
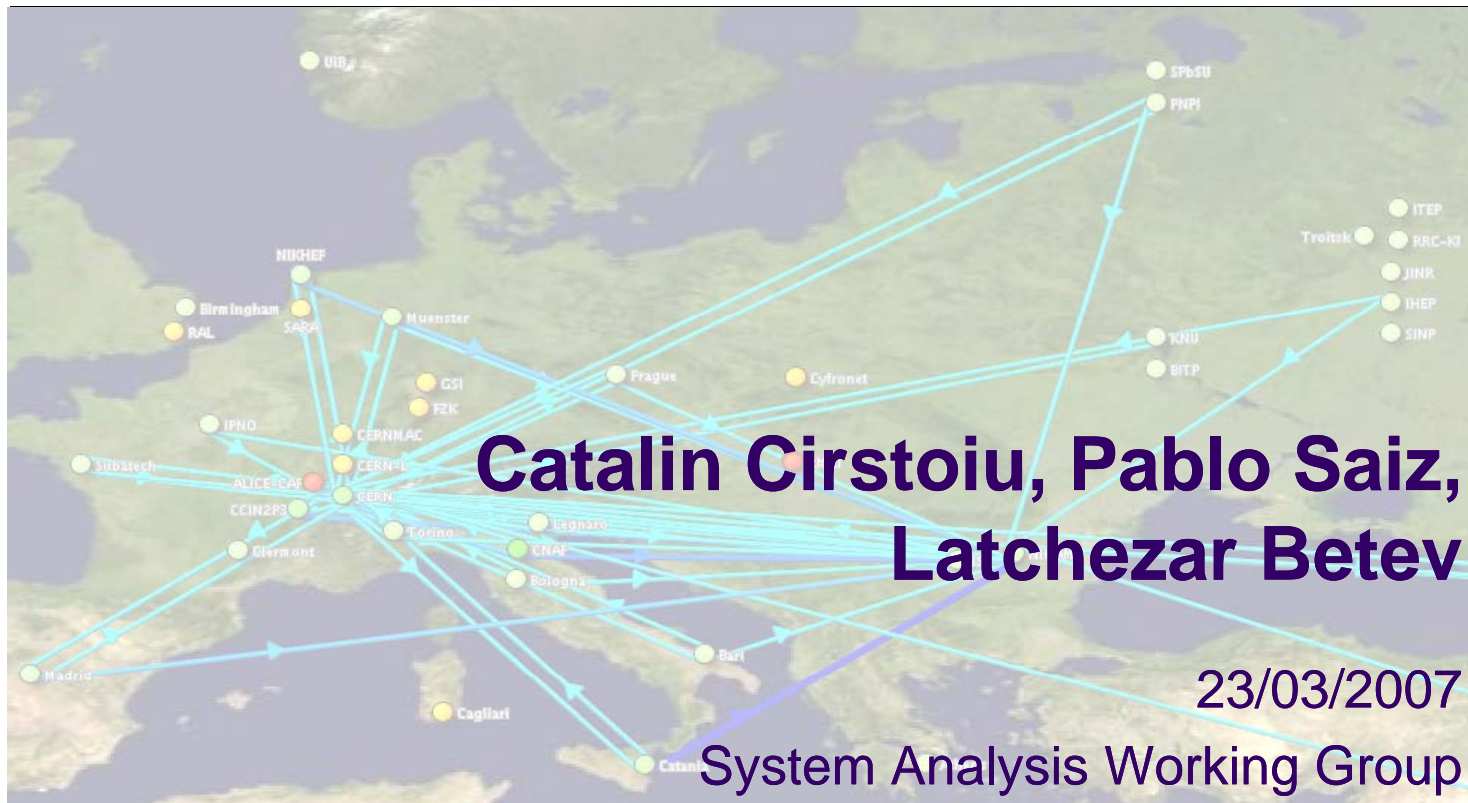


# Overview of ALICE monitoring



AliEn<sup>2</sup>  
@GRID





# Contents

- AliEn Overview
- Jobs workflow
- Direct Job Monitoring
- Monitoring requirements
- MonALISA overview
- What we monitor
- Monitoring architecture in AliEn
- Actions framework
- Feature examples



# AliEn overview

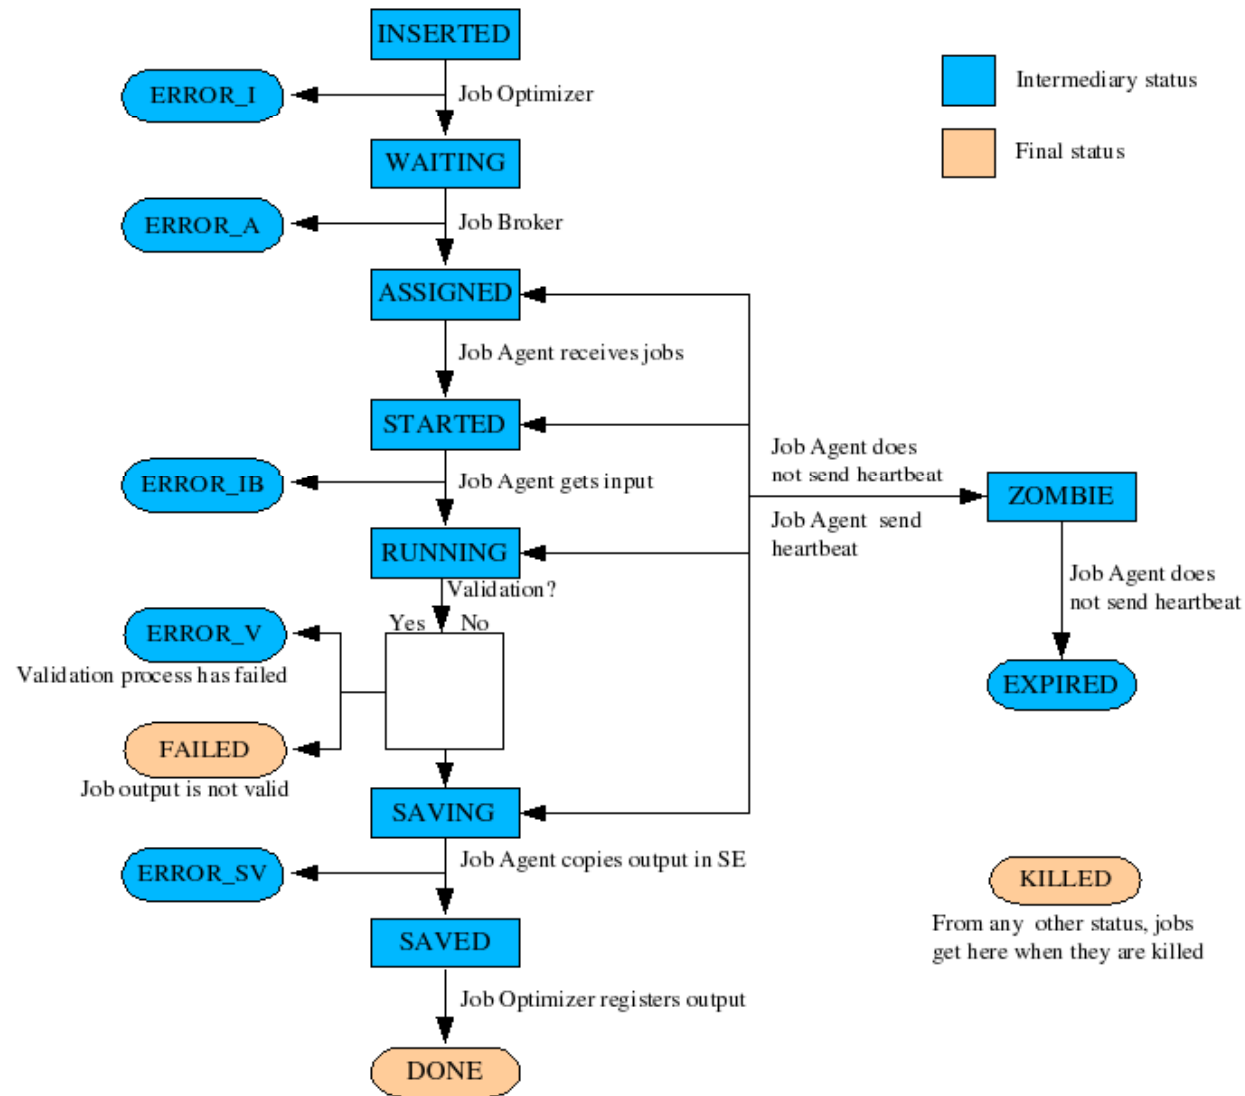
- Worldwide distributed system for running ALICE jobs
- Set of central services
  - TaskQueue, File Catalogue, Job Optimizer, Transfer Optimizer, User authentication, Logger Configuration...
- Set of site services – running on the VoBOX
  - Site proxy, Computing element, Storage Adapter, Package Manager;
  - Pilot jobs (Job Agents) submitted to the WNs
- It provides a high-level single interface to the Grid, separating the user from
  - Grid flavors
  - Complexity of the Grid itself
  - Platforms, operating system flavors, software environment



