

# DISTRIBUTED ANALYSIS USING GANGA ON THE EGEE/LCG INFRASTRUCTURE

Johannes Elmsheuser

Ludwig-Maximilians-Universität München, Germany

03 September 2007/CHEP 2007



- ① DISTRIBUTED ANALYSIS MODEL IN ATLAS
- ② DISTRIBUTED ANALYSIS WITH GANGA
- ③ CONCLUSIONS

- Heterogeneous grid environment based on 3 grid infrastructures:



- Grids have different middle-ware, replica catalogs and tools to submit jobs


⇒ Hide differences from the ATLAS user

# DISTRIBUTED ANALYSIS MODEL I

The distributed analysis model is based the ATLAS computing model

- Data is distributed in Tier1/Tier-2 facilities by default  
available 24/7
- user jobs are sent to the data  
large input datasets (100 GB up to several TB)
- Results must be made available to the user  
potentially already during processing
- Data is added with meta-data and bookkeeping in catalogs

# ATLAS DATA REPLICATION AND DISTRIBUTION



- Event Filter Farm at CERN
  - Located near the Experiment, assembles data into a stream to the Tier0
- Tier0 at CERN
  - Raw data → Mass storage at CERN and to Tier 1s
  - Swift production of Event Summary Data (ESD) and Analysis Object Data (AOD)
  - Ship ESD, AOD to Tier1s → Mass storage at CERN
- Tier1s distributed worldwide (10 centers)
  - Re-reconstruction of raw data, producing new ESD, AOD
  - Scheduled, group access to full ESD and AOD
- Tier2s distributed worldwide (~ 30 centers)
  - MC Simulation, producing ESD, AOD → Tier 1s
  - On demand user physics analysis
- CERN Analysis Facility
  - Analysis
  - Heightened access to ESD and RAW/calibration data on demand
- Tier3s distributed worldwide

# DISTRIBUTED ANALYSIS MODEL II

Need for: **Distributed Data Management (DDM)**

- Managed by DDM system DQ2 (Don-Quijote 2)
- Automated file management, distribution and archiving throughout the whole grid using a Central Catalog, FTS, LFCs
- Random access needs a pre-filtering of data of interest  
e.g. Trigger or ID streams or TAGs (event-level meta data)

**Current situation and implementation**

- Data from MC Production System is currently consolidated by DDM-operations team on all Tier1 and then all Tier2 sites
- Analysis model foresees Athena analysis of AODs/ESDs and interactive use of Athena-aware-ROOT tuples
- Analysis tuple format(s) in enhancement

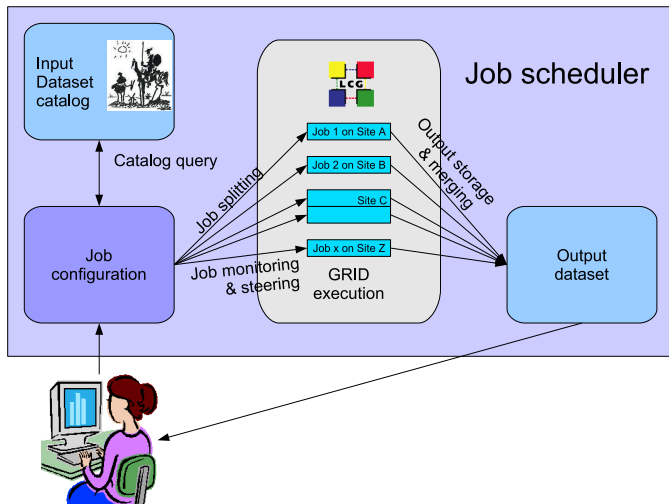
Naive assumption: Grid  $\approx$  large batch system

- Provide complicated job configuration jdl file (Job Description Language)
- Find suitable Athena software, installed as distribution kits in the Grid
- Locate the data on different storage elements
- Job splitting, monitoring and book-keeping
- etc.

