

# Grid applications at CREATIS

7 juillet 2009


Tristan Glatard (CR CNRS 2008), Sorina Pop (IR CNRS 2007)  
SA1 meeting

<http://www.creatis.insa-lyon.fr>

# Applications

- GATE radiotherapy simulation


- ~250 tasks per run
- Days -> hours
- /O=GRID-FR/C=FR/O=CNRS/OU=CREATIS/CN=David Sarrut



ThIS Simulation


- FIELD ultrasound simulation

- matlab
- ~1000/2000 tasks
- 16h -> 3h
- /O=GRID-FR/C=FR/O=CNRS/OU=CREATIS/CN=Carlos Gines Fuster




- CAVIAR cardiac analysis

- ~100/500 tasks
- /O=GRID-FR/C=FR/O=CNRS/OU=CREATIS/CN=Christopher Casta




Segmentation



Motion estimation

- SIMRI MRI simulation


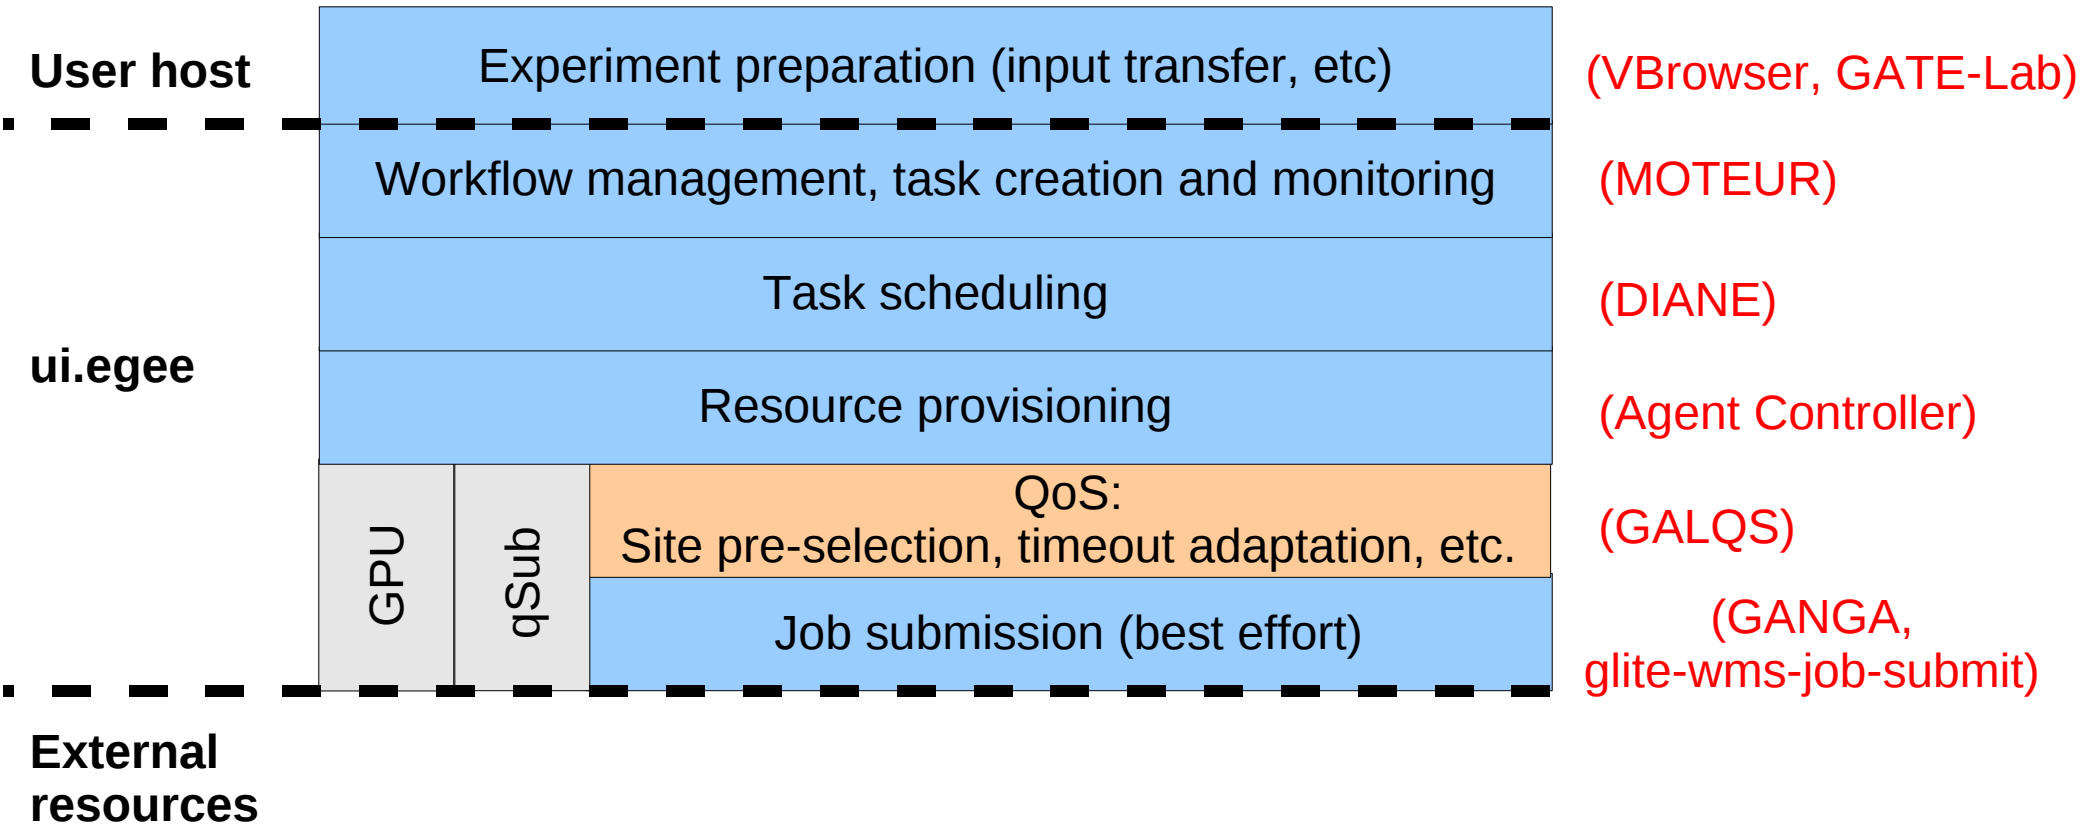
- MPI
- ~20% success
- /O=GRID-FR/C=FR/O=CNRS/OU=CREATIS/CN=Hugues Benoit-Cattin