# AliEn Jobs workflow

- User submits job to AliEn
  - Job is registered in the AliEn TaskQueue
  - User receives back a job ID to track his job
- Various optimizers run on the queued jobs
  - Optimizer splits the job (if necessary) to run in several sub-jobs, usually close to the location of the data
  - Job priority and quotas
- Jobs are matched to the resources needed for their execution
- Job Agents
  - Picks a job, prepares input files, runs the job, saves the output
  - And reports on the status of the job through its lifetime

# Job status flow chart





AliEn<sup>2</sup>  
@GRID

# Direct job monitoring

- From the AliEn shell (second part of the talk – MonALISA)

- Show all my jobs

- Show details for job 2325337 (split job with 943 sub-jobs)

- Show details for 2477057 (another split job)

```
[aliendb06a.cern.ch:3307] /alice/cern.ch/user/a/aliprod/ > ps -x
aliprod    2325337  WS    -      0      /alice/bin/aliproot
aliprod    2462183  RS    -      0      /alice/bin/aliproot
aliprod    2467498  RS    -      0      /alice/bin/aliproot
aliprod    2468832  RS    -      0      /alice/bin/aliproot
aliprod    2470231  RS    -      0      /alice/bin/aliproot
aliprod    2471497  RS    -      0      /alice/bin/aliproot
aliprod    2472663  RS    -      0      /alice/bin/aliproot
aliprod    2474071  RS    -      0      /alice/bin/aliproot
aliprod    2475596  RS    -      0      /alice/bin/aliproot
aliprod    2477057  RS    -      0      /alice/bin/aliproot
aliprod    2481104  RS    -      0      /alice/bin/aliproot
aliprod    2488112  RS    -      0      /alice/bin/aliproot
[aliendb06a.cern.ch:3307] /alice/cern.ch/user/a/aliprod/ > jobinfo 2325337
=====
      ERROR_V :      13
      DONE    :     900
  SPLITTING  :        1
      ERROR_E :      30
=====
[aliendb06a.cern.ch:3307] /alice/cern.ch/user/a/aliprod/ > jobinfo 2477057
=====
      ERROR_IB :      26
      DONE    :     400
      SPLIT   :        1
      ERROR_SV :        1
      SAVED   :        1
      RUNNING :     355
      STARTED :      20
      EXPIRED :        5
      ERROR_V :     162
      ERROR_E :      14
      SAVING  :      14
      ZOMBIE  :        2
=====
[aliendb06a.cern.ch:3307] /alice/cern.ch/user/a/aliprod/ > █
```

# Direct job monitoring (2)

- Show all sub-jobs from master job 2325337 in a specific error condition

```
[aliendb06a.cern.ch:3307] /alice/cern.ch/user/a/aliproduct/ > top -split 2325337 -status ERROR_V
JobId  Status      Command name      Submitthost
2364804 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2364812 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2364813 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2364815 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2364817 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2364822 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2364826 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2365865 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2365866 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2365867 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2365869 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2365870 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
2365871 ERROR_V     /alice/bin/aliproduct      aliproduct@pcalimonitor.cern.ch
[aliendb06a.cern.ch:3307] /alice/cern.ch/user/a/aliproduct/ > █
```



# Direct job monitoring (3)

- Show full tracelog of one of the sub-jobs in ERROR\_V

```
[aliendb06a.cern.ch:3307] /alice/cern.ch/user/a/aliproduct/ > ps trace 2365866 all
001 Wed Mar 14 05:40:36 2007 [state      ]: Job 2365866 inserted from aliproduct@pcalimonitor.cern.ch [Master Job is 2325337]
002 Wed Mar 14 05:41:31 2007 [state      ]: Job state transition from INSERTING to WAITING
003 Wed Mar 14 06:18:35 2007 [state      ]: Job state transition to STARTED      |=|  procinfotime: 1173849515 site: ALICE::CNAF::LCG started
: 1173849515 node: wn-05-01-08-a.cr.cnaf.infn.it
004 Wed Mar 14 06:18:59 2007 [state      ]: Job state transition from STARTED    to RUNNING      |=|  procinfotime: 1173849539 site: ALICE::C
NAF::LCG started: 1173849539 spyurl: wn-05-01-08-a.cr.cnaf.infn.it:8085 node: wn-05-01-08-a.cr.cnaf.infn.it
005 Wed Mar 14 06:19:06 2007 [trace      ]: The job has been taken by the jobagent 6740_15115( https://egee-rb-01.cnaf.infn.it:9000/Euar9S
IW1e1j6QjNeANTjw )
006 Wed Mar 14 06:19:06 2007 [trace      ]: The job needs 72000 seconds
007 Wed Mar 14 06:19:06 2007 [trace      ]: Defining the environment variable ALIEN_JDL_OUTPUTDIR=/alice/sim/2006/pp_minbias/5139/427
008 Wed Mar 14 06:19:06 2007 [trace      ]: Defining the environment variable ALIEN_JDL_PACKAGES=VO_ALICE@APISCONFIG::V2.3##VO_ALICE@ROOT::
v5-14-00##VO_ALICE@GEANT3::v1-6-2##VO_ALICE@AliRoot::v4-04-Rev-10
009 Wed Mar 14 06:19:06 2007 [trace      ]: Creating the working directory /home/alicesgm/globus-tmp.wn-05-01-08-a.822.0/WMS_wn-05-01-08-a
01307_https_3a_2f_2fegee-rb-01.cnaf.infn.it_3a9000_2fEuar9SIW1e1j6QjNeANTjw/alien-job-2365866
010 Wed Mar 14 06:19:06 2007 [trace      ]: Request 4000 * 1 MB, found 102638 MB free!
011 Wed Mar 14 06:19:06 2007 [trace      ]: Downloading input file: /alice/bin/aliproduct
012 Wed Mar 14 06:19:06 2007 [trace      ]: Downloading input file: /alice/cern.ch/user/a/aliproduct/prod2006_2/configs_pp/validation.sh
013 Wed Mar 14 06:19:06 2007 [proc       ]: 00:00:00 0 0 0 1152 4448 4 15 3066.841 0.00 0 0
014 Wed Mar 14 06:19:06 2007 [trace      ]: Downloading input file: /alice/cern.ch/user/a/aliproduct/prod2006_2/configs_pp/CheckESD.C
015 Wed Mar 14 06:19:06 2007 [trace      ]: Downloading input file: /alice/cern.ch/user/a/aliproduct/prod2006_2/configs_pp/Config.C
016 Wed Mar 14 06:19:06 2007 [trace      ]: Downloading input file: /alice/cern.ch/user/a/aliproduct/prod2006_2/configs_pp/rec.C
017 Wed Mar 14 06:19:06 2007 [trace      ]: Downloading input file: /alice/cern.ch/user/a/aliproduct/prod2006_2/configs_pp/sim.C
018 Wed Mar 14 06:19:06 2007 [trace      ]: Downloading input file: /alice/cern.ch/user/a/aliproduct/prod2006_2/configs_pp/simrun.C
019 Wed Mar 14 06:19:06 2007 [trace      ]: Downloading input file: /alice/cern.ch/user/a/aliproduct/prod2006_2/configs_pp/tag.C
020 Wed Mar 14 06:19:06 2007 [trace      ]: Installing package VO_ALICE@APISCONFIG::V2.3
021 Wed Mar 14 06:19:06 2007 [trace      ]: Installing package VO_ALICE@ROOT::v5-14-00
022 Wed Mar 14 06:19:06 2007 [trace      ]: Installing package VO_ALICE@GEANT3::v1-6-2
023 Wed Mar 14 06:19:06 2007 [trace      ]: Installing package VO_ALICE@AliRoot::v4-04-Rev-10
024 Wed Mar 14 06:31:11 2007 [proc       ]: 00:11:38 698 28.67 10.7 405 450792 1265800 4 15 3066.841 1242.07 301171 1113865

042 Wed Mar 14 11:05:49 2007 [state      ]: Job state transition from RUNNING    to SAVING      |=|  procinfotime: 1173866749 site: ALICE::
NAF::LCG error:
043 Wed Mar 14 11:06:16 2007 [proc       ]: 00:00:00 0 87.73 0 0 1025624 1867792 4 15 3066.841 49198.26 533490 1345572
044 Wed Mar 14 11:06:16 2007 [trace      ]: Validating the output
045 Wed Mar 14 11:06:48 2007 [state      ]: Job state transition to ERROR_V      |=|  procinfotime: 1173866808 site: ALICE::CNAF::LCG spyurl
jdl:
046 Wed Mar 14 11:07:17 2007 [trace      ]: After the validation ERROR_V
047 Wed Mar 14 11:07:17 2007 [error      ]: The file AliESDs.root doesn't exist
048 Wed Mar 14 11:07:17 2007 [error      ]: The file AliESDfriends.root doesn't exist
049 Wed Mar 14 11:07:17 2007 [error      ]: The file check.root doesn't exist
050 Wed Mar 14 11:07:17 2007 [trace      ]: Using the guides from guid.txt
051 Wed Mar 14 11:07:17 2007 [trace      ]: Saving the files in the SE
052 Wed Mar 14 11:07:17 2007 [error      ]: The job didn't create raw.root
053 Wed Mar 14 11:07:17 2007 [trace      ]: Registering log_archive in Alice::CERN::CASTOR2
054 Wed Mar 14 11:07:17 2007 [trace      ]: Registering root_archive.zip in Alice::CERN::castor2
```





# Direct job monitoring (4)

- Show full tracelog of one of the sub-jobs in ERROR\_IB

```
001 Thu Mar 22 21:57:19 2007 [state   ]: Job 2491934 inserted from aliproduct@pcalimonitor.cern.ch [Master Job is 2481104]
002 Thu Mar 22 21:58:34 2007 [state   ]: Job state transition from INSERTING to WAITING
003 Thu Mar 22 23:24:41 2007 [state   ]: Job state transition to STARTED   |=|   procinfotime: 1174602281 site: ALICE::SARA::LCG started
: 1174602281 node: mu27.matrix.sara.nl
004 Thu Mar 22 23:24:49 2007 [trace   ]: The job has been taken by the jobagent 24653_2114( https://mu3.matrix.sara.nl:9000/J54yXmEfq8K
C-r00A97kUQ )
005 Thu Mar 22 23:24:49 2007 [trace   ]: The job needs 72000 seconds
006 Thu Mar 22 23:24:49 2007 [trace   ]: Defining the environment variable ALIEN_JDL_OUTPUTDIR=/alice/sim/2006/pp_minbias/5192/594
007 Thu Mar 22 23:24:49 2007 [trace   ]: Defining the environment variable ALIEN_JDL_PACKAGES=VO_ALICE@APISCONFIG::V2.3##VO_ALICE@ROOT::
v5-14-00##VO_ALICE@GEANT3::v1-6-2##VO_ALICE@AliRoot::v4-04-Rev-10
008 Thu Mar 22 23:24:49 2007 [trace   ]: Creating the working directory /scratch/259198.mu6.matrix.sara.nl/WMS_mu27_017072_https_3a_2f_2
fmu3.matrix.sara.nl_3a9000_2fJ54yXmEfq8KC-r00A97kUQ/alien-job-2491934
009 Thu Mar 22 23:24:49 2007 [trace   ]: Request 4000 * 1 MB, found 98113 MB free!
010 Thu Mar 22 23:30:53 2007 [trace   ]: Downloading input file: /alice/bin/aliroot
011 Thu Mar 22 23:30:53 2007 [trace   ]: Error downloading input file: /alice/bin/aliroot (trying again)
012 Thu Mar 22 23:30:53 2007 [trace   ]: Error downloading input file: /alice/bin/aliroot (trying again)
013 Thu Mar 22 23:30:53 2007 [error    ]: Could not download the input file: /alice/bin/aliroot (into /scratch/259198.mu6.matrix.sara.nl/
WMS_mu27_017072_https_3a_2f_2fmu3.matrix.sara.nl_3a9000_2fJ54yXmEfq8KC-r00A97kUQ/alien-job-2491934/command)
014 Thu Mar 22 23:39:45 2007 [state   ]: Job state transition to ERROR_IB   |=|   procinfotime: 1174603184 site: ALICE::SARA::LCG spyurl:
finished: 1174603184
```

# Direct job monitoring (5)

- All error conditions can be traced through this method
- Highly detailed log of the job progress
- Allows for further debugging of what went wrong
  - Example ERROR\_IB – SE not working, etc...
- In addition to already finished jobs – monitoring of currently running jobs

```
[aliendb06a.cern.ch:3307] /alice/cern.ch/user/a/aliprod/ > spy 2493328 workdir
total 136
drwxrwxrwx    2 alisgm02 alicesgm    4096 Mar 22 23:34 .
drwx-----   4 alisgm02 alicesgm    4096 Mar 22 23:34 ..
-rw-rw-r--    1 alisgm02 alicesgm   21974 Mar 22 23:34 CheckESD.C
-rwxr-xr-x    1 alisgm02 alicesgm     669 Mar 22 23:34 command
-rw-rw-r--    1 alisgm02 alicesgm   27838 Mar 22 23:34 Config.C
-rw-r--r--    1 alisgm02 alicesgm    3297 Mar 22 23:34 galice.root
-rw-rw-r--    1 alisgm02 alicesgm     535 Mar 22 23:34 rec.C
-rw-r--r--    1 alisgm02 alicesgm   25473 Mar 22 23:34 .rootrc
-rw-rw-r--    1 alisgm02 alicesgm     425 Mar 22 23:34 sim.C
-rw-rw-rw-    1 alisgm02 alicesgm   14839 Mar 22 23:50 sim.log
-rw-rw-r--    1 alisgm02 alicesgm    1835 Mar 22 23:34 simrun.C
-rw-rw-rw-    1 alisgm02 alicesgm         0 Mar 22 23:34 stderr
-rw-rw-rw-    1 alisgm02 alicesgm    2926 Mar 22 23:34 stdout
-rw-rw-r--    1 alisgm02 alicesgm    3158 Mar 22 23:34 tag.C
-rw-rw-r--    1 alisgm02 alicesgm    2129 Mar 22 23:34 validation.sh
[aliendb06a.cern.ch:3307] /alice/cern.ch/user/a/aliprod/ > █
```



# Direct job monitoring (6)

- Spy on any file in the sandbox as the job is running

```
[aliendb06a.cern.ch:3307] /alice/cern.ch/user/a/aliprod/ > spy 2493328 stdout
Test: ClusterMonitor is at erf.nikhef.nl:8084
Execution machine: wn-lui2-004.farm.nikhef.nl
*****
* APISCONFIG V2.3
* Setting up close SE ....
* Setting up API endpoints .....
* Setting up API PATH and LD_LIBRARY_PATH for shipped library.. *
* PATH => /data/esia/alice/slc3/packages/VO_ALICE/APISCONFIG/V2.3/a
pi/bin
* LD_LIBRARY_PATH => /data/esia/alice/slc3/packages/VO_ALICE/APISCONFIG/V2.3/a
pi/lib
* GCLIENT_NOGSI => 1
* GCLIENT_NOPROMPT => 1
* GCLIENT_COMMAND_MAXWAIT => 3600
* GCLIENT_COMMAND_RETRY => 50
* GCLIENT_SERVER_RESELECT => 4
* GCLIENT_SERVER_RECONNECT => 2
* GCLIENT_RETRY_DAMPING => 1.5
* GCLIENT_RETRY_SLEEPTIME => 2
*****
GCLIENT_SERVER_LIST => pcapiserv06.cern.ch:10000|pcapiserv07.cern.ch:10000|pcapi
serv05.cern.ch:10000|
*****
alien_CLOSE_SE => ALICE::NIKHEF::LCG
*****
Setting the environment for ROOT
Setting ROOTSYS to /data/esia/alice/slc3/packages/VO_ALICE/ROOT/v5-14-00/v5-14-00
Setting GEANT3
Setting AliROOT to /data/esia/alice/slc3/packages/VO_ALICE/AliRoot/v4-04-Rev-10
total used free shared buffers cached
Mem: 8199452 1798232 6401220 0 110904 1360976
-/+ buffers/cache: 326352 7873100
Swap: 3148700 52 3148648

HOME IS /tmp/332832.tbn20.nikhef.nl/WMS_wn-lui2-004_028915_https_3a_2f_2fbosheks.nikhef.n
l_3a9000_2fPQEFgyXyFfillw6BPE7OImg
alien-job-2493328
dg-submit.8893.sh
std.err
std.out
```



# Direct job monitoring (7)

- Additional information in the logfiles after the job has finished or ended in an error condition
- The sub-jobs with errors can be resubmitted for execution after the error has been identified in the same format as in the beginning (preserves even the job ID)
  - Simple command 'resubmit <job ID>
- ***Services status can also be monitored through the AliEn interface***
- For statistical overview of the job and services status and error conditions, ALICE uses MonALISA
  - Second part of the presentation




# Monitoring requirements

- Global view of the entire distributed system
  - Non-intrusive
  - Accurate
- Providing
  - Near real-time information
  - Long-term history of aggregated data
- On key parameters like
  - System status
  - Resource usage
- Helping with
  - Correlating events
  - System debugging
  - Generating reports
- Taking automated actions based on the monitored data

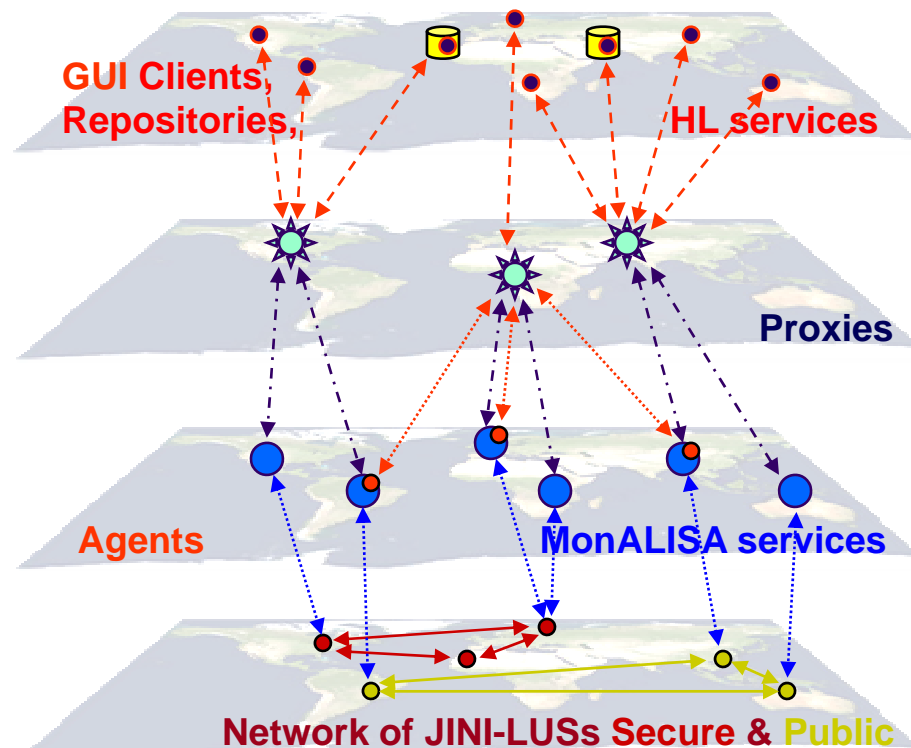
# MonALISA overview

- MonALISA is a Dynamic, Distributed Service Architecture capable to collect any type of information from different systems, to analyze it in near real time and to provide support for automated control decisions and global optimization of workflows in complex grid systems.



# ML Discovery System & Services

- The framework is based on a hierarchical structure of loosely coupled agents acting as distributed services which are independent & autonomous entities able to discover themselves and to cooperate using a dynamic set of proxies or self describing protocols.



**Global Services or Clients**

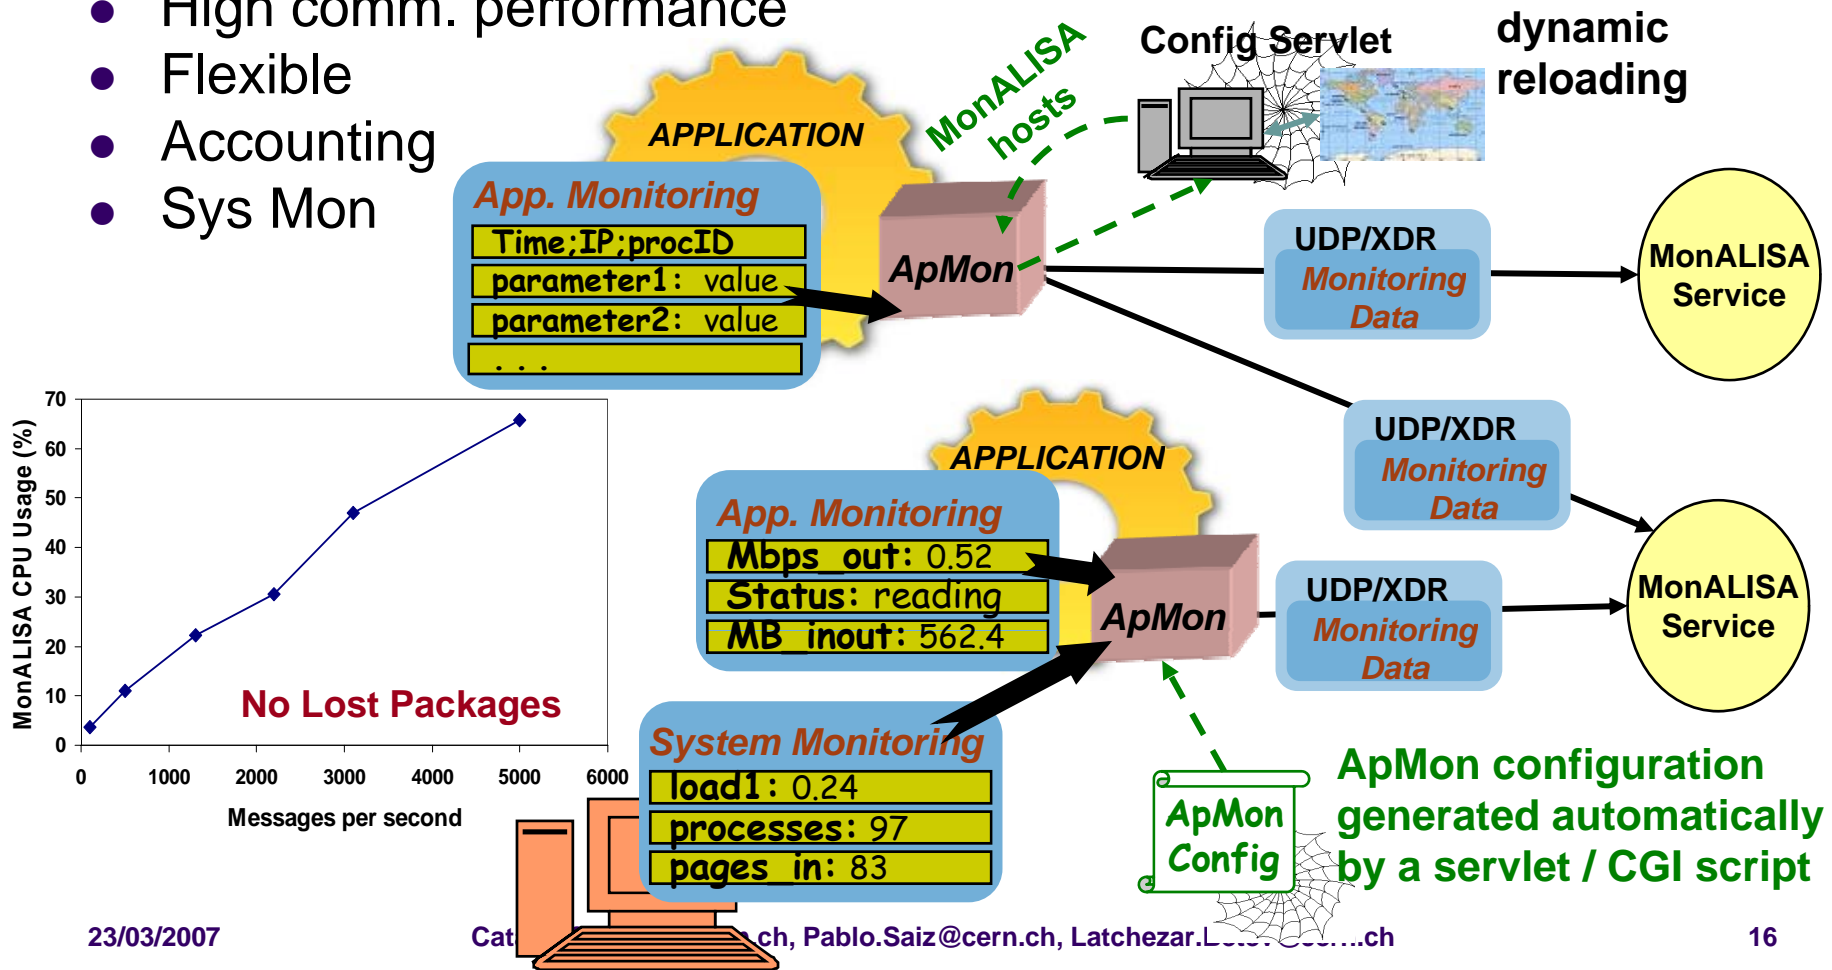
**Dynamic load balancing  
Scalability & Replication  
Security AAA for Clients**

**Distributed System  
for gathering and  
Analyzing Information**

**Distributed Dynamic  
Discovery-based on a lease  
Mechanism and REN**

# ApMon – Application Monitoring

- Lightweight library of APIs (C, C++, Java, Perl, Python) that can be used to send any information to MonALISA Services
- High comm. performance
- Flexible
- Accounting
- Sys Mon



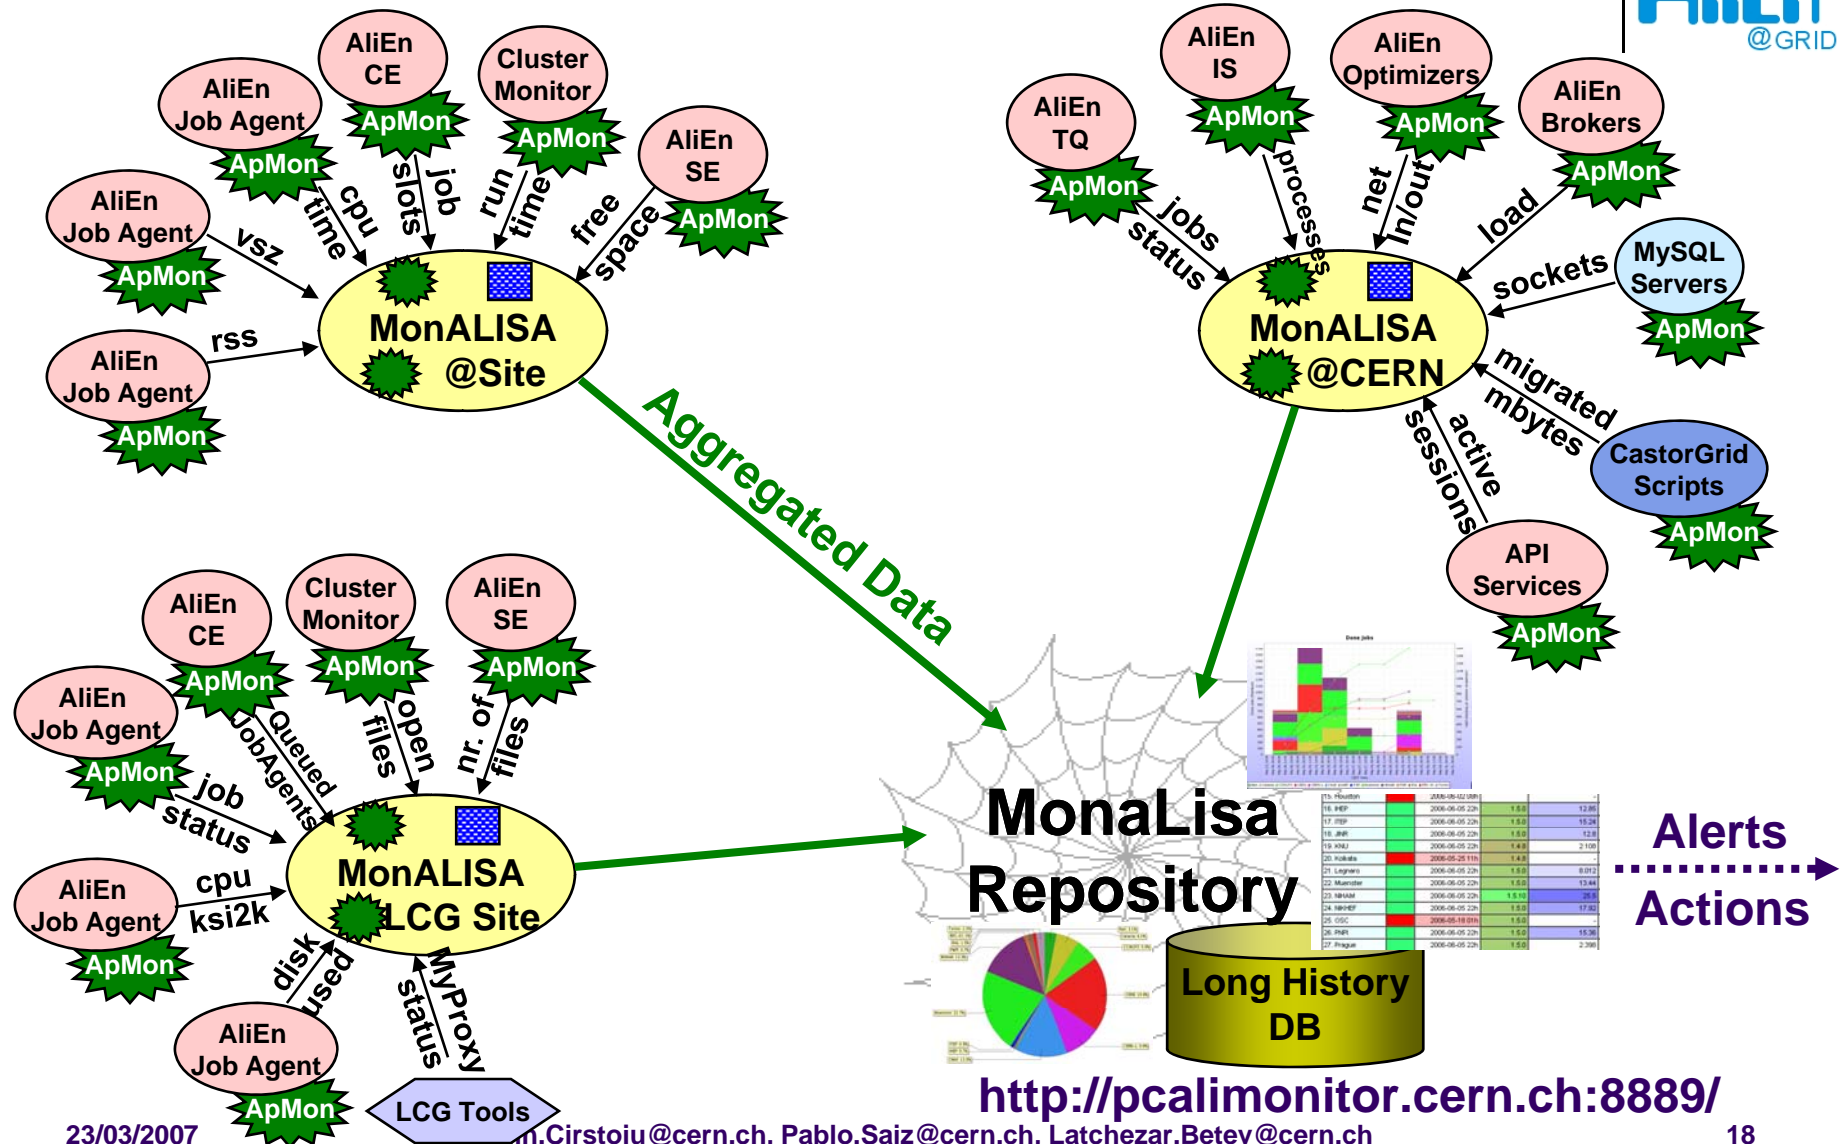




# What we monitor

- AliEn Components
  - Central Services
    - Task Queue, Information Service, Optimizers, API etc.
  - Site Services
    - Cluster Monitor, Computing & Storage Elements
    - Job Agents
  - Jobs status & resource usage
- Other Services
  - CastorGrid staging & migration, Xrootd, MySQL
- Nodes
  - Central, site, worker nodes
- Network traffic – inter & intra site
  - Via Xrootd
  - Via FTD

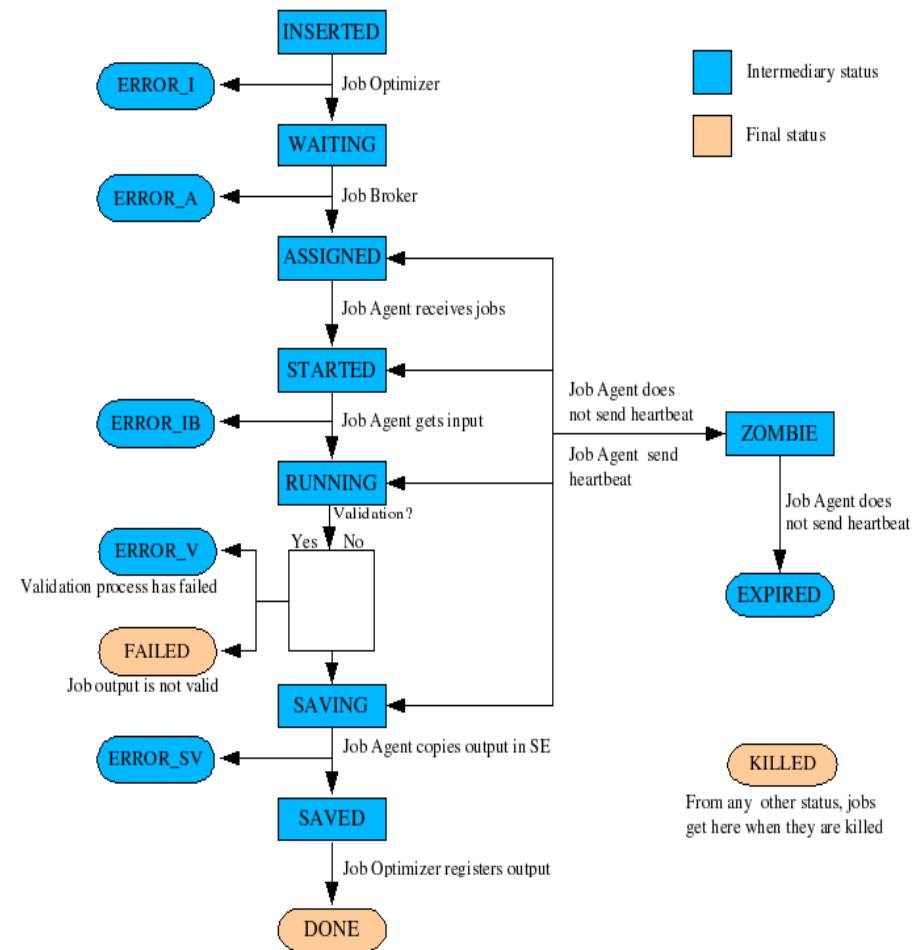
# Monitoring architecture in AliEn



# Job status monitoring



- Global summaries
  - For each/all conditions
  - For each/all sites
  - For each/all users
  - Running & cumulative
- Error status
- From job agents
- From central services
- Real-time map view
- Integrated pie charts
- History plots



# Job status & traffic - real-time map



MonALISA Repository  
**ALICE**

**MonALISA Client**  
Click on the button below to start the Monalisa Client.

[Client](#)

**ALICE Repository**

- Interactive Map
- Running trend
- Job Information
- SE Information
- Services
- Network Traffic
- FTD Transfers
- CAF Monitoring
- SHUTTLE
- Admin section

close all

**ALICE Reports**

Jobs

**Running jobs trend**

→ 24h
 → 12h
 → 6h
 → 1h

(Click arrows for detailed view)

Job Status
  Utilisation: Busy / Free Nodes

● Running Jobs
 ■ Zombie Jobs
 ■ No Active Jobs
 ■ ML Service Down

**SINP**

36 current running jobs

[1 hour](#)
[1 day](#)
[1 week](#)
[1 month](#)
[1 year](#)

Mar 2007

■ SINP

23/03/2007

Catalin.Cirstoiu@cern.ch, Pablo.Saiz@cern.ch, Latchezar.Betev@cern.ch

20

# Job status – integrated pie charts

Job status:  Assigned  Started  Running  Saving  Saved  Killed  Failed  Zombie  Error

Done jobs views:


Sites:  Athens  Bari  Birmingham  BITP  Bologna  Cagliari  Catania  CCIN2P3  CERN  CERN-L  CERNMAC  Clermont  CNAF  Cyfronet  FZK  GSI  Houston  IHEP  IPNO  ISS  ITEP  JINR  KNU  Kolkatta  LBL  Legnaro  Madrid  Muenster  NIHAM  NIKHEF  OSC  PNPI  Prague  RAL  RRC-KI  SARA  Sejong  SINP  SPbSU  Subatech  Torino  Troitsk  UIB  (check all | uncheck all)

Function:

Interval selection:  or  23:00 -  09:00 Image size:

Annotations:  Show annotation text:

Done jobs statistics





# Job status – history plots

Job status:

Running jobs views:


Sites: (check all | uncheck all)

Athens  Bari  Birmingham  BITP  Bologna  Cagliari  Catania  CCIN2P3  CERN  CERN-L  
 CERNMAC  Clermont  CNAF  Cyfronet  FZK  GSI  Houston  IHEP  IPNO  ISS  ITEP  
 JINR  KNU  Kolkatta  LBL  Legnaro  Madrid  Muenster  NIHAM  NIKHEF  OSC  PNPI  
 Prague  RAL  RRC-KI  SARA  Sejong  SINP  SPbSU  Subatech  Torino  Troitsk  UiB  
 SUM

Interval selection: - choose - or « 29-9-2006 15:00 . 30-9-2006 15:00 » Image size: 800x550

Annotations:  Show annotation text:

[Annotations](#) [What is this about?](#)


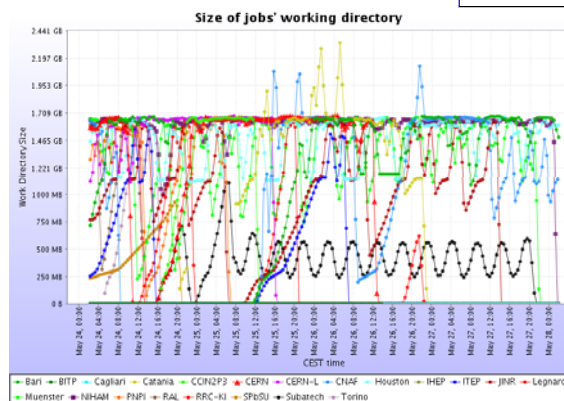




Running Jobs				
Farm	Last value	Min	Avg	Max
▲ Athens	0	0	4.746	15
▲ Bari	13	13	37.05	60
▲ Birmingham	0	0	0	0
▲ BITP	16	0	16.02	35
▲ Bologna	5	0	4.291	5
▲ Cagliari	22	13	21.13	23
▲ Catania	135	63	125.2	144
▲ CCIN2P3	180	84	164.1	234
▲ CERN	0	0	0.041	1
▲ CERN-L	413	238	390.8	591
▲ Clermont	50	20	42.56	50
▲ CNAF	176	106	224.7	340
▲ GSI	111	1	66.9	172
▲ Houston	7	0	35.44	94
▲ IHEP	0	0	0	0
▲ IPNO	25	13	23.29	25
▲ ISS	0	0	4.508	20
▲ ITEP	18	3	16.25	20
▲ JINR	7	0	4.365	8
▲ KNU	0	0	0	0
▲ Kolkatta	20	0	15.47	20
▲ LBL	248	0	133	262
▲ Legnaro	93	41	86.98	102
▲ Madrid	60	21	57.56	68
▲ Muenster	0	0	0	0
▲ NIHAM	93	39	87.2	95
▲ OSC	0	0	6.076	34
▲ PNPI	37	9	33.99	40
▲ Prague	48	0	37.88	67
▲ RAL	0	0	0	0
▲ RRC-KI	27	1	21.67	31
▲ Sejong	2	1	6.375	10
▲ SINP	27	12	30.19	45
▲ SPbSU	3	1	2.995	4
▲ Subatech	15	7	13.8	15
▲ Torino	22	17	37.53	53
▲ Troitsk	6	5	7.72	10
▲ UiB	19	4	17.04	24
<b>Total</b>	<b>1898</b>	<b>1776</b>		



# Job resource usage monitoring


- Cumulative parameters
  - CPU Time & CPU KSI2K
  - Wall time & Wall KSI2K
  - Read & written files
  - Input & output traffic (xrootd)
- Running parameters
  - Resident memory
  - Virtual memory
  - Open files
  - Workdir size
  - Disk usage
  - CPU usage
- Aggregated per site



# Job network traffic monitoring



- Based on the xrootd transfer from every job
- Aggregated statistics for
  - Sites (incoming, outgoing, site to site, internal)
  - Storage Elements (incoming, outgoing)
- Of what
  - Read and written files
  - Transferred MB/s



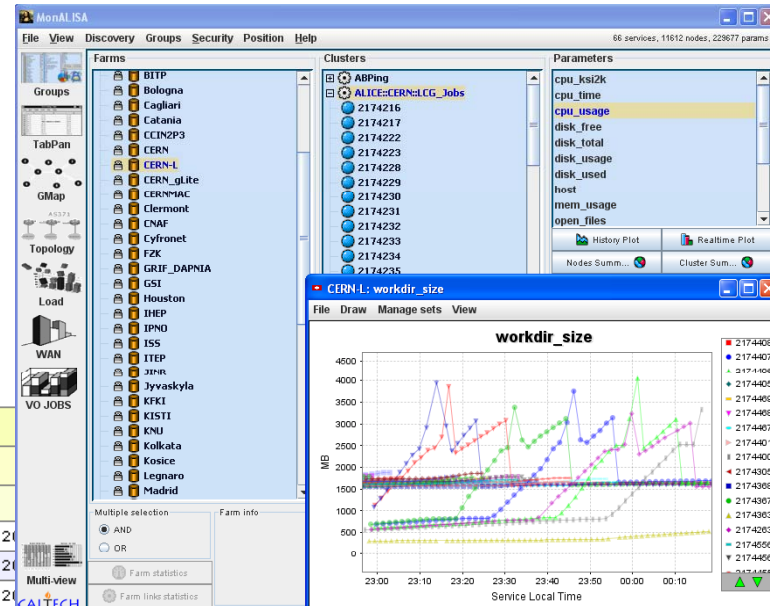
Incoming Traffic (Writing to SEs)						Outgoing Traffic (Reading from SEs)					
Farm	Last value	Min	Avg	Max	Total	Farm	Last value	Min	Avg	Max	Total
Alice::CERN::Castor2	61.5 MB/s	2.758 MB/s	46.5 MB/s	105.6 MB/s	3.812 TB	Alice::CERN::Castor2	0 B/s	0 B/s	0 B/s	0 B/s	0 B
Alice::CERN::scratch	0 B/s	0 B/s	0 B/s	0 B/s	0 B	Alice::CERN::scratch	0 B/s	0 B/s	0 B/s	0 B/s	0 B
Alice::CERN::se01	0 B/s	0 B/s	0 B/s	0 B/s	0 B	Alice::CERN::se01	0 B/s	0 B/s	89.27 KB/s	18.61 KB/s	7.319 MB
Alice::CERN::se03	0 D/s	0 D/s	0 D/s	0 D/s	0 D	Alice::CERN::se03	0 D/s	0 D/s	3.029 KB/s	14.16 KB/s	279.5 MD
Alice::ISS::File	0 B/s	0 B/s	0 B/s	0 B/s	0 B	Alice::NIHAM::File	0 B/s	0 B/s	0 B/s	0 B/s	0 B
Alice::NIHAM::File	0 B/s	0 B/s	6.121 KB/s	15.61 KB/s	513.9 MB						
<b>Total</b>	<b>61.51 MB/s</b>		<b>46.51 MB/s</b>			<b>Total</b>	<b>3.295 KB/s</b>		<b>3.416 KB/s</b>		<b>286.8 MB</b>



# Individual job tracking

- Based on AliEn shell cmds.
  - top, ps, spy, jobinfo, masterjob
  - Interaction with file catalogue
- Using the GUI ML Client
  - Status, resource usage, per job

Job parameters				Application software			
Run#	PID	Owner	Events	ROOT	ALIROOT	GEANT	Date
		aliproduct (1023)					last year
5188	2472663	aliproduct		v5-14-00	v4-04-Rev-10	v1-6-2	21.03.2007 23:07 /alice/sim/2
5187	2471497	aliproduct		v5-14-00	v4-04-Rev-10	v1-6-2	21.03.2007 21:07 /alice/sim/2
5186	2470231	aliproduct	8500	v5-14-00	v4-04-Rev-10	v1-6-2	21.03.2007 19:08 /alice/sim/2

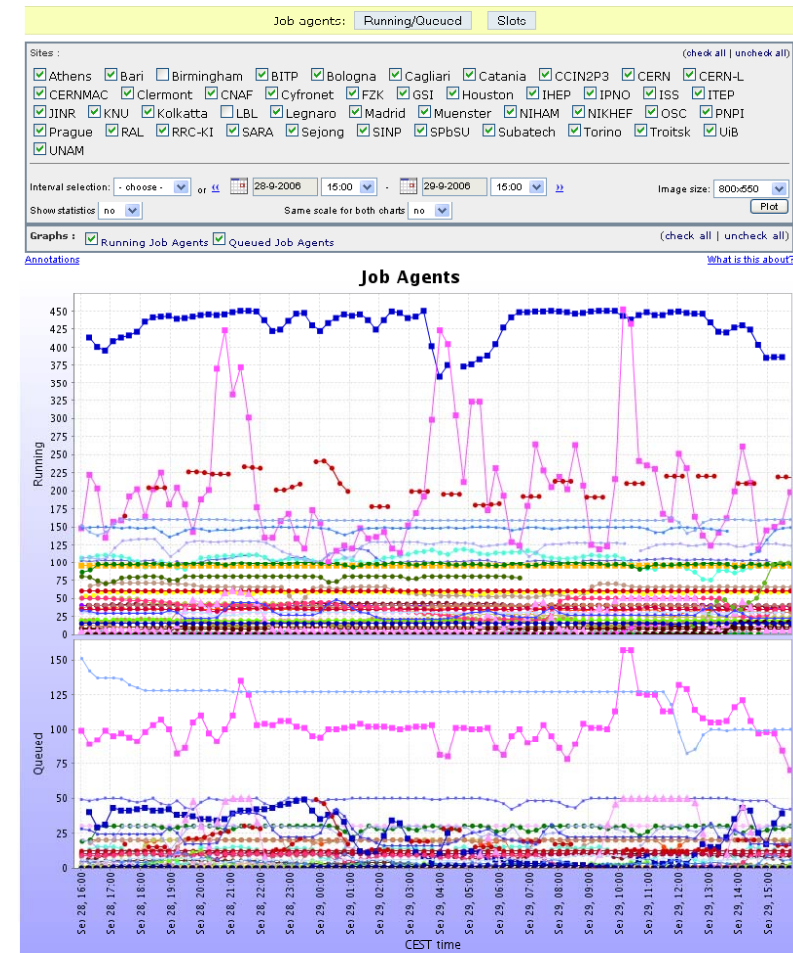


pid	owner	first seen	last seen	subjobs	Job states														
					SPLIT	WAITING	STARTED	RUNNING	SAVING	DONE	ERRORS	ERROR_V	ERROR_SW						
2472663	aliproduct	21.03.2007 23:07	22.03.2007 00:07	1000	1	98%	986	0%	4	0%	9								
2471497	aliproduct	21.03.2007 21:07	22.03.2007 00:07	1000	1			0%	1	91%	918				8%	81	1%	15	
2470231	aliproduct	21.03.2007 19:08	22.03.2007 00:07	1000	1					90%	903	0%	2	8%	85	0%	8	0%	2
2468832	aliproduct	21.03.2007 16:06	22.03.2007 00:07	1000	1			0%	1	58%	587	3%	39	32%	328	3%	35	0%	4
2467498	aliproduct	21.03.2007 13:06	22.03.2007 00:07	1000	1					38%	380	0%	9	59%	592	1%	18	0%	2
2463641	aliproduct	21.03.2007 10:06	22.03.2007 00:07	1000	1					15%	150	0%	4	83%	833	1%	13	0%	1
2462183	aliproduct	21.03.2007 06:06	22.03.2007 00:07	1000	1					13%	132	0%	5	85%	854	0%	9		
2460887	aliproduct	21.03.2007 02:06	22.03.2007 00:07	1000	1					6%	62	0%	1	92%	929	0%	8		
2459472	aliproduct	21.03.2007 00:06	22.03.2007 00:07	1000	1					4%	44			94%	944	1%	12	0%	2

# Job agents monitoring



- From Job Agent itself
  - Requesting job
  - Installing packages
  - Running job
  - Done
  - Error statuses
- From Computing Element
  - Available job slots
  - Queued Job Agents
  - Running Job Agents



# AliEn & LCG Services monitoring



- AliEn services
  - Periodically checked
  - PID check + SOAP call
  - Simple functional tests
  - SE space usage
  - Efficiency

AliEn services												
Service	VOBox		Monitoring script	AliEn services					AliEn tests			
	Address	AliEn version		CE	SE	PackMan	Monitor	FTD	add	get	rm	whereis
1. aliendb1.cern.ch	aliendb1.cern.ch			-	-		-	-	-	-	-	-
2. Athens	xg010.inp.demokritos.gr	v2-11_102						-	Oct ...			
3. Bari	alicegrid6.ba.infn.it	Failed wit...			DEAD...			-	Fail...	Fail...	Fail...	Fail...
4. Birmingham	epbf006.ph.bham.ac.uk	Failed wit...						-	Fail...	Fail...	Fail...	Fail...
5. BITP	alice9.bitp.kiev.ua	v2-11_102						-	Oct ...			

- LCG environment and tools
  - Proxy, gsiscp, LCG CE/SE, Job submission, BDII, Local catalog, sw. area etc.
  - Error messages in case of failure
  - Efficiency


LCG services																			
Service	Site			Proxy						LCG CE					BDII			Local catalog	Software area
	Name	Included in LDAP	Machine's proxy	Renewal	Server	Connection	Query	User proxy registration	VOBOX registration	gsiscp	Publication of SE	LCG CE	Job submission	LFC def. in LDAP	