$\Rightarrow$  Need for automation and integration of various different components

# DISTRIBUTED ANALYSIS

How to combine all these: Job scheduler/manager: GANGA





# FRONT-END CLIENT: GANGA



- A **user-friendly** job definition and management tool.
- Allows simple switching between testing on a **local batch system** and large-scale data processing on distributed resources (**Grid**)
- Developed in the context of **ATLAS** and **LHCb** :
  - For **ATLAS**, have built-in support for applications based on **Athena** framework, for Production System **JobTransforms**, and for **DQ2** data-management system
- **Component** architecture readily allows extension
- Python framework

# WHO IS GANGA ?

- Ganga is an ATLAS - LHCb joint project



- Support for development work from UK (PPARC/GridPP), Germany (D-Grid) and EU (EGEE/ARDA)



- Core team

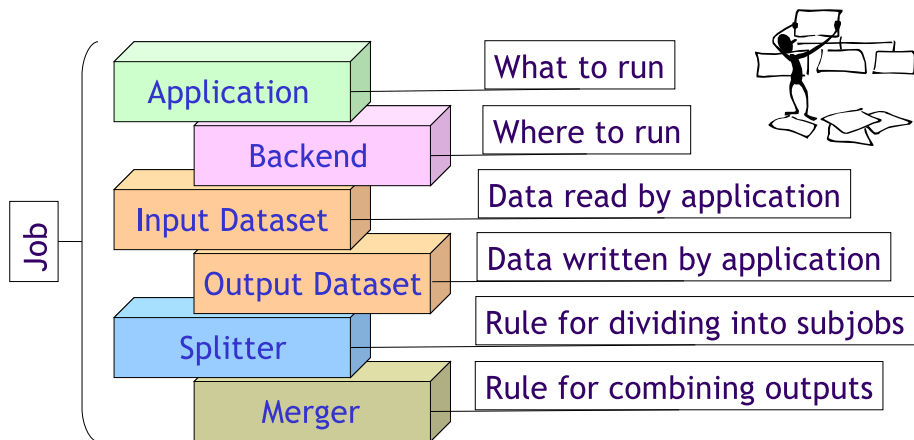
F.Brochu (Cambridge), U.Egede (Imperial), J. Elmsheuser (Munich), K.Harrison (Cambridge), H.C.Lee (ASGC Taipei), D.Liko (CERN), A.Maier (CERN), J.T.Moscicki (CERN), A.Muraru (Bucharest), W.Reece (Imperial), A.Soroko (Oxford), CL.Tan (Birmingham)



Ganga is now under GPL

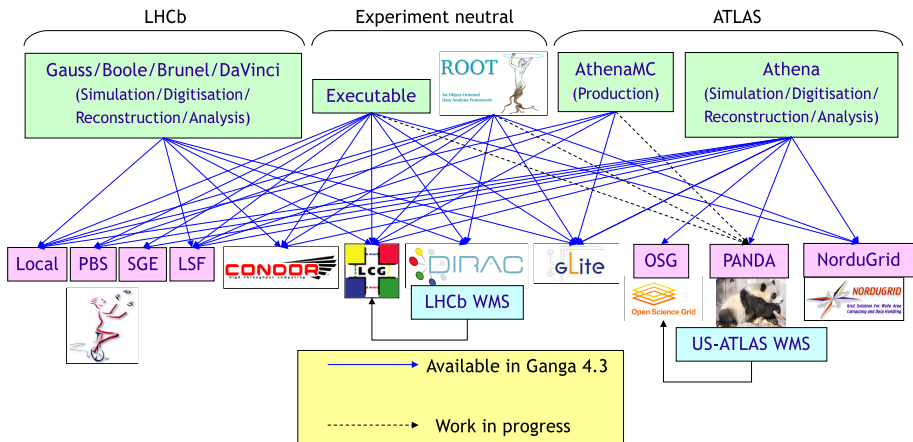
# GANGA JOB

- GANGA is based on a simple, but flexible, job abstraction
- A job is constructed from a set of building blocks, not all required for every job



# GANGA BACKENDS AND APPLICATIONS

- GANGA simplifies running of ATLAS (and LHCb) applications on a variety of Grid and non-Grid back-ends



# JOB DEFINITION USING ATLAS SOFTWARE I

GANGA offers three ways of user interaction:

- Shell command line
- Interactive IPython shell
- Graphical User Interface

Job definition at command line for GRID submission:

