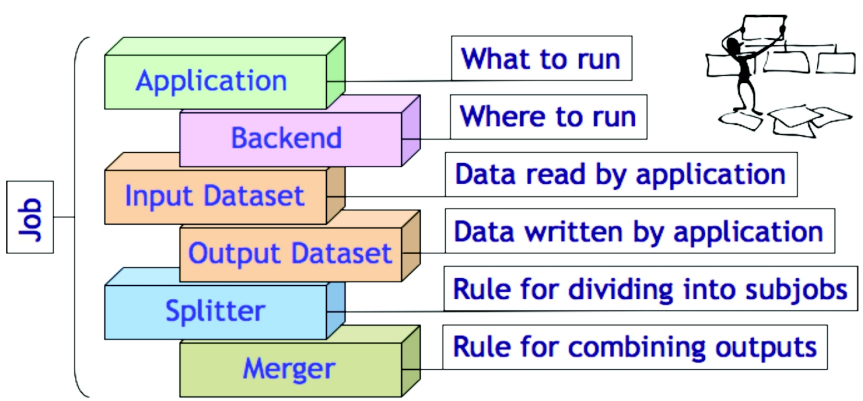


**ganga athena**
AnalysisSkeleton_topOptions.py
––panda
––inDS fdr08_run2.0052283.physics_Muon
.merge.AOD.o3_f8_m10
––outputdata AnalysisSkeleton.aan.root
––split 3

# Monitoring Ganga Jobs

- Interactive IPython shell:
  - just setup and start ganga:

```
setupATLAS
localSetupGanga
ganga
```


More Ganga features:

- resubmit functionality
- submit more jobs when others are completed

- Graphical Web Interface (not Ganga specific, but you can use it to monitor your jobs): https://dashb-atlas-task.cern.ch/templates/task-analysis/

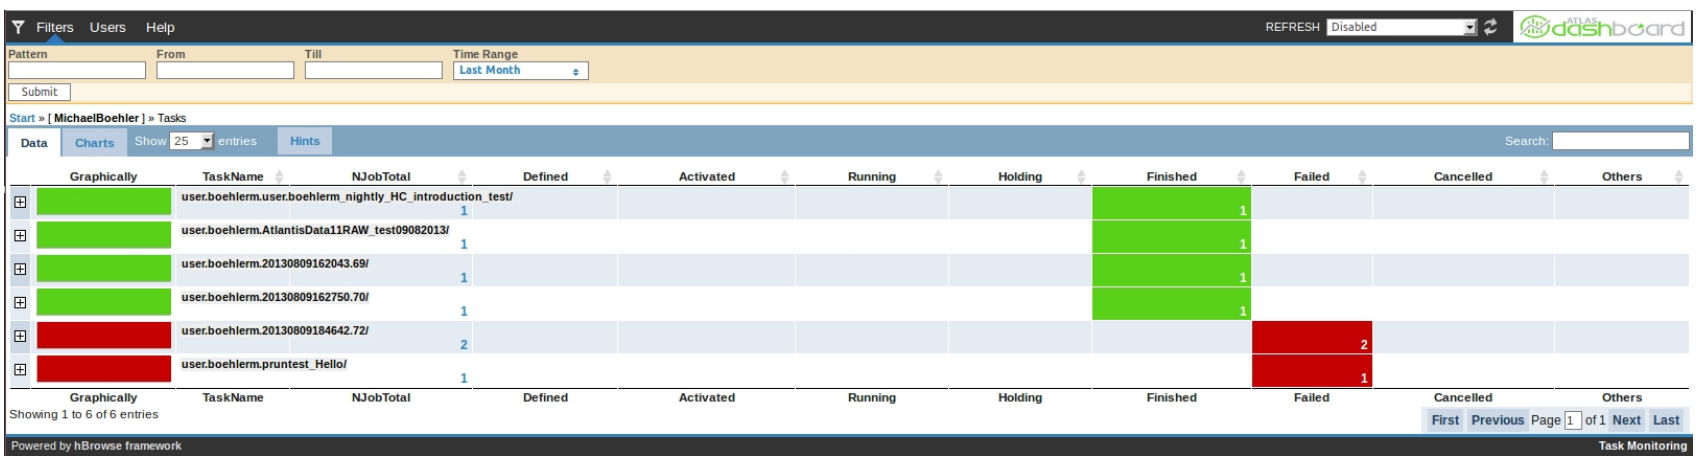

## Which (Grid) tools do I need for my analysis?

1. Download Signal sample of my analysis from the GRID $\rightarrow$ dq2-get (next hands-on session)

- This is the plan of the day!

## Which (Grid) tools do I need for my analysis?

1. Download Signal sample of my analysis from the GRID → dq2-get (next hands-on session)
2. Implement an analysis and run on small test samples locally (this afternoon with MANA)

- This is the plan of the day!

## Which (Grid) tools do I need for my analysis?

1. Download Signal sample of my analysis from the GRID → dq2-get (next hands-on session)
2. Implement an analysis and run on small test samples locally (this afternoon with MANA)
3. Large scale analysis of many backgrounds

- This is the plan of the day!

## Which (Grid) tools do I need for my analysis?

1. Download Signal sample of my analysis from the GRID → dq2-get (next hands-on session)
2. Implement an analysis and run on small test samples locally (this afternoon with MANA)
3. Large scale analysis of many backgrounds
   - either on a local batch system, or on the GRID → this you can do with Panda or with Ganga (in the hands-on session with MANA)

- This is the plan of the day!

## Which (Grid) tools do I need for my analysis?

1. Download Signal sample of my analysis from the GRID → dq2-get (next hands-on session)
2. Implement an analysis and run on small test samples locally (this afternoon with MANA)
3. Large scale analysis of many backgrounds
   - either on a local batch system, or on the GRID → this you can do with Panda or with Ganga (in the hands-on session with MANA)
4. You found a couple of interesting/problematic events

- This is the plan of the day!

## Which (Grid) tools do I need for my analysis?

1. Download Signal sample of my analysis from the GRID → dq2-get (next hands-on session)
2. Implement an analysis and run on small test samples locally (this afternoon with MANA)
3. Large scale analysis of many backgrounds
   - either on a local batch system, or on the GRID → this you can do with Panda or with Ganga (in the hands-on session with MANA)
4. You found a couple of interesting/problematic events
   - check the event by eye make your own event display (last session of the day)

- This is the plan of the day!

## More Help:

- Distributed Data Management
  - DDM Twiki:
    https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/DistributedDataManagement
- Panda
  - Where to get help when you need it!
    hn-atlas-dist-analysis-help@cern.ch (Distributed Analysis Help)
  - CERN Tutorial:
    https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/SoftwareTutorialUsingTheGrid
  - Distributed Analysis with Panda (with FAQ)
    https://twiki.cern.ch/twiki/bin/view/Atlas/DAonPanda
  - Find your jobs in the PanDA monitor
    http://panda.cern.ch
  - Panda wiki page
    https://twiki.cern.ch/twiki/bin/view/AtlasComputing/Panda
- Ganga:
  - All analysis-tools related problems → The Distributed Analysis List
    hn-atlas-dist-analysis-help@cern.ch
  - The General Ganga Manual
    http://ganga.web.cern.ch/ganga/
  - The Full GangaAtlas Tutorial
    https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/FullGangaAtlasTutorial
  - A GangaAtlas Quick Start Guide
    https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/GangaAtlasQuickReferenceGuide
  - GangaAtlas FAQ
    https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/DAGangaFAQ