



# Grid Computing in LHCb

Or:  
How  
I  
Learned  
To  
Stop  
Worrying  
And  
Love  
Distributed  
Analysis




## Outline

- The LHCb Computing Model
- Ganga
- Examples and user experience

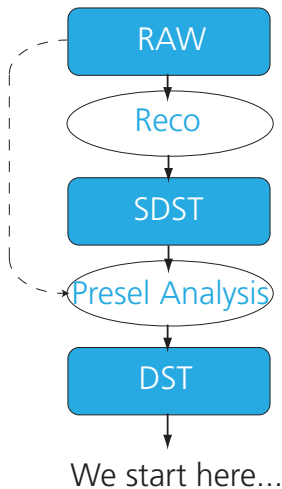


## LHCb Computing Model



- Every Tier 2 belongs to one Tier 1 (in first order).
- No contact of user with Tier 2.
- #Tier 1 : #Tier 2 = 1 : 2

## From RAW to DST



What comes out of the detector.

RAW data is reconstructed  
(Calo clusters, tracks, ...).

Stripping DST contain information  
to cut on physical properties.

Preselection Analysis is run to split  
data into subsets and throw away  
uninteresting events.

DSTs are starting point for analysis.

## Ganga - GAudi aNd Grid Alliance

- Ganga is a frontend to submit jobs.
- "Configure once, run anywhere", meaning: Once you have decided what job to run, you can send it wherever you want.
- Shared between LHCb and ATLAS.
- Written in python, easily adaptable by users.

Possible places to run the job:

- Local: Interactive or Background.
- Batch System: LSF, PBS, ...
- LCG via Dirac.



# Ganga user statistics

From: 2010-08-11

To: 2010-09-01

Experiment: LHCb

Total number sessions: 7718

Number unique users: 208


Number of sites: 25

Sessions summary

Users summary


Time charts

## Ganga Usage Statistics



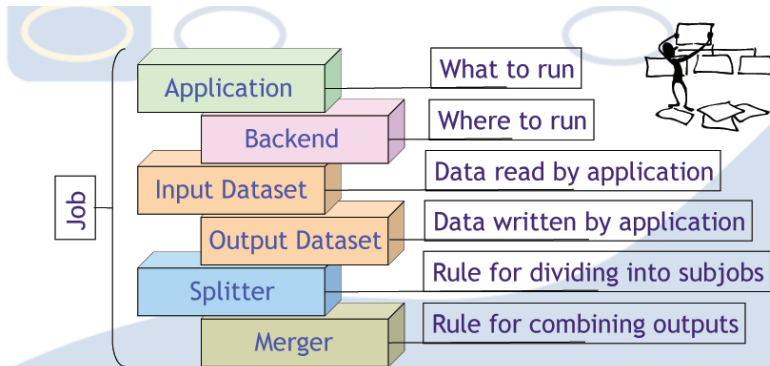
## Ganga installations

Ganga Sessions of Non-CERN Ganga sites



## Ganga options

Decide the following:



## Example: Reconstruct $J/\Psi \rightarrow \mu\mu$

```
#####  
# Script to submit job to the Grid via Ganga  
#####  
  
# Application  
t = JobTemplate( application = DaVinci( version = "v25r7p1" ))  
  
# Job steering options  
t.application.optsfile = [  
"/home/hep/decianm/cmtuser/DaVinci_v25r7p1/gridtest/gridtest.py"]  
  
# Give as input the datafiles which should be run over  
t.application.optsfile.append("/home/hep/decianm/cmtuser/DaVinci_v25r7p1/gridtest/DiMuonReco05Stripping08LFN.py")  
  
# Platform  
t.application.platform = "x86_64-slc5-gcc43-opt" # if submitting to  
the GRID (Dirac)  
  
# Create a job for the DIRAC backend  
j = Job( t, backend = Dirac())  
  
# Specify the output data  
j.outputdata = [ 'Jpsi.root' ] # if you want to return data via the  
outputdir  
  
# Split the job in subjobs  
j.splitter=DiracSplitter(filesPerJob=30)  
  
# Submit the job!  
j.submit()
```