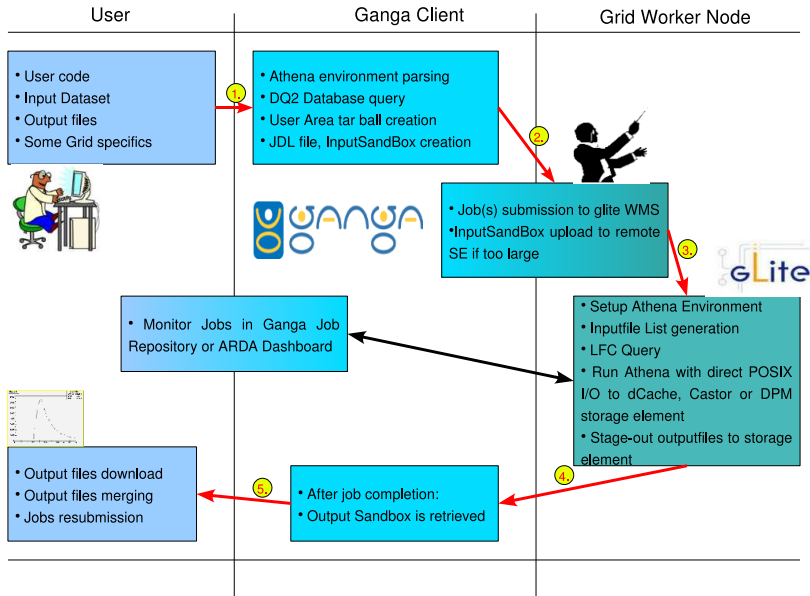
```
ganga athena
--inDS csc11.005320.PythiaH170wwll.recon.AOD.v11004107
--outputdata AnalysisSkeleton.aan.root
--split 3
--lcg
AnalysisSkeleton_topOptions.py
```

## JOB DEFINITION USING ATLAS SOFTWARE II

Job definition within GANGA IPython shell:

```
j = Job()
j.application=Athena()
j.application.prepare(athena_compile=False)
j.application.option_file='$HOME/athena/12.0.5/InstallArea/jobOptions
j.splitter=AthenaSplitterJob()
j.splitter.numsubjobs = 3
j.merger=AthenaOutputMerger()
j.inputdata=DQ2Dataset()
j.inputdata.dataset='csc11.005145.PythiaZmumu.recon.AOD.v11004103'
j.inputdata.match_ce=True
j.outputdata=DQ2OutputDataset()
j.outputdata.outputdata=['AnalysisSkeleton.aan.root']
j.backend=LCG()
j.submit()
```


# JOB WORKFLOW



# GANGA JOBS MONITORED BY THE DASHBOARD

<http://dashb-atlas-job.cern.ch/dashboard/request.py/jobsummary>

JOB SUMMARY



user:	current status							grid exit status				application exit status			overall	
	Sub	Unk	Pend	Run	Term	Done	Canc	Abort	Unk	Grid%	Succ	Fail	Unk	App%	D/S	Overall%
Bitte Domnik/unknown	604	0	209	21	374	267	2	105	0	71.77	213	36	126	85.54	211	56.42
Brigitte Epp/unknown	10	0	0	10	10	0	0	0	0	100	2	8	0	20	2	20
Carlos Escobar Ibanez/unknown	4	0	0	0	4	4	0	0	0	100	4	0	0	100	4	100
Christos Anastopoulos/unknown	21	0	0	0	21	17	4	0	0	100	9	9	3	50	9	42.86
Elena Oliver Garcia/unknown	3	0	0	0	3	3	0	0	0	100	3	0	0	100	3	100
Esteban Fullana Torregrosa/unknown	6	0	0	0	6	6	0	0	0	100	0	8	0	0	0	0
Farida Fassi/unknown	6	0	0	0	6	6	0	0	0	100	6	0	0	100	6	100
Folkert Koetsveld/unknown	3	0	0	0	3	0	0	3	0	0	2	1	0	66.67	0	0
Grant Gorfine 9015/unknown	4	0	0	0	4	4	0	0	0	100	1	3	0	25	1	25
Jean-Baptiste De vivie de regie/unknown	13	0	0	0	13	13	0	0	0	100	10	3	0	76.92	10	76.92
Johannes Ebke/unknown	119	0	0	0	119	116	0	3	0	97.48	21	15	89	58.33	21	17.65
Johannes Elmsheuser/unknown	14	0	0	0	14	14	0	0	0	100	11	3	0	78.57	11	78.57
total	6924	0	209	176	6539	7666	602	434	87	94.65	4927	1403	2209	77.64	4686	56.99



