

Scaling bio-analyses from computational clusters to grids

George Byelas

University Medical Centre Groningen, the Netherlands

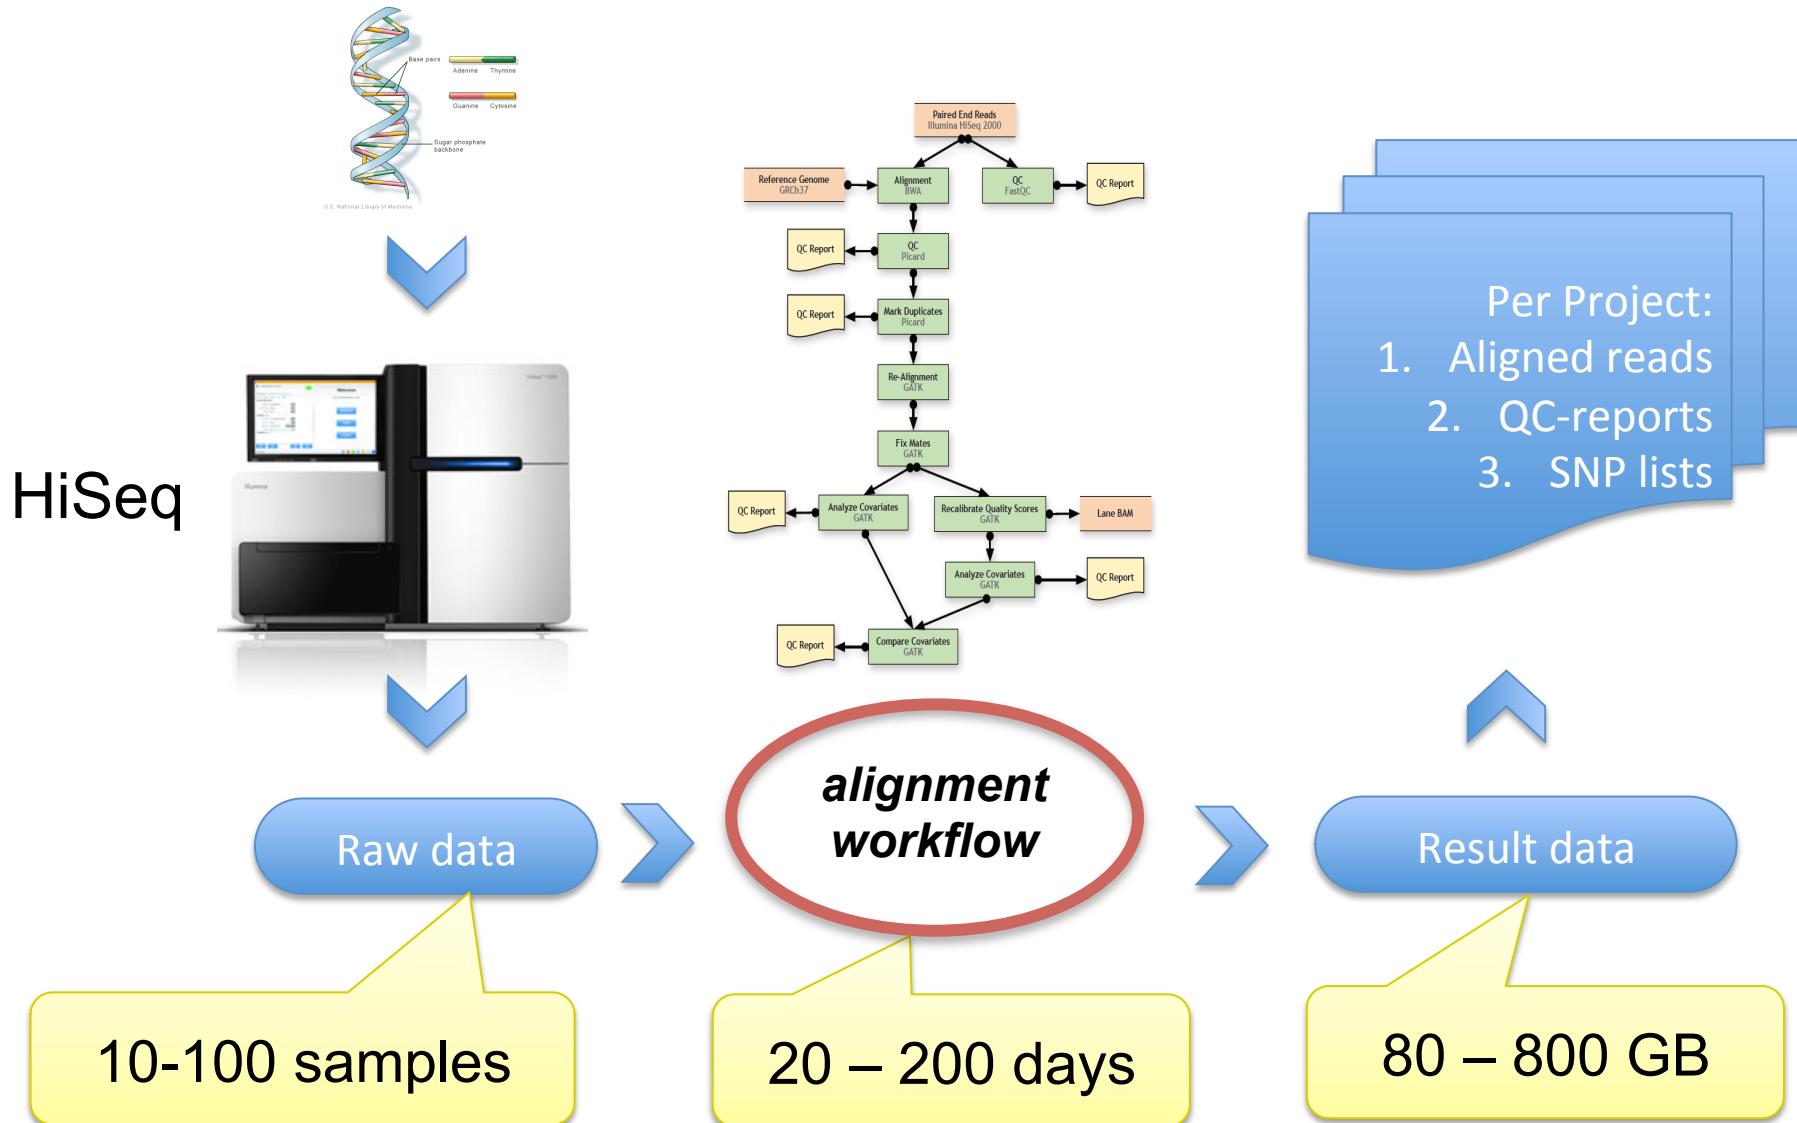
IWSG-2013, Zürich, Switzerland, June 3rd, 2013

Content

- Bio-workflows
- Workflow
 - Modeling
 - Job Generation
 - Tool deployment
 - Data management
 - Workflow execution
 - Implementation detail
 - Recent developments
- Conclusion and further steps

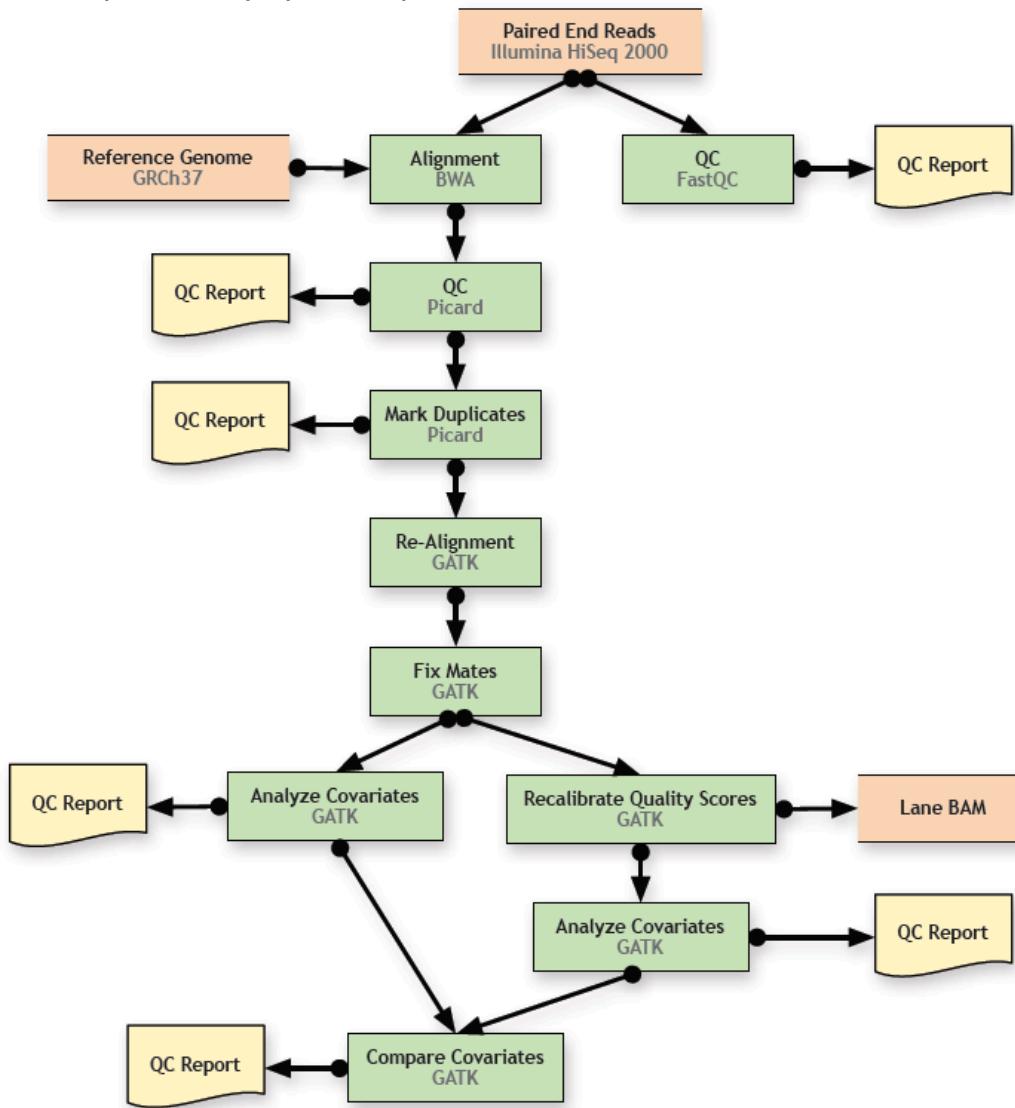
Bio-workflows

Example: NGS alignment workflow



Alignment & SNP calling workflow

31 steps, ≥ 2 days per sample



- **Input**
 - Analysis protocols
 - Sample DNA data
 - Reference DNA data
- **Analysis**
 - Scripts are generated and executed
- **Output**
 - Aligned DNA and QC reports

An analysis job (script) generated from a protocol

```
#!/bin/bash
#PBS -q test
#PBS -l nodes=1:ppn=4
#PBS -l walltime=08:00:00
#PBS -l mem=6gb
#PBS -e $GCC/test_compute/projects/batch4/intermediate/test1/err/err_test1_BwaElement1A102a_FC81D90ABXX_L7.err
#PBS -o $GCC/test_compute/projects/batch4/intermediate/test1/out/out_test1_BwaElement1A102a_FC81D90ABXX_L7.out
mkdir -p $GCC/test_compute/projects/batch4/intermediate/test1/err
mkdir -p $GCC/test_compute/projects/batch4/intermediate/test1/out
printf "test1_BwaElement1A102a_FC81D90ABXX_L7_started " >>$GCC/test_compute/projects/batch4/intermediate/test1/log_test1.txt
date "+DATE: %m/%d/%y%tTIME: %H:%M:%S" >>$GCC/test_compute/projects/batch4/intermediate/test1/log_test1.txt
date "+start time: %m/%d/%y%t %H:%M:%S" >>$GCC/test_compute/projects/batch4/intermediate/test1/
test1_BwaElement1A102a_FC81D90ABXX_L7.txt
echo running on node: `hostname` >>$GCC/test_compute/projects/batch4/intermediate/test1/
test1_BwaElement1A102a_FC81D90ABXX_L7.txt
```

backend specific

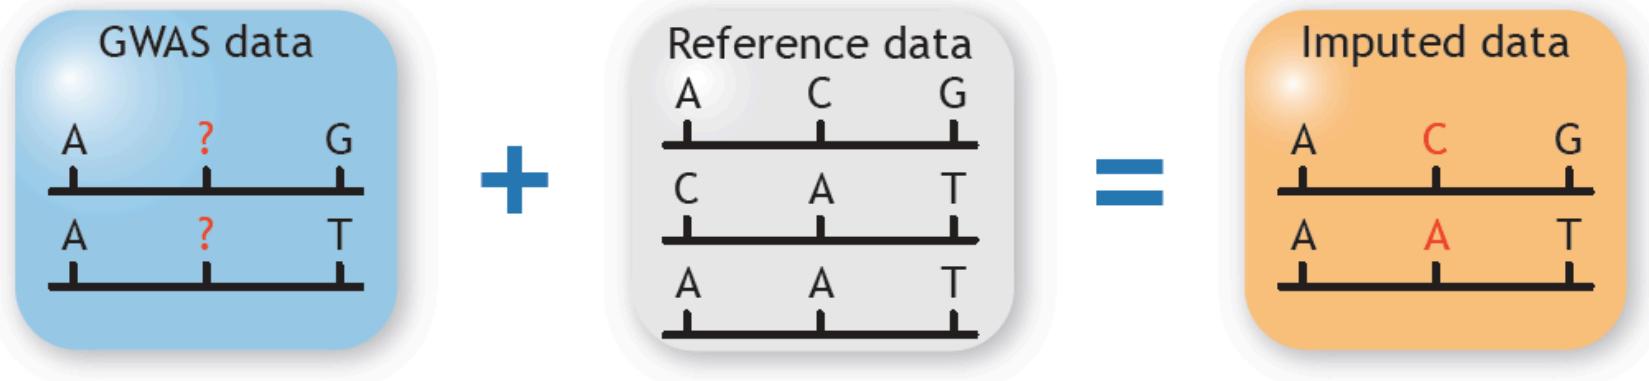
```
/target/gpfs2/gcc/tools//bwa-0.5.8c_patched/bwa aln \
/target/gpfs2/gcc/resources/hg19/indices/human_g1k_v37.fa \
$GCC/test_compute/projects/batch4/rawdata/110121_I288_FC81D90ABXX_L7_HUMrutRGADIAAPE_1.fq.gz \
-t 4 \
-f $GCC/test_compute/projects/batch4/intermediate/A102a_110121_I288_FC81D90ABXX_L7_HUMrutRGADIAAPE_1.fq.gz.sai
```

analysis specific

```
printf "test1_BwaElement1A102a_FC81D90ABXX_L7_finished " >>$GCC/test_compute/projects/batch4/intermediate/test1/log_test1.txt
date "+finish time: %m/%d/%y%t %H:%M:%S" >>$GCC/test_compute/projects/batch4/intermediate/test1/
test1_BwaElement1A102a_FC81D90ABXX_L7.txt
date "+DATE: %m/%d/%y%tTIME: %H:%M:%S" >>$GCC/test_compute/projects/batch4/intermediate/test1/log_test1.txt
```

Imputation workflow

- Imputation:



- Number of jobs

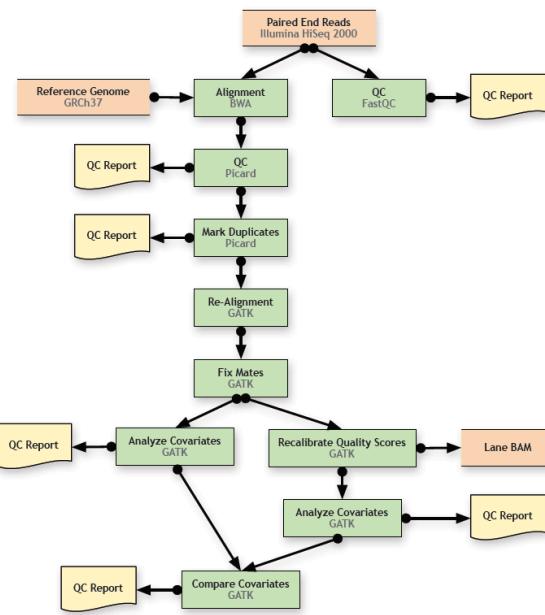
$$\text{number_jobs} = \sum_{chr=1..22} \frac{\text{Length}_{chr}}{5\text{-megabase}} * \frac{\text{number_of_samples}}{500}$$

- One run:



Bio-workflow complexity

- Many analysis steps
 - Many analysis jobs
 - Different analysis tools and their dependencies
- Large various data involved
- Heterogeneous resources

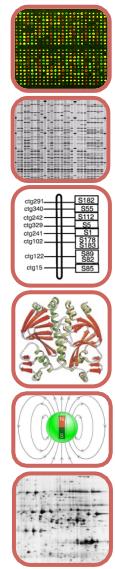


Workflow design and generation

MOLGENIS approach

- Model
- Generate
- Use

Analyses...



nbioassist

B B M R I • N L
• • • • •

EMBL-EBI

B B M R I • N L
• • • • •

university of
groningen



university of
groningen

WAGENINGEN UR
For quality of life

umcg

umcg

University of
Leicester

University of
Leicester

Projects ...



Species...

GEN2PHEN

GEN2PHEN

LIFELINES

FIMM

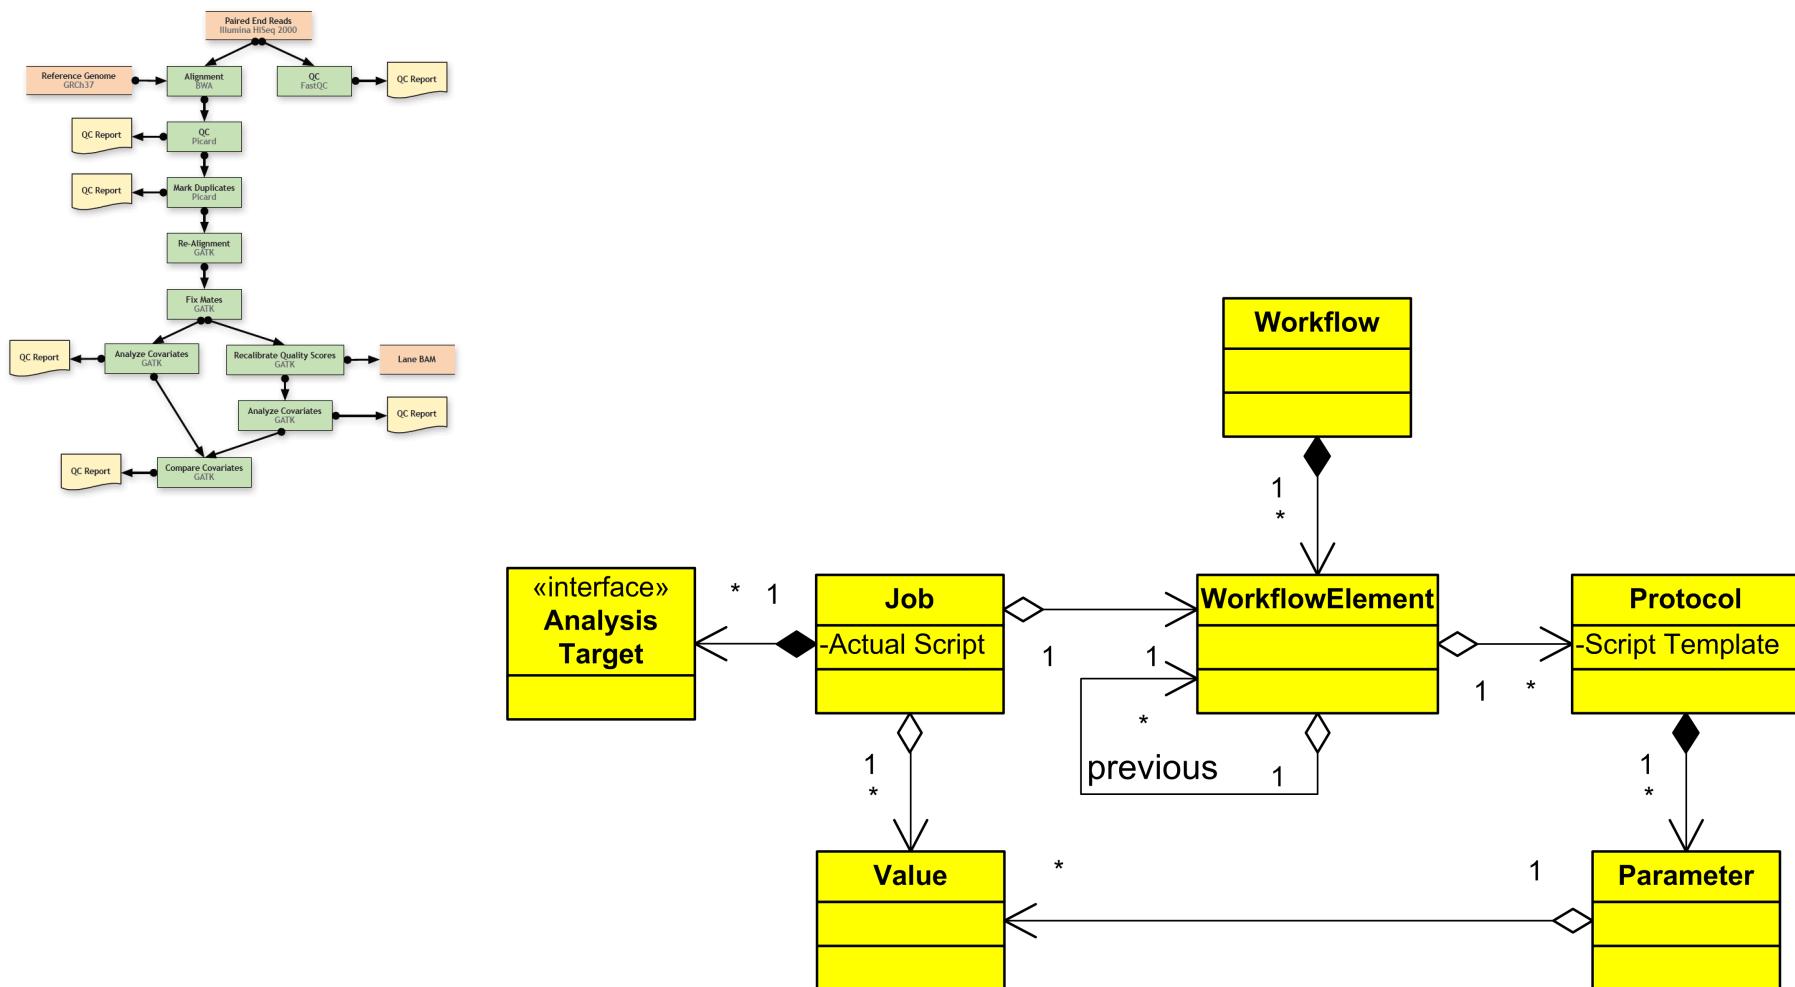
SYSGENET

CASIMIR

IOP

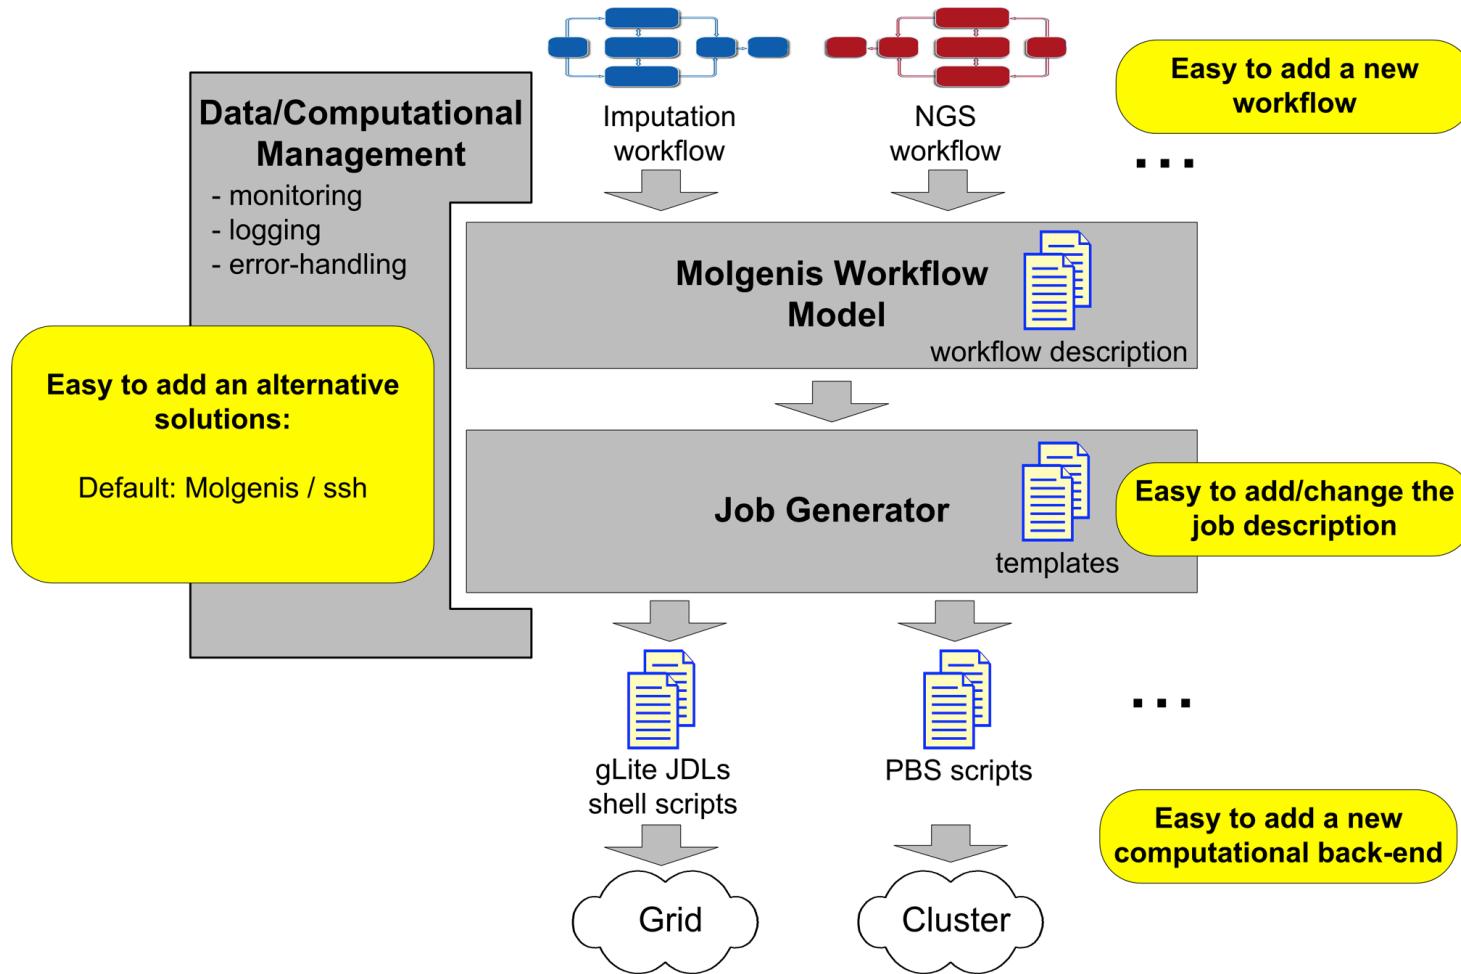
PANACEA

Workflow design



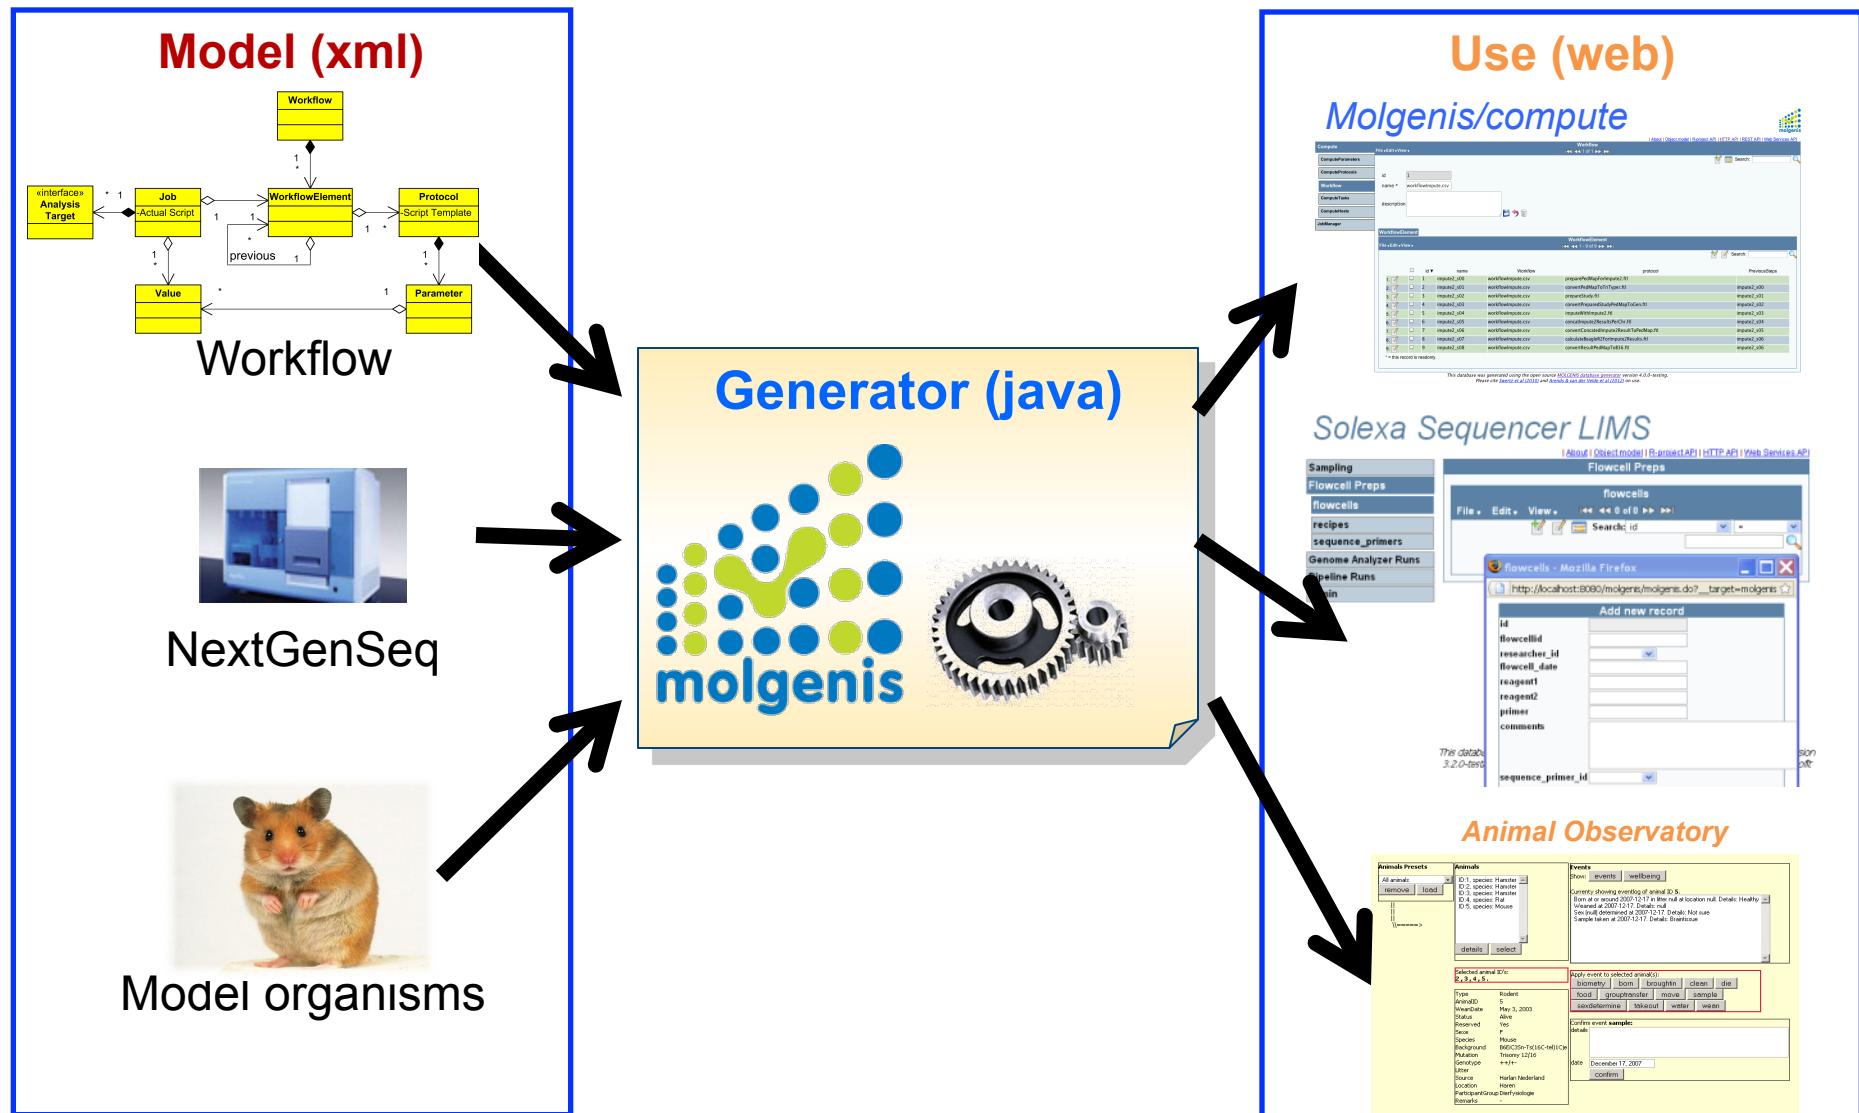
- Jobs are generated from the model
- Every job has an analysis target (e.g. Genome region)

Command-line generator (Demo @ IWSG-2012)

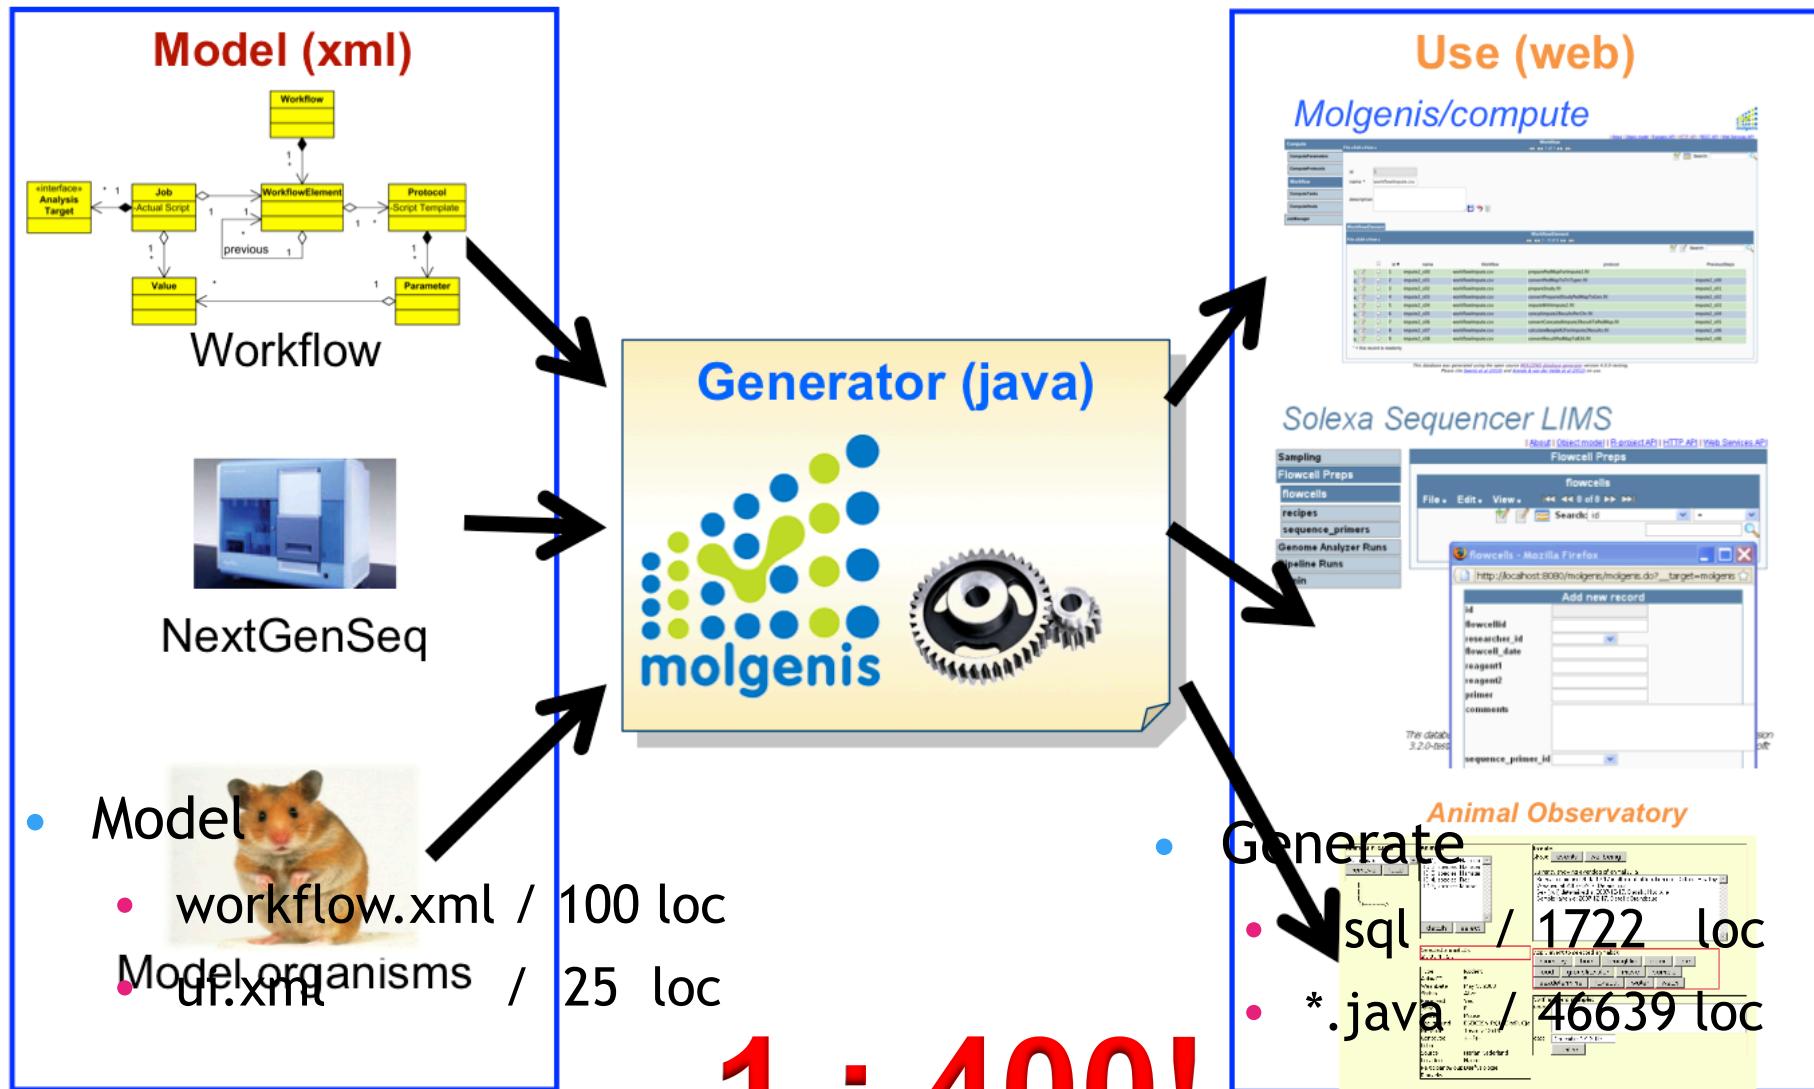


- Generates jobs (scripts) from model described in files
- Suitable for **workflows** (PBS cluster) and **single jobs** (gLite grid)

Database solution with MOLGENIS software toolkit (1)



Database solution with MOLGENIS software toolkit (2)



Workflow design view in the generated Molgenis web-UI

Compute Testing



Workflow

WorkflowElement

WorkflowElement

Protocol

PreviousSteps

	ID	Name	Workflow	Protocol	PreviousSteps
1.	impute2_s00	workflowImpute.csv	preparePedMapForImpute2.ftl		impute2_s00
2.	impute2_s01	workflowImpute.csv	convertPedMapToTriTyper.ftl		impute2_s01
3.	impute2_s02	workflowImpute.csv	prepareStudy.ftl		impute2_s02
4.	impute2_s03	workflowImpute.csv	convertPreparedStudyPedMapToGen.ftl		impute2_s03
5.	impute2_s04	workflowImpute.csv	imputeWithImpute2.ftl		impute2_s04
6.	impute2_s05	workflowImpute.csv	concatImpute2ResultsPerChr.ftl		impute2_s05
7.	impute2_s06	workflowImpute.csv	convertConcatenatedImpute2ResultToPedMap.ftl		impute2_s06
8.	impute2_s07	workflowImpute.csv	calculateBriggleR2ForImpute2Results.ftl		impute2_s07
9.	impute2_s08	workflowImpute.csv	convertImpute2ResultsToB36.ftl		impute2_s08

* = this record is readonly.

This was generated by the generator version 4.0.0-testing. See [the source code](#) for more information.

workflow
step

analysis
protocol

previous
steps

Workflow run-time view (analysis jobs)

Compute

- [ComputeParameters](#)
- [ComputeProtocols](#)
- [Workflow](#)
- [ComputeTasks](#)
- [ComputeHosts](#)
- [JobManager](#)

File • Edit • View •

ComputeTasks

1 - 20 of 21 >> >>

id	name	ComputeScript	RunLog	WorkflowElement	Interpreter	PrevSteps	requirements	Status	Code
1.	impute2_s00_test1_1349185453555253000	##### BEFORE ##### touch SPBS_O_WORKDIR/jobname.out source \$WORKDIR/tools/scripts/import.sh before	TASKID:impute2_s00_test1_1349185453555253000 touch: cannot touch '/opt/glite/var/tmp/jobname.out': N_-	impute2_s00	bash	impute2_s00	1	done	done
2.	impute2_s01_test1_1349185453792052000	##### BEFORE ##### touch SPBS_O_WORKDIR/jobname.out source \$WORKDIR/tools/scripts/import.sh before	TASKID:impute2_s01_test1_1349185453792052000 touch: cannot touch '/opt/glite/var/tmp/jobname.out': N_-	impute2_s01	bash	impute2_s01	1_13491854537920520001	done	done
3.	impute2_s02_test1_1349185453886012000	##### BEFORE ##### touch SPBS_O_WORKDIR/jobname.out source \$WORKDIR/tools/scripts/import.sh before	TASKID:impute2_s02_test1_1349185453886012000 touch: cannot touch '/opt/glite/var/tmp/jobname.out': N_-	impute2_s02	bash	impute2_s01	1_13491854537920520001	done	done
4.	impute2_s03_test1_1349185454031690000	##### BEFORE ##### touch SPBS_O_WORKDIR/jobname.out source \$WORKDIR/tools/scripts/import.sh before	TASKID:impute2_s03_test1_1349185454031690000 touch: cannot touch '/opt/glite/var/tmp/jobname.out': N_-	impute2_s03	bash	impute2_s02	1_13491854538860120001	done	done
5.	impute2_s04_test1_134918545425596								
6.	impute2_s04_test1_1349185								
7.	impute2_s04_test1_1349185								
8.	impute2_s04_test1								
9.	impute2_s04_test								
10.	impute2_s04_test								
11.	impute2_s04_test								
12.	impute2_s04_test								
13.	impute2_s04_test								
14.	impute2_s04_test								
15.	impute2_s04_test								
16.	impute2_s04_test								
17.	impute2_s04_test								
18.	impute2_s05_test								
19.	impute2_s06_test								
20.	impute2_s07_test								

* this record is readonly.

Task Log:

```

TASKID:BwaAlignLeft_run01_1355507698165482000
Running on node: wn-car-004.farm.nikhef.nl
srmcp -server_mode=passive srm://srm.grid.sara.nl/pnfs/grid.sara.nl/data/bbmri.nl/RP2/resources/hg19/indices/human_g1k_v37.chr1.fa
file:///tmp/jobdir/28851242.stro.nikhef.nl/resources/hg19/indices/human_g1k_v37.chr1.fa
Picked up _JAVA_OPTIONS: -Xmx512M
server_mode is specified as passive, setting number of streams to 1
srmcp -server_mode=passive srm://srm.grid.sara.nl/pnfs/grid.sara.nl/data/bbmri.nl/RP2/resources/hg19/indices/human_g1k_v37.chr1.fa.amb
file:///tmp/jobdir/28851242.stro.nikhef.nl/resources/hg19/indices/human_g1k_v37.chr1.fa.amb
Picked up _JAVA_OPTIONS: -Xmx512M
server_mode is specified as passive, setting number of streams to 1
srmcp -server_mode=passive srm://srm.grid.sara.nl/pnfs/grid.sara.nl/data/bbmri.nl/RP2/resources/hg19/indices/human_g1k_v37.chr1.fa.ann
file:///tmp/jobdir/28851242.stro.nikhef.nl/resources/hg19/indices/human_g1k_v37.chr1.fa.ann
Picked up _JAVA_OPTIONS: -Xmx512M
server_mode is specified as passive, setting number of streams to 1
srmcp -server_mode=passive srm://srm.grid.sara.nl/pnfs/grid.sara.nl/data/bbmri.nl/RP2/resources/hg19/indices/human_g1k_v37.chr1.fa.bwt
file:///tmp/jobdir/28851242.stro.nikhef.nl/resources/hg19/indices/human_g1k_v37.chr1.fa.bwt
Picked up _JAVA_OPTIONS: -Xmx512M
server_mode is specified as passive, setting number of streams to 1
srmcp -server_mode=passive srm://srm.grid.sara.nl/pnfs/grid.sara.nl/data/bbmri.nl/RP2/resources/hg19/indices/human_g1k_v37.chr1.fa.fai
file:///tmp/jobdir/28851242.stro.nikhef.nl/resources/hg19/indices/human_g1k_v37.chr1.fa.fai
Picked up _JAVA_OPTIONS: -Xmx512M
server_mode is specified as passive, setting number of streams to 1
srmcp -server_mode=passive srm://srm.grid.sara.nl/pnfs/grid.sara.nl/data/bbmri.nl/RP2/resources/hg19/indices/human_g1k_v37.chr1.fa.pac
file:///tmp/jobdir/28851242.stro.nikhef.nl/resources/hg19/indices/human_g1k_v37.chr1.fa.pac
Picked up _JAVA_OPTIONS: -Xmx512M
server_mode is specified as passive, setting number of streams to 1
srmcp -server_mode=passive srm://srm.grid.sara.nl/pnfs/grid.sara.nl/data/bbmri.nl/RP2/resources/hg19/indices/human_g1k_v37.chr1.fa.rbwt
file:///tmp/jobdir/28851242.stro.nikhef.nl/resources/hg19/indices/human_g1k_v37.chr1.fa.rbwt
Picked up _JAVA_OPTIONS: -Xmx512M
server_mode is specified as passive, setting number of streams to 1

```

statusCode
done

Failed jobs overview

chr: 4 from: 185000001 to: 190000001

```
grep "ERROR: There are no type 2 SNPs after applying the command-line settings for awk '{print $1}'" == 1 ] then if [ ! -f ${tmpOutput}_info ] then echo "Impute2 found empty output" echo "Touching file: ${tmpOutput}" echo 'Touching file: ${tmpOutput}_info' fi
```

Running on node: v33-45.gina.sara.nl

```
Error: terminate called after throwing an instance of 'std::bad_alloc'  
what(): St9bad_alloc
```

How much memory:
virtual memory (kbytes, -v) 4194304

Search results where: **StatusCode = failed** ✖

Interpreter	PrevSteps	requirements	StatusCode
bash		1	failed

its where: **StatusCode = failed** ✖

requirements StatusCode failed

echo: 190000001

... we now create a touching file:

TaskHistory |◀◀◀ 1 - 10 of 16 ▶▶▶|

Running on node: v33-45.gina.sara.nl

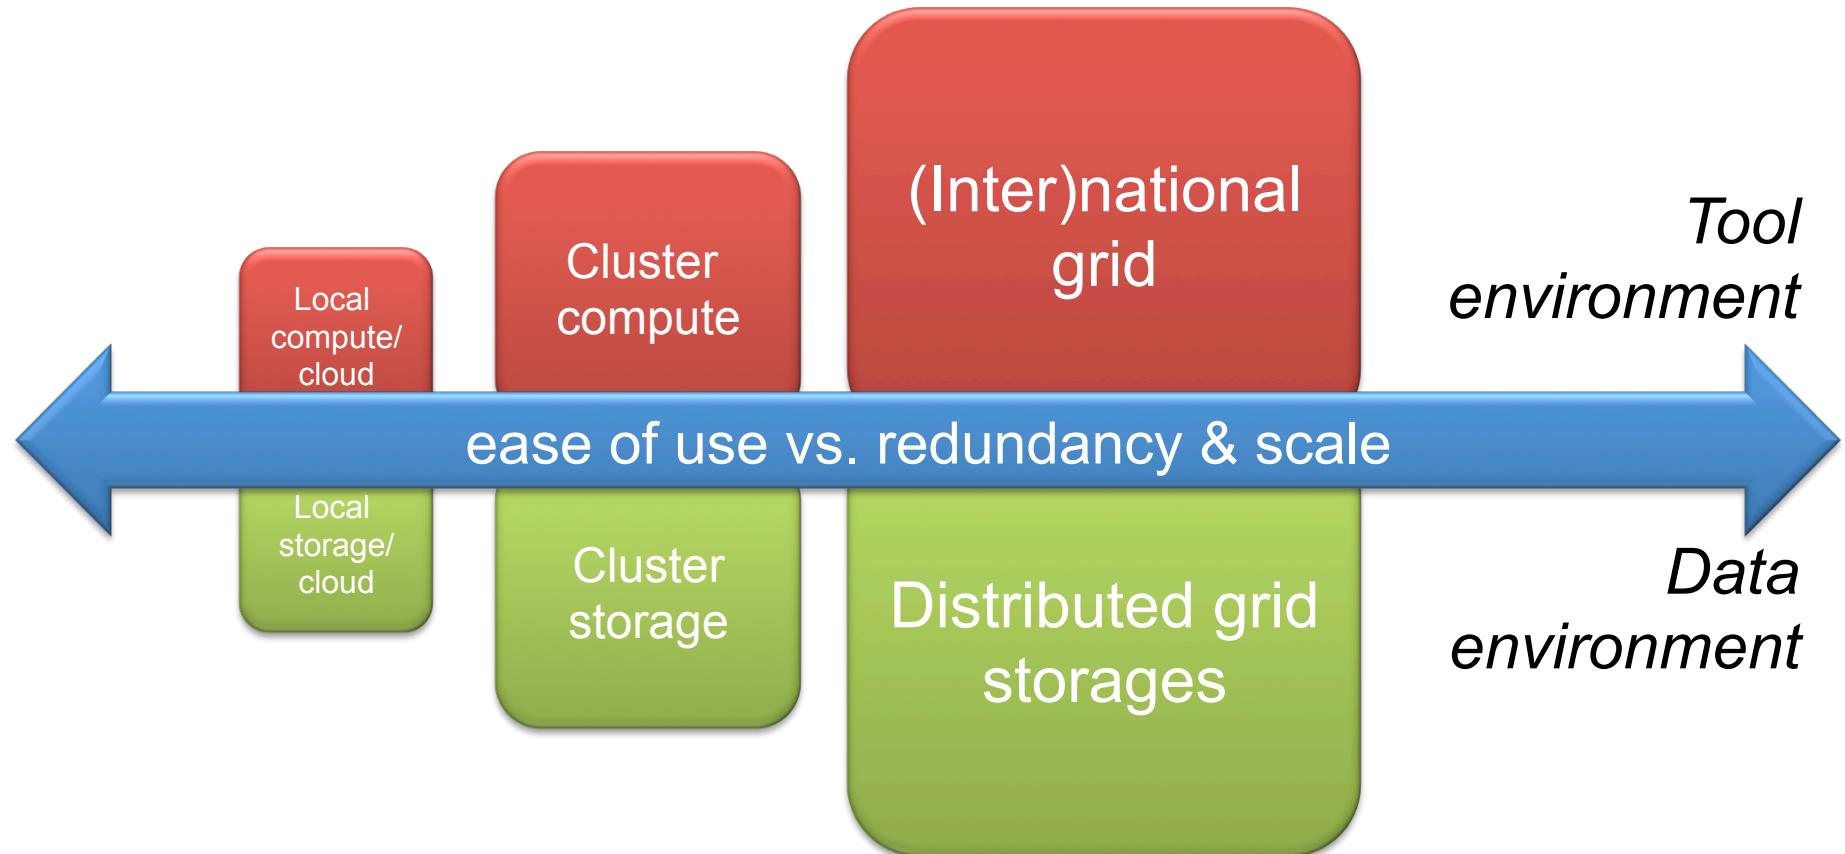
	id	ComputeTask	ComputeWorker	StatusCode	StatusTime	
1.						
2.						
3.						
4.						
5.						
6.						
7.						
8.	6740	imputation_run01_6664677232470394		failed	April 6, 2013, 07:47:18	TASKID:imputation_run01_6664677232470394
9.	6951	imputation_run01_6664677232470394		failed	April 6, 2013, 14:31:48	TASKID:imputation_run01_6664677232470394
10.	7134	imputation_run01_6664677232470394		failed	April 6, 2013, 19:04:43	TASKID:imputation_run01_6664677232470394

Error: terminate called after throwing an instance of 'std::bad_alloc'
what(): St9bad_alloc

TaskHistory
|<< << 1 - 10 of 16 >> >>|

Workflow deployment

Computational environments



“Harmonized” tool management

Tool in input sandbox
“getFile(‘tool.zip’)”

In \$WORKDIR
• Download

Tool deployed as
“load module”

In \$VO_BBMRI_NL_SW_DIR
• Download
• Build
• Configure

Simple download

vs.

On-site build deployment

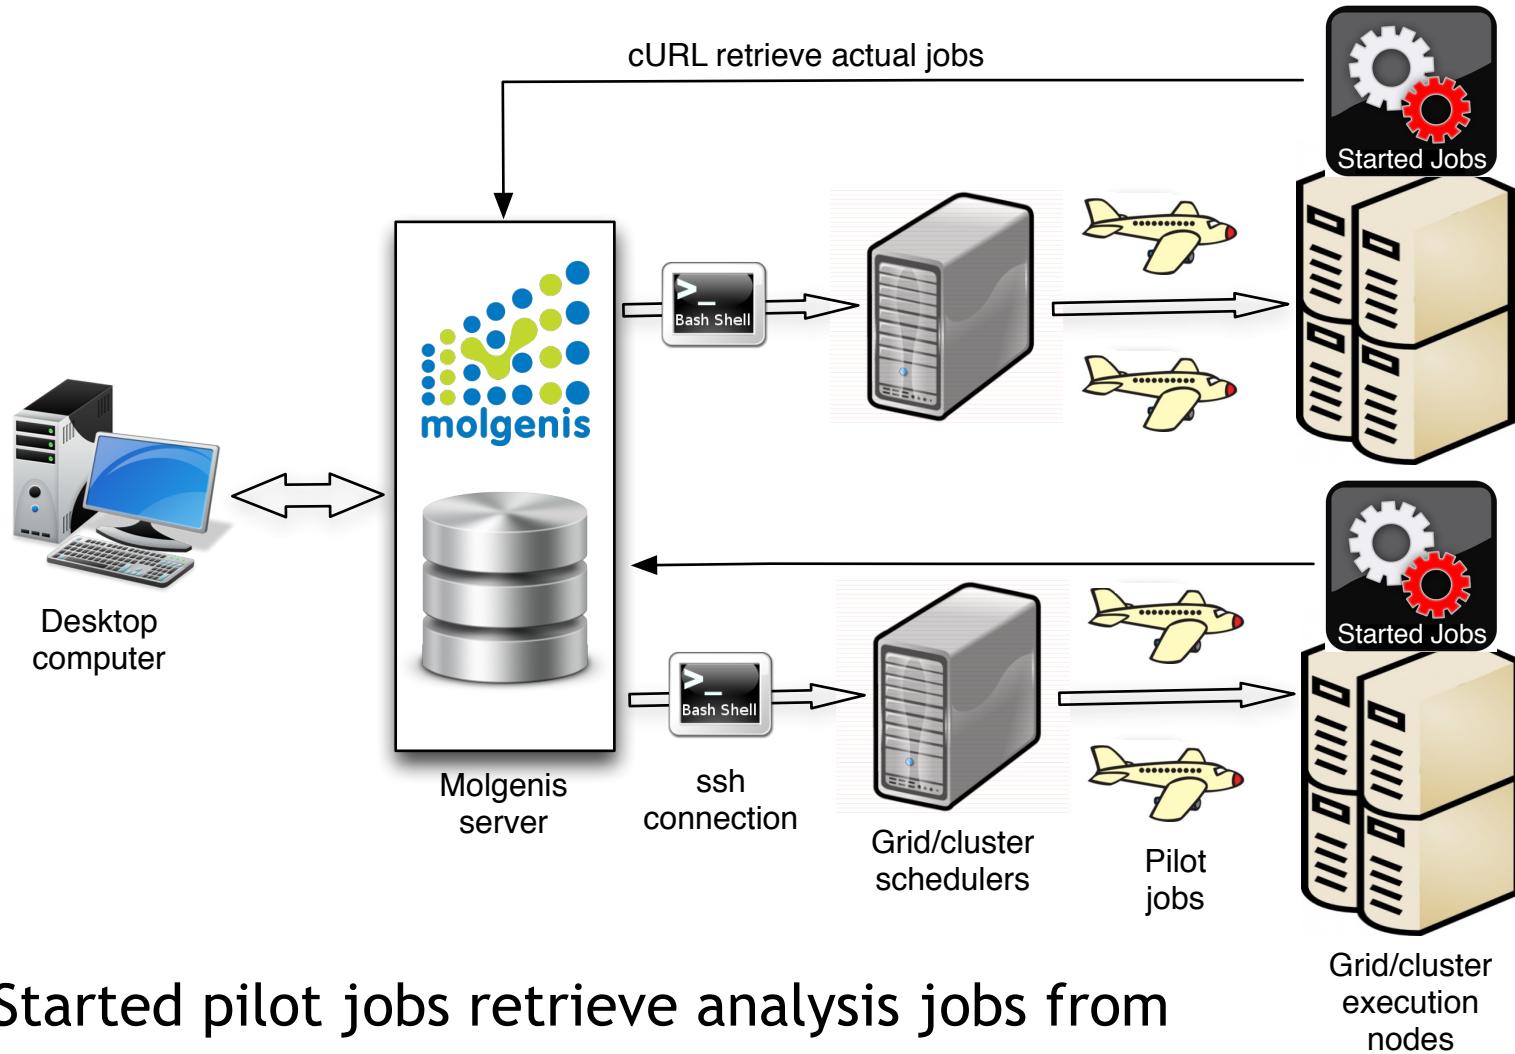
‘Harmonized’ tool management: modules

- Build using standard ‘modulecmd’ tool
- Software should be deployed at all grid sites
- Module file should be added to all sites
- <http://www.bbmriwiki.nl/svn/ebiogrid/modules/>

```
#  
%Module1.0#####  
##  
## bwa 0.5.8c_patched modulefile  
  
proc ModulesHelp { } {  
    global toolversion  
    global toolname  
  
        puts stderr "Set up the environment for $toolname version $toolversion  
\n"  
}  
  
# for Tcl script use only  
set toolname          bwa  
set toolversion        0.5.8c_patched  
set tooldir           "$env(VO_BBMRI_NL_SW_DIR)/tools/$toolname-$toolversion"
```

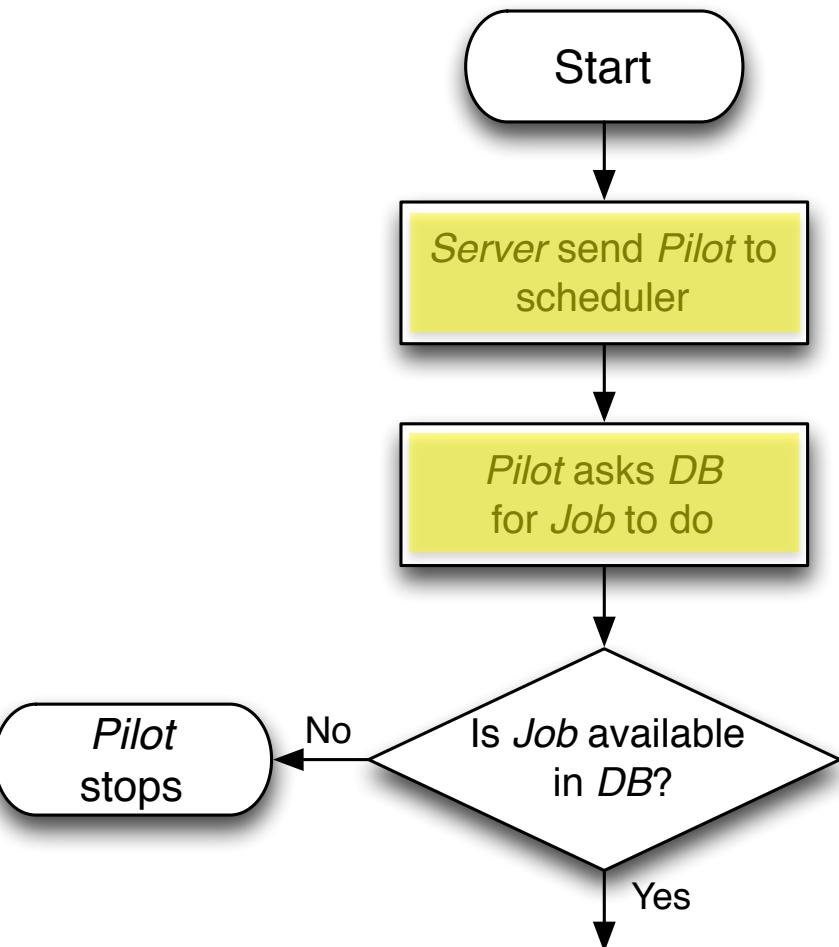
Workflow execution

Execution topology



- Started pilot jobs retrieve analysis jobs from Molgenis server

Workflow execution with pilots (1)



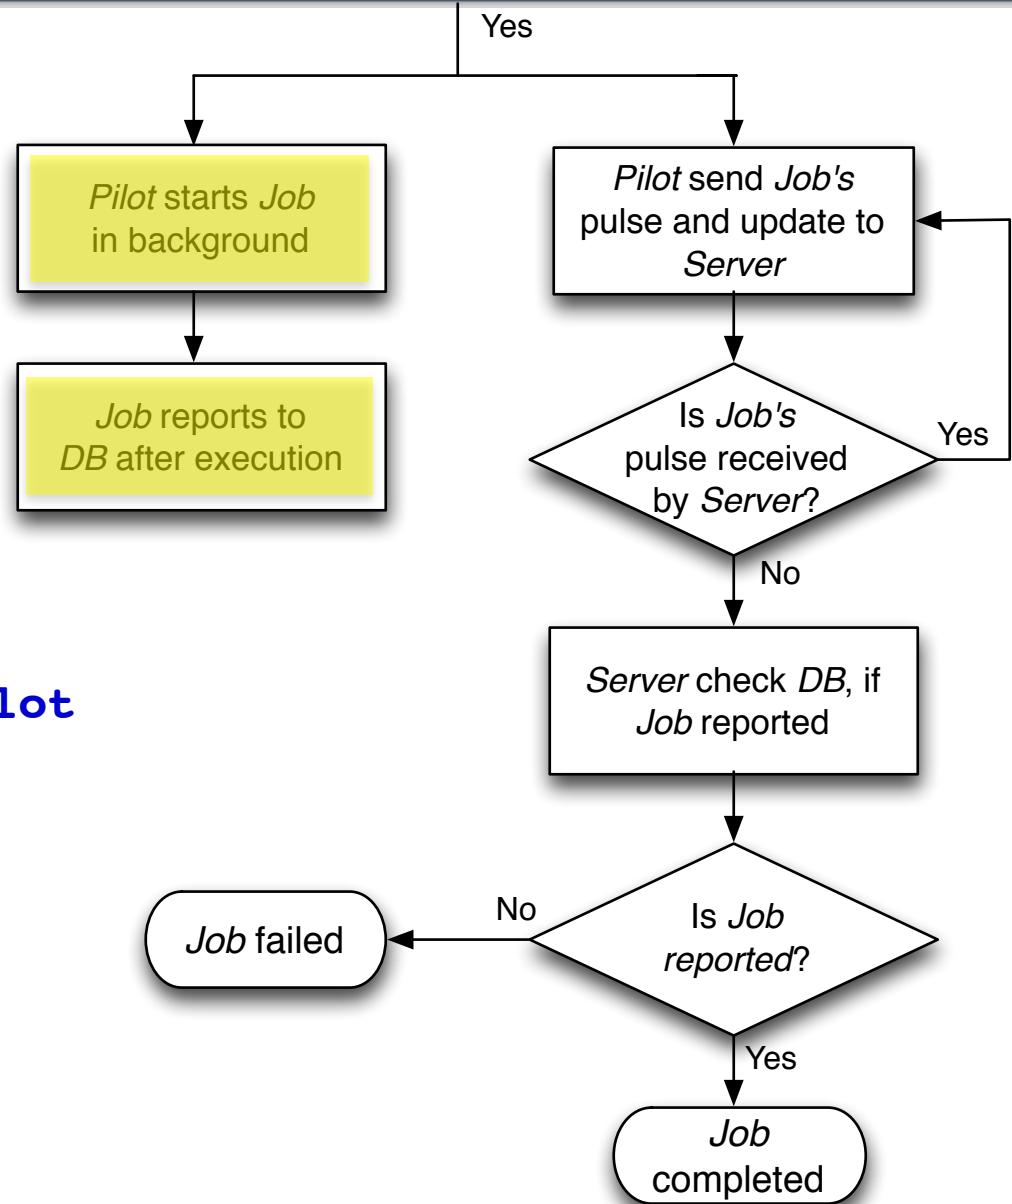
```
glite-wms-job-submit \
-d $USER ... $HOME/maverick.jdl
```

```
curl ... -F status=started \
-F backend=ui.grid.sara.nl \
http://$SERVER:8080/api/pilot > script
```

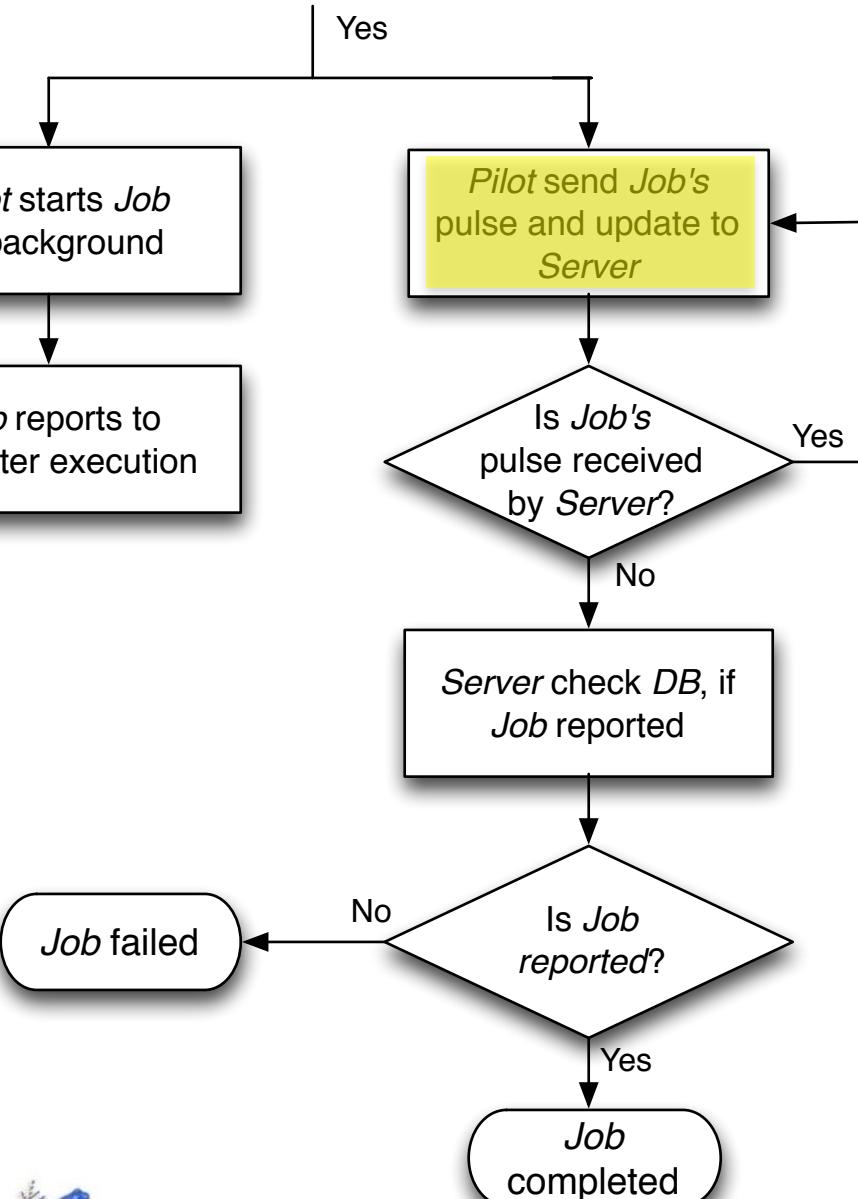
Workflow execution with pilots (2)

```
bash -l script.sh 2>&1 \
| tee -a log.log &
```

```
curl ... -F status=done \
-F log_file=@done.log \
http://$SERVER:8080/api/pilot
```



Workflow execution with pilots (3)



```

while [ 1 ] ; do
...
check_process "script.sh"
CHECK_RET=$?
if [ $CHECK_RET -eq 0 ] ;
then
...
curl ... -F status=nopulse \
-F log_file=@inter.log ...
...
elif
...
curl ... -F status=pulse \
-F log_file=@inter.log ...
...

```

Back-end independent analysis templates

Template structure

```
//header
    #MOLGENIS walltime=15:00 nodes=1 cores=4 mem=6

//tool management
    module load bwa/${bwaVersion}

//data management
    getFile ${indexfile}
    getFile ${leftbarcodefqgz}

//template of the actual analysis
    bwa aln \
        ${indexfile} ${leftbarcodefqgz} \
        -t ${bwaaligncores} -f ${leftbwaout}

//data management
    putFile ${leftbwaout}
```

Data transfer

- getFile and putFile
 - are back-end functions
 - now, we
 - check if file exists
 - do srm commands
- Input

```
getFile $1
```

```
#----- data transfer

getRemoteLocation()
{
  ARGS=($@)
  myFile=${ARGS[0]}
  remoteFile=srm://srm.grid.sara.nl/pnfs/grid.sara.nl/data/bbmri.nl/RP2${myFile}`expr length $TMPDIR`}
  echo $remoteFile
}

getFile()
{
  ARGS=($@)
  NUMBER="${#ARGS[@]}";
  if [ "$NUMBER" -eq "1" ]
  then

    myFile=${ARGS[0]}
    remoteFile=`getRemoteLocation $myFile`

    # 1. myPath = getPath( myFile ) will strip off the file name and return the path
    mkdir -p $(dirname "$myFile")

    # 2. cp srm:.../remoteFile myFile
    echo "srmcp -server_mode=passive $remoteFile file:///${myFile}"
    srmcp -server_mode=passive $remoteFile file:///${myFile}
    chmod 755 ${myFile}

  else
    echo "Example usage: getData \"\$TMPDIR/datadir/myfile.txt\""
  fi
}
```

- Generated output

```
getFile $1
imputation
chr20.map
```

```
putFile()
{
  ARGS=($@)
  NUMBER="${#ARGS[@]}";
  if [ "$NUMBER" -eq "1" ]
  then
    myFile=${ARGS[0]}
    remoteFile=`getRemoteLocation $myFile`
    echo "srmmr $remoteFile"
    srmmr $remoteFile
    echo "srmcop -server_mode=passive file:///${myFile} $remoteFile"
    srmcop -server_mode=passive file:///${myFile} $remoteFile
    returnCode=$?

    echo "srmcopy: ${returnCode}"

    if [ $returnCode -ne 0 ]
    then
      exit 1
    fi
  else
    echo "Example usage: putData \"\$TMPDIR/datadir/myfile.txt\""
  fi
}
```

Generated back-end independent script

```
//header
#MOLGENIS walltime=15:00 nodes=1 cores=4 mem=6
//tool management
module load bwa/0.5.8c_patched
//data management
getFile $WORKDIR/resources/hg19/indices/human_g1k_v37.fa
getFile $WORKDIR/groups/gcc/projects/cardio/run01/rawdata/
121128_SN163_0484_AC1D3HACXX_L8_CAACTT_1.fq.gz
//template of the actual analysis
bwa aln \
human_g1k_v37.fa 121128_SN163_0484_AC1D3HACXX_L8_CAACTT_1.fq.gz -t 4 \
-f 121128_SN163_0484_AC1D3HACXX_L8_CAACTT_1.bwa_align.human_g1k_v37.sai
//data management
putFile $WORKDIR/groups/gcc/projects/cardio/run01/results/
121128_SN163_0484_AC1D3HACXX_L8_CAACTT_1.bwa_align.human_g1k_v37.sai
```

Current developments

Pilots Dashboard

- During execution

run08@ui.grid.sara.nl

(May 31, 2013 11:31:31 AM)

Active

Inactivate

Stop Submitting Pilots

Jobs generated	0	Jobs done	4
Jobs ready	2	Jobs failed	0
Jobs running	0		
Pilots submitted	17	Pilots started	10

x

- Workflow is completed

run08@ui.grid.sara.nl

(May 31, 2013 11:31:31 AM)

Completed

Jobs generated	0	Jobs done	6
Jobs ready	0	Jobs failed	0
Jobs running	0		
Pilots submitted	21	Pilots started	12

x

Dash-board for jobs monitoring (work-in-progress)

Pilot dashboard						
run05@ui.grid.sara.nl	(May 30, 2013 3:43:51 PM)	Jobs generated	6	Jobs done	0	
Not active		Jobs ready	0	Jobs failed	0	
<input type="button" value="Activate"/>		Jobs running	0			
		Pilots submitted	0	Pilots started	0	
<hr/>						
run04@ui.grid.sara.nl	(May 30, 2013 3:43:33 PM)	Jobs generated	6	Jobs done	0	
Active		Jobs ready	0	Jobs failed	0	
<input type="button" value="Inactivate"/>		Jobs running	0			
<input type="text" value="Username"/>	<input type="text" value="Password"/>	<input type="button" value="Submit Pilots"/>	Pilots submitted	0	Pilots started	0
<hr/>						
run03@ui.grid.sara.nl	(May 30, 2013 3:43:18 PM)	Jobs generated	2	Jobs done	0	
Active		Jobs ready	4	Jobs failed	0	
<input type="button" value="Inactivate"/>		Jobs running	0			
<input type="button" value="Stop Submitting Pilots"/>		Pilots submitted	0	Pilots started	0	
<hr/>						
run02@ui.grid.sara.nl	(May 30, 2013 3:43:02 PM)	Jobs generated	5	Jobs done	0	
Not active		Jobs ready	0	Jobs failed	1	
<input type="button" value="Activate"/>		Jobs running	0			
<input type="button" value="Resubmit failed jobs"/>		Pilots submitted	0	Pilots started	0	
<hr/>						
run01@ui.grid.sara.nl	(May 30, 2013 3:42:45 PM)	Jobs generated	0	Jobs done	6	
Completed		Jobs ready	0	Jobs failed	0	
		Jobs running	0			
		Pilots submitted	0	Pilots started	0	

Enhancements and further steps

- What if not all parameters are known at the generation time
 - Run-time parameters passing to DB from previous steps (**implemented**)
- Advanced pilot management
 - Pilots re-using
- Better workflow visualization
 - Showing workflow elements and their properties
H. Byelas and M. Swertz, “Visualization of bioinformatics workflows for ease of understanding and design activities,” *Proceedings of the BIOSTEC BIOINFORMATICS-2013 conference*, pp. 117-123, 2013.

Conclusion

Conclusion

- One protocol template style that is suitable for different back-ends
- Workflow tools deployment using module system
- Hidden in scripts data management
- Workflow execution using pilot jobs

All available as open source

<http://www.molgenis.org>

<http://www.molgenis.org/wiki/ComputeStart>

h.v.byelas@gmail.com m.a.swertz@gmail.com

The screenshot shows a GitHub profile for the user 'molgenis'. The profile features a large blue starburst graphic containing the text 'Thank you! Questions?'. The GitHub interface includes a sidebar with the 'biosoftware platform' logo, location ('The Netherlands, Europe'), website link ('http://www.molgenis.org'), and join date ('Joined on Apr 28, 2012'). Below this, repository statistics are displayed: 4 public repos, 0 private repos, and 23 members. The main area lists four repositories:

- molgenis_apps**: MOLGENIS Advanced Application and Computation Framework. Last updated 14 hours ago.
- molgenis**: The MOLGENIS Software generator tool for creating Dynamic Software Infrastructure used in the Life Sciences. Last updated 17 hours ago.
- molgenis_test**: The MOLGENIS Test suite, fork this if you plan to contribute to MOLGENIS apps, so you can test commits before sending them in. Last updated 7 days ago.
- molgenis_distro**: The empty MOLGENIS Distribution, fork this and create your own MOLGENIS web application. Last updated a month ago.

On the right side of the GitHub interface, there is a vertical column of text fragments:

- analysing this data. This
- compute architecture by
- ing to 23 analysis steps.
- s converted to a binary

At the bottom of the slide, a note states: "• Realignment: in this part of the pipeline duplicate reads are marked using Picard. Afterwards realignment around known