# Submit the jobs...

```

sah
--jedi-- 1 decima hep 0810 Aug 28 09:25 job1000
--jedi-- 1 decima hep 1020 Aug 28 17:30 gr4000_0y
--jedi-- 1 decima hep 2411 Aug 25 15:55 job1001r00f
--jedi-- 1 decima hep 2649 Aug 25 15:56 job1001r00f
--jedi-- 1 decima hep 2654 Aug 25 15:56 job1002r00f
--jedi-- 1 decima hep 2697 Aug 25 15:56 job1003r00f
--jedi-- 1 decima hep 2152 Aug 25 15:57 job1003r00f
--jedi-- 1 decima hep 1868 Aug 28 15:58 [iljxw@ecdf03stripping@Prescaled.FH.y
--jedi-- 1 decima hep 1870 Aug 28 15:58 gr1000_0y
[decima@grid-01]~/vafuser/Software/200701/gr1000_0y$ ganga -i gangaub.y

*** Welcome to Ganga ***
Version: Ganga-6.5-0
Documentation and support: http://cern.ch/ganga
Type help() or help('index') for online help.

This is free software (GPL), and you are welcome to redistribute it
under certain conditions; type license() for details.

Hello from Ganga at University of Zurich.
Ganga.GPDev.Lib_300      1 INFO      submitting job 40
Ganga.GPDev.Adapters    1 INFO      job 40 status changed to "submitting"
Ganga.GPDev.Lib_300      1 INFO      submitting job 40.0 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.0 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.1 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.1 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.2 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.2 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.3 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.3 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.4 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.4 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.5 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.5 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.6 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.6 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.7 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.7 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.8 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.8 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.9 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.9 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.10 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.10 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.11 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.11 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.12 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.12 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.13 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.13 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.14 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.14 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.15 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.15 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.16 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.16 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.17 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.17 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.18 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.18 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.19 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.19 status changed to "submitting"
Ganga.GPDev.Adapters    1 INFO      submitting job 40.20 to Siroc backend
Ganga.GPDev.Lib_300      1 INFO      job 40.20 status changed to "submitting"
Ganga.GPDev.Lib_300      1 INFO      job 40 status changed to "submitted"

```

15:10: All jobs submitted



# ...wait until they have finished...

```

file = ["./srcroot"],
(location =
)
file = jobfile <
timeofapp =
Time (UTC) Status
-----
2010/09/20 15:06:20 - new
2010/09/20 15:06:20 - submitting
2010/09/20 15:07:21 - submitted
2010/09/20 15:11:02 - backend_running
2010/09/20 15:14:00 - running
2010/09/20 15:41:30 - backend_completing
2010/09/20 15:41:41 - backend_fini
2010/09/20 15:41:45 - completing
2010/09/20 15:46:00 - final
}
splitter = DirectSplitter (
filesPerJob = 5,
jobFiles = -1,
ignoreMissing = false
)
subJobs = Registry Bise1 jobs(40) subJobs (21 objects)

```

jobId	status	name	subJobs	application	backend	backend.actualICE
job_0	completed			001	001	001
job_1	completed			001	001	001
job_2	completed			001	001	001
job_3	completed			001	001	001
job_4	completed			001	001	001
job_5	completed			001	001	001
job_6	completed			001	001	001
job_7	completed			001	001	001
job_8	completed			001	001	001
job_9	completed			001	001	001
job_10	completed			001	001	001
job_11	completed			001	001	001
job_12	completed			001	001	001
job_13	completed			001	001	001
job_14	completed			001	001	001
job_15	completed			001	001	001
job_16	completed			001	001	001
job_17	completed			001	001	001
job_18	completed			001	001	001
job_19	completed			001	001	001
job_20	failed			001	001	001
job_21	failed			001	001	001
job_22	failed			001	001	001
job_23	failed			001	001	001
job_24	failed			001	001	001
job_25	failed			001	001	001
job_26	failed			001	001	001
job_27	failed			001	001	001
job_28	failed			001	001	001
job_29	failed			001	001	001
job_30	failed			001	001	001
job_31	failed			001	001	001
job_32	failed			001	001	001
job_33	failed			001	001	001
job_34	failed			001	001	001
job_35	failed			001	001	001
job_36	failed			001	001	001
job_37	failed			001	001	001
job_38	failed			001	001	001
job_39	failed			001	001	001
job_40	failed			001	001	001


```

backend = Direct
status = None
actualICE = None
withName = ("DIRT1", 172088)
name3DTime = None
diracData = ""
statusInfo = ""
id = None
inputBackendNo = []
}
}
CompGPIDevLib-300 : INFO job 40-10 status changed to "Completed"
CompGPIDevLib-300 : INFO job 40-10 status changed to "Failed"
CompGPIDevLib-300 : INFO job 40-20 status changed to "Completing"
CompGPIDevLib-300 : INFO job 40-20 status changed to "Completed"
CompGPIDevLib-300 : INFO job 40 status changed to "Failed"

```

17:46: All jobs completed or failed

## ...and download the data.



- After downloading the tuples from the SEs we can plot the result.
- Summary:
  - 20 jobs sent to the Grid.
  - Got tuples back 2.5 hours later.
  - 10 % of jobs failed.
  - However, this was a rather easy test...

## The reality

Let's start with the positive:

- If distributed analysis works as foreseen, it is fast and convenient. It is (almost) no difference to a job running on your local machine.
- Ganga is now convenient to use and quite stable.
- People are very helpful and things have improved a lot over the last year.

## The reality

However...

- When a job crashes locally (and it's your own fault):  $\mathcal{O}(\min)$  to fix it.
- When a job crashes on the Grid (and it's your own fault):  $\mathcal{O}(h)$  to fix it.
- When a job crashes (and it's not your own fault) it is very hard to figure out the problem / solution.
  - Was the Storage Element / Computing Element not available (and the job got sent there anyway)?
  - Is there a problem with the database?
  - ...
- With no major problem, 90-95 % of the jobs finish successfully, rest has to be resubmitted. It rarely happens that all jobs finish successfully.
- Will it ever converge to a really stable operation?

**"Marge, I agree with you - in theory. In theory, communism works. In theory."**