# GANGA USAGE ON THE GRID

- GANGA has been used in the last 4 month on LCG at:

Tier	0	1	2	3
Usage [%]	8	37	40	15

where at Tier1: 48 % Lyon and 36% FZK

⇒ After successful dataset replication, following the ATLAS user analysis computing model

- NorduGrid: backend for ARC submission is used, integrated with Athena and DQ2
  - OSG/Panda: recently integrated a backend for OSG/Panda
- ⇒ Now all ATLAS Grid flavours are supported

# DISTRIBUTED ANALYSIS TUTORIALS AND SUPPORT


<https://twiki.cern.ch/twiki/bin/view/Atlas/GangaTutorial43>

- Edinburgh (February 1st-2nd )
- Milan (February 5th-6th )
- Lyon (March 5th-7th)
- Munich (March 29th-30th )
- Toronto (April 18th)
- Bergen (April 27th)
- Valencia (May 3rd-4th)
- DESY Zeuthen (Sep 18-19th)
- **Ganga User support and Feedback via Hypernews**  
[hn-atlas-GANGAUserDeveloper@cern.ch](mailto:hn-atlas-GANGAUserDeveloper@cern.ch)



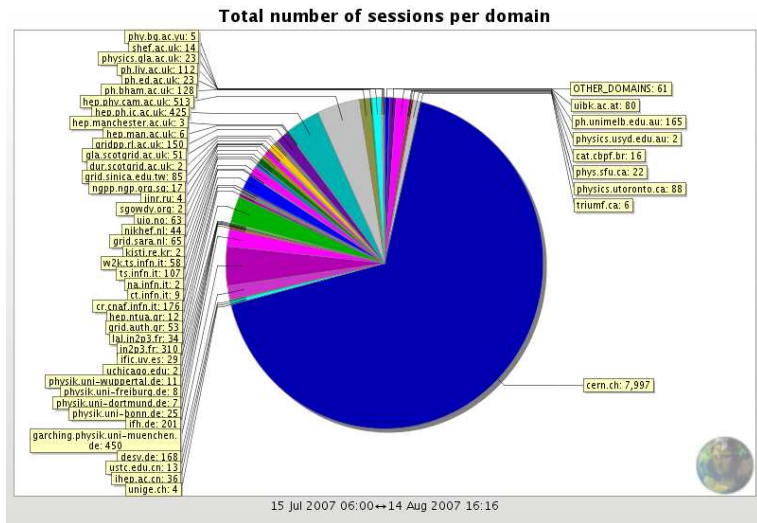
# USAGE STATISTICS

- Over 820 unique users since beginning of the year
- About 500 ATLAS users have tried Ganga at least once
- About 60 ATLAS Ganga users per week



# USAGE OF GANGA AT REMOTE SITES

- GANGA has been used at over 50 different sites in the last month



# RECENT ASPECTS OF DISTRIBUTED ANALYSIS USING GANGA

- Many groups are able to use GANGA for day to day work
- User feedback drives the development
- Increasing stability and functionality of the middle-ware, but still room for improvements
  
- Data distribution with complete datasets is a key issue
- Problems with setup during SLC4 migration
- Site specific problems need to be chased up

# DISTRIBUTED ANALYSIS NEEDS AND CONCLUSIONS

For the distributed analysis it is vital to have:

- Easy interface that does not scare off physicists
- A reliable and robust service of many components

Conclusions

- Growing number of users are using the Grid for analysis,  
Still some room to grow
- Data Management is a central issue for Distributed Analysis

Homepage:

- <http://cern.ch/ganga>

References:

- GANGA depends on many components - for details see talks this week