# Outline

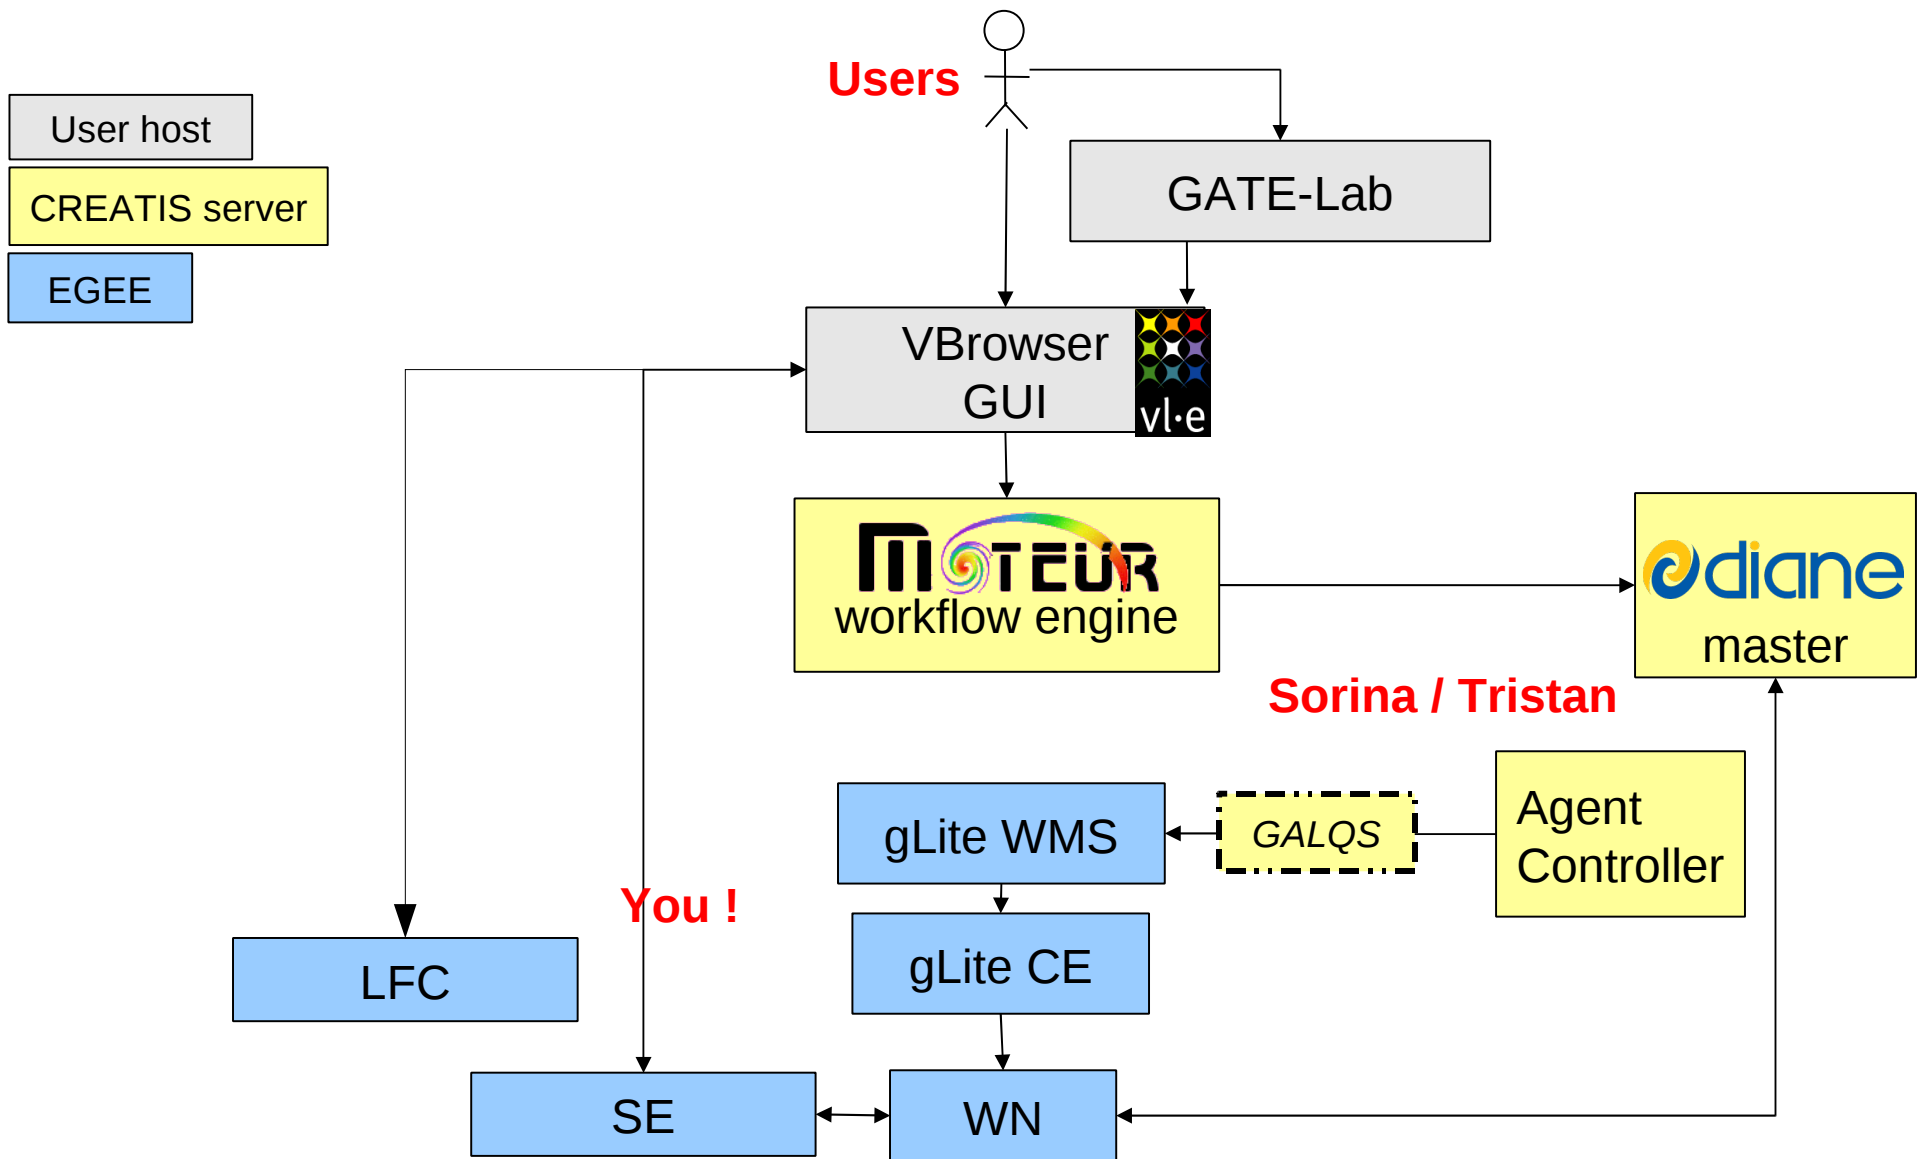
- Grid environment
  - Execution environment
  - Monitoring
  - Application porting
- Activity report
  - Some statistics
  - Errors / main issues

# Execution stack



(meta-)data ?

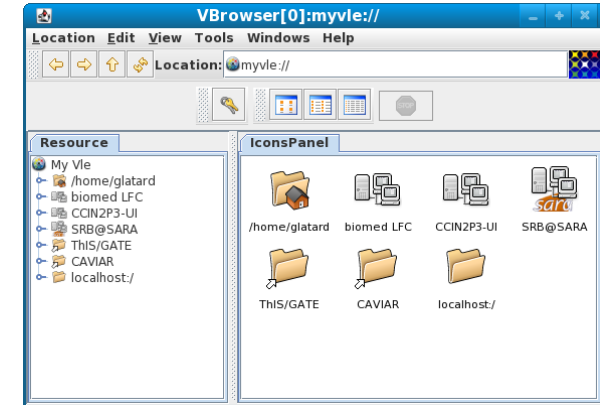
# Execution environment



# Graphical interface to the grid

- VBrowser

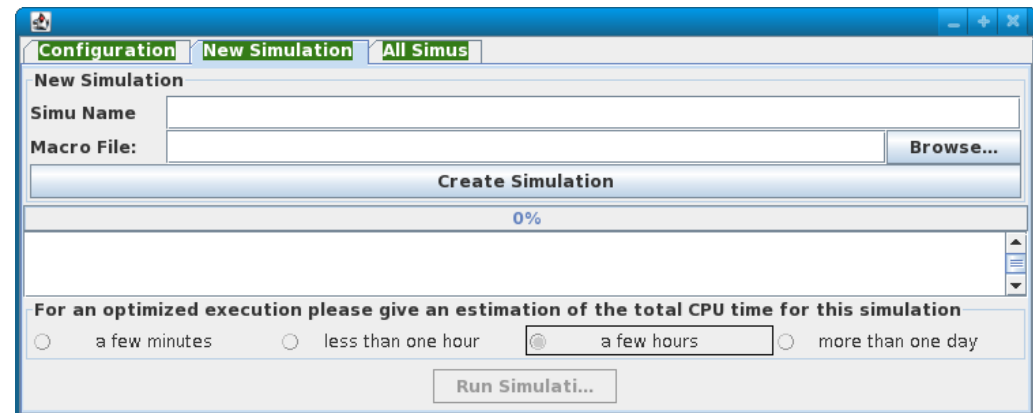
- ↗ Provides user interface to the grid
- ↗ Create VOMS proxy from certificate
- ↗ Read/write files on LFC+SE
- ↗ Browses input/result files
- ↗ Launches workflows (calls moteur server)
- ↗ <http://www.vl-e.nl/vbrowser>



VBrowser

- Application-specific: GATE-Lab

- ↗ GATE-Lab client (VBrowser plugin)
  - Parses simulation (mac) file, zip inputs
  - Stores inputs on the grid, submits workflow
  - Keeps track of simulation history



Gate-Lab client

**Resource**

- My Vle
  - /home/glatard
  - biomedLFC
  - SRB@sara
  - viz-login
  - gwendia\_cardiac
  - vlemed
    - alex
    - amygdLR
    - amygdLR-enc
    - atlas
    - demo-VLFMs
    - generated
  - glatard
  - jalkemade
  - jeroene
  - kboulebiar
  - martin
  - matthan
  - mdm
  - piter.t.de.boer
  - remi
    - data\_storage
    - db
    - joblogs
    - masks
    - output
    - output-may-2008
    - output\_26-08
    - scripts
    - workflows
      - group
      - individual
        - iaps
        - nback
      - groupAnalyzes
      - individualAnalyzes
      - inputs
      - flirt\_and\_roi\_indiv.s
  - http
    - results
    - silvia
    - testVFSLFC
    - testVFSLFC2
    - tristan
    - wibisono
    - garbage.sh
    - hello-1228916611960604456.t
  - ccU
  - Desktop
  - applisCreatis

**Grid files  
browsing,  
reading and  
writing**

**CobraViewer**

Status [Services](#) [Input](#) [Results](#) [Info](#)

**Workflow monitoring**

```

    graph TD
      indivAnalysis --> roi
      indivAnalysis --> flirtIndiv
      flirtIndiv --> zstat2standard
      roi --> roiIndiv
      flirtIndiv --> roiIndiv
  
```

roiIndiv  
done:0  
running:10  
failed:0

flirtIndiv  
done:2  
running:0  
failed:0

zstat2standard

**Job monitoring**

**JOB STATUS:workflow-PTIv75**

Configuration JobStatus

Update Periodic Refresh(minute): 1

Job Actions

Retrieve Output Cancel jobs Select/Unselect

N#	JobID	JobStatus	link Out	S
1	https://rb.grid.sara.nl:9000...	DONE (SUC...	Not yet...	
2	https://rb.grid.sara.nl:9000...	DONE (SUC...	Not yet...	
3	https://rb.grid.sara.nl:9000...	SCHEDULED	Not yet...	
4	https://rb.grid.sara.nl:9000...	READY	Not yet...	
5	https://rb.grid.sara.nl:9000...	READY	Not yet...	
6	https://rb.grid.sara.nl:9000...	READY	Not yet...	
7	https://rb.grid.sara.nl:9000...	WAITING	Not yet...	

**Workflow inputs**

Load from file Save to file Add Parameter List Add Parameter Range Add Parameter Tag Path Delete selected

<input type="checkbox"/>	Name:	indivAnalysis	Group:	Value:	lfn://lfc.grid.sara.nl/grid/vlemed/remi/output/feat-dofhigh-12
<input type="checkbox"/>	Name:	roi	Group:	Value:	lfn://lfc.grid.sara.nl/grid/vlemed/remi/amygdLR_bin.nii.gz

Run Workflow Web service URL: https://ws1.grid.sara.nl/~glatard/workflow/workflow-PTIv75/workf

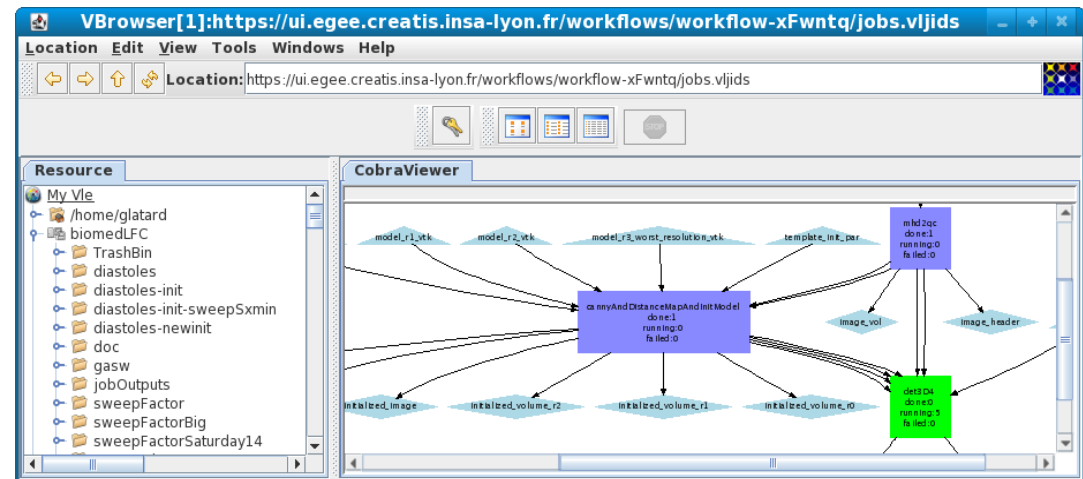
*Creatis*

# Workflow execution

- Application description
  - XML document (Scufl, Taverna)
  - Describes data and control links between components
    - Independent from middleware (DIANE, WMS, GALQS, ARC, G5K)
  - Intrinsic parallelism
  - Structures results (provenance, annotation)



- Engine configuration parameters
  - Middleware
  - Timeout
  - Retry count
  - SE where to write results

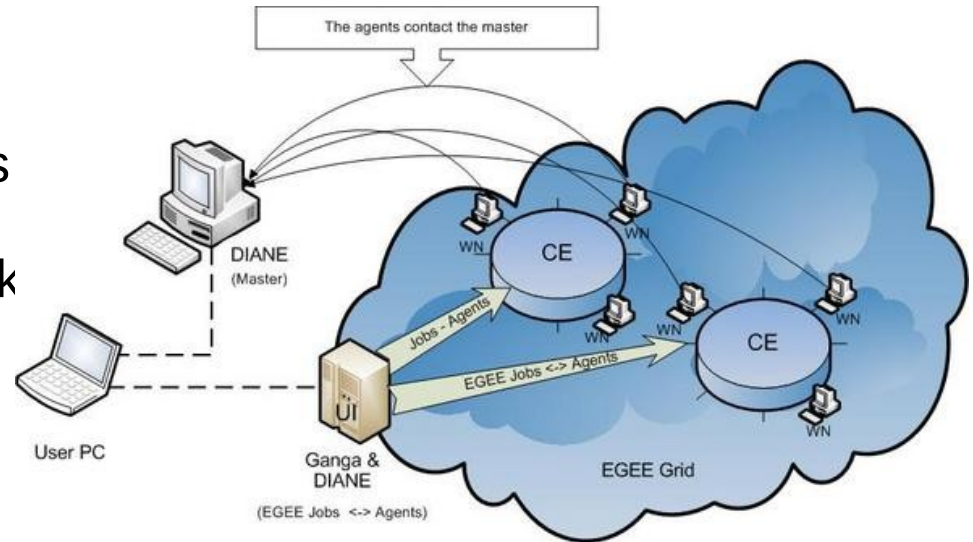


- Execution
  - Engine installed on UI machine
  - Web-Service interface with proxy delegation
  - Called by VBrowsers plugin



# Reliability and responsiveness

- DIANE master
  - **1 master per user**
  - Registers/removes agents
  - Schedules tasks on agents
  - Stdout/err transfers
  - Unregister agent when task



- Agent controller
  - Submits agents to WMS using GANGA
  - Submission algorithm

```
init=10, defaultSleep=400, s=defaultSleep, maxSub=300, factor=5
start DIANE master
submit init pilots
while master is alive do
  sleep s seconds
  n = number of scheduled tasks in master
  submit sub=min(maxSub,n) pilots
  s = defaultSleep+sub*factor
end while
```

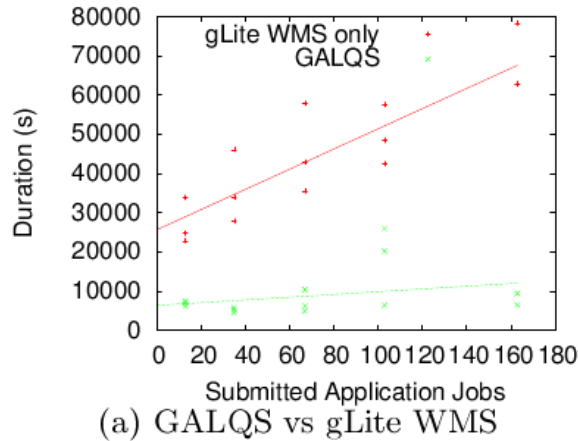
# GALQS (in progress)

- Grid Application-Level QoS Server
- Monitoring
  - Sends probe jobs to check application execution time/correctness
  - Maintain database of error ratios and latencies per CEs
- Site pre-selection
  - Queried before submission to WMS -> JDL adaptation
  - Application feedback to be included

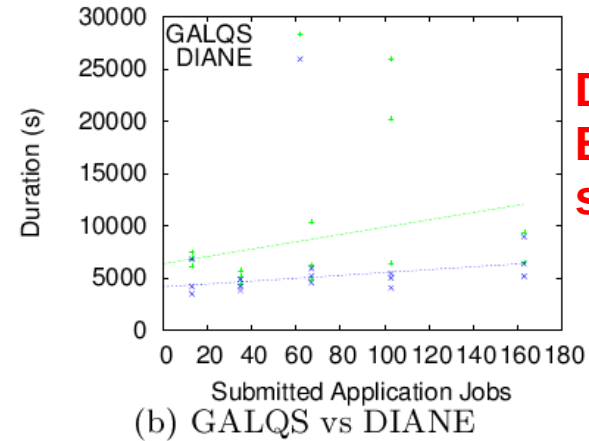
# GALQS evaluation

- Application: CAVIAR (~12-min jobs)
- Metric: time to reach 100% success

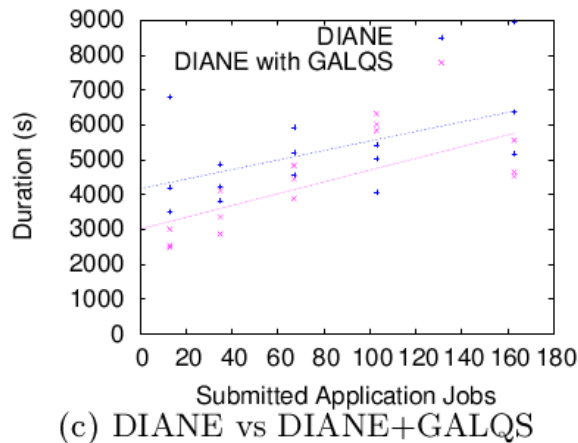
**WMS vs GALQS  
Throughput x 5**




**DIANE vs GALQS  
Equivalent in  
some cases**



**DIANE vs  
DIANE+GALQS**



**Strong error  
reduction**

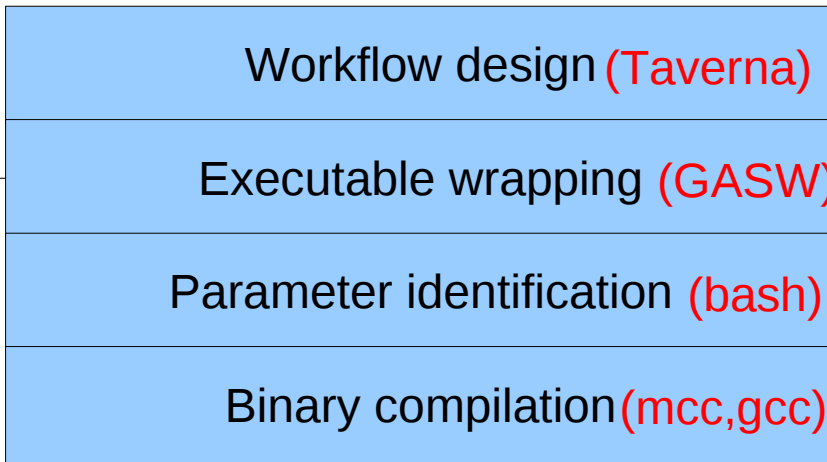
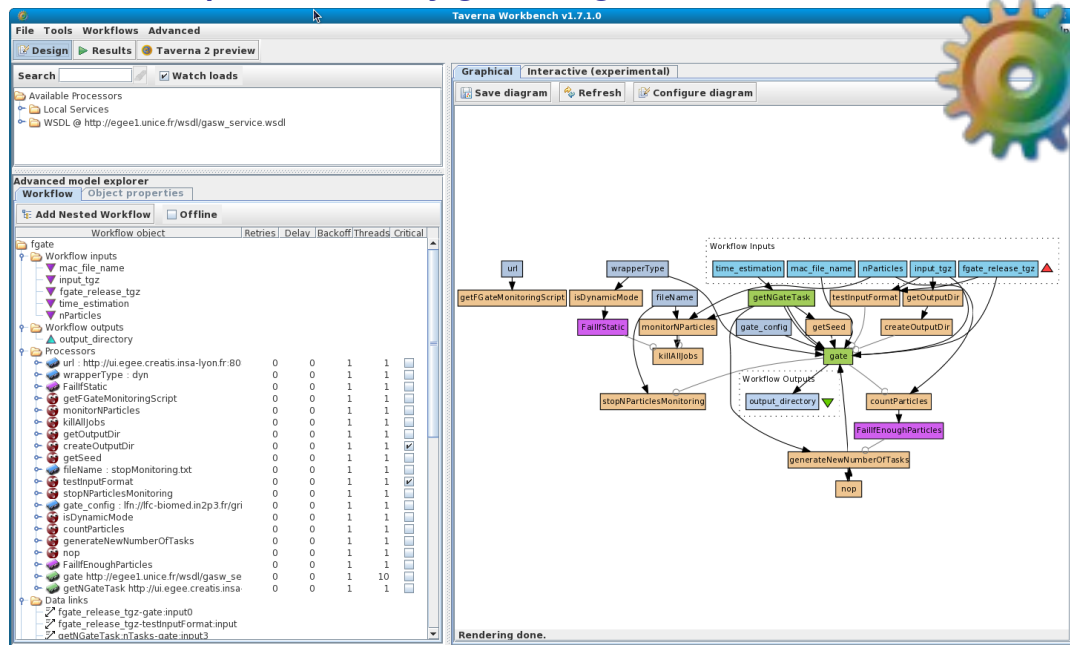


# Application porting stack

<http://www.mygrid.org.uk/tools/taverna/>

```

<description>
<executable name="CrestLines.pl">
  <access type="URL">
    <path value="http://colors.unice.fr:80/" />
  </access>
  <value value="CrestLines.pl" />
  <input name="image" option="-im1">
    <access type="LFN" />
  </input>
  <input name="scale" option="-s" />
  <output name="crest_lines" option="-c2">
    <access type="LFN" />
  </output>
  <sandbox name="convert8bits">
    <access type="URL">
      <path value="http://colors.unice.fr:80/" />
    </access>
    <value value="Convert8bits.pl" />
  </sandbox>
</executable>
</description>
  
```



# Grid activity monitoring

- DIANE masters dashboard

Actions	User	MASTERS				AGENTS				TASKS					STATS			
		Date launched	Workspace	Port	Alive	Submitted	Running	Registered	Removed	Submitted	Scheduled	Running	Done	Failed	Worktime	Running	Download	Upload
Kill!	Tristan Glatard/	Jul 7 08:53	/var/www/diane/runs/0392	23003	true	26	11	18	7	0	0	4	21	4	11h:4m:49s	95.5%	1.6%	2.7%
Kill!	David Sarrut/	Jul 2 11:18	/var/www/diane/runs/0380	23005	true	757	0	713	713	0	0	0	759	99	133h:38m:42s	79.4%	9.7%	10.7%
Kill!	Sorina Camarasu/	Jun 29 13:40	/var/www/diane/runs/0379	23004	true	260	0	172	172	0	0	0	184	37	85h:13m:13s	87.9%	8.0%	3.9%
Cleanup	Carlos Gines Fuster/	Jun 26 16:32	/var/www/diane/runs/0366	23005	false	30	0	1	1	0	0	0	5	0	0h:19m:16s	21.7%	74.7%	3.5%

- Submitted / running agents per user
- Tasks statuses
- Transfer VS running time

# Workflow dashboard


- Current user activity

Report from Jul 1th to the last day of Jul


Actions	Date	User	Workflow ID	Total grid jobs	Done application jobs	Failed application jobs	Global success ratio	Application success ratio
Kill	Jul 7 08:54	Tristan Glatard/	workflow-XZ73UQ	29	21	0	72.00%	100.00%
Cleanup	Jul 6 18:44	Tristan Glatard/	workflow-DQ2ZGT	59	50	0	84.00%	100.00%
Cleanup	Jul 6 17:01	Tristan Glatard/	workflow-UUhZJF	52	1	0	1.00%	100.00%

- Detailed statistics


## Job flow



## Up/download histograms

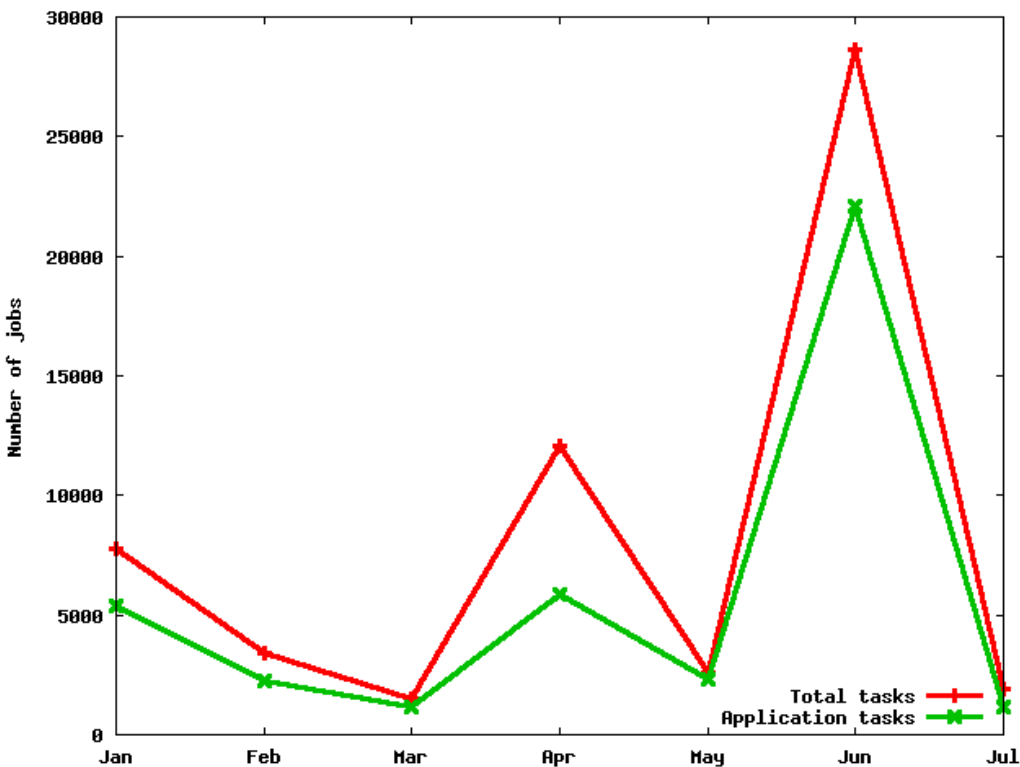


## Running times

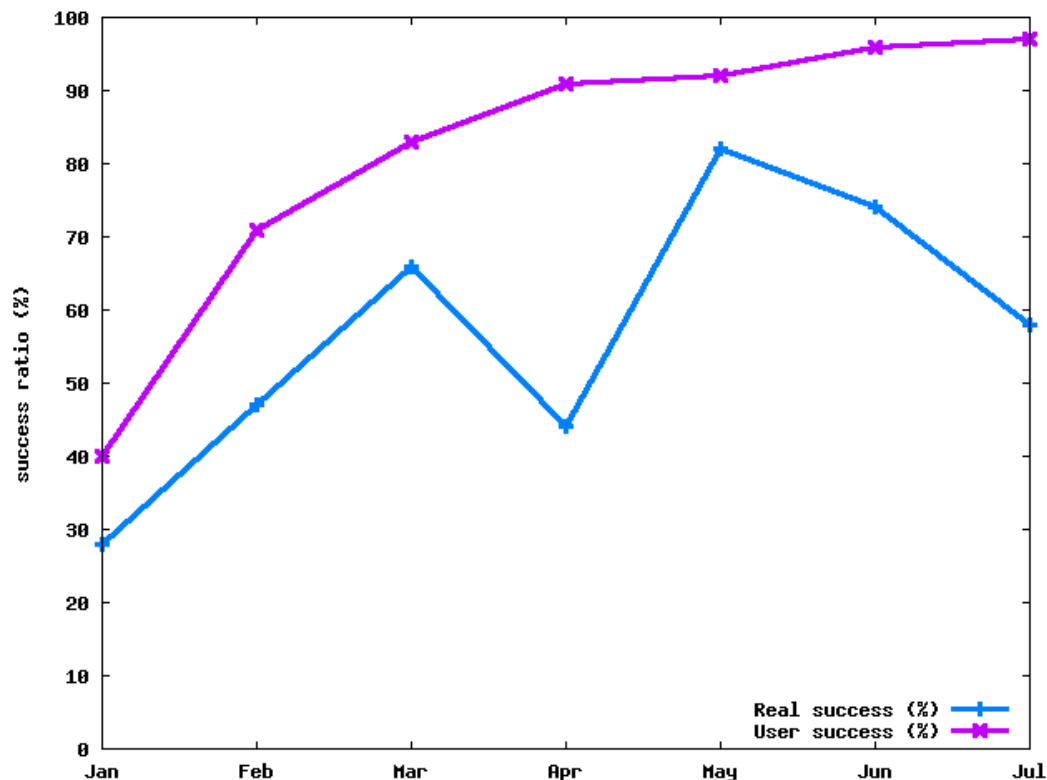


# Activity report (1/1/09 - 7/7/09)

- Job numbers



## Success rate



- 1285 workflows in total

# Main issues / challenges

- Access to the grid
  - Certificate creation and VO registration
- Site heterogeneity
  - Data transfers from A to B fail
  - SE not in BDII
  - Application-specific errors (unzip not found, not enough RAM, segfault, etc)
  - => application-specific dynamic CE selection
- Scalability
  - LFC (should be solved now)
  - Disk quota / number of connections to SEs
  - Slow agent registration (<250 concurrent agents)



# Conclusion

- Increasing grid activity / job success rate
  - Starts being usable for end-users
  - Projects starting based on this environment
- Administrative cost of execution environment
  - Deploy applications
  - Monitor user activity
  - Submit / follow-up GGUS tickets
  - Improve setup

=> closer applications / SA1 interaction ?

# Merci !

## Questions ? Suggestions ?

- Acknowledgment:
  - Applications: David Sarrut, Patrick Clarysse, Denis Friboulet, Hervé Liebgott, Hugues Benoit-Cattin, Christopher Casta
  - Grid environment: Carlos Gines Fuster, Alejandro Tovar de Duenas, Wen-Jun Tan
  - All the anonymous teams behind <https://gus.fzk.de>

# Discussion with CPPM

- Connections

- CPPM: “please reduce your number of concurrent connections to our SE”
- Creatis: “what is the limit and how could I control it ?”

- Disk quota

- CPPM: “You are using 833 GB on our SE, this is too important”
- Creatis: “Where is the quota defined and how could I control it ?”

- Central SE

- CPPM: “You should read/write from/on close SEs”
- Creatis: “1. Read: users put their data on some SEs and *then* submit jobs ; 2. Write: writing on > 200 SEs greatly increases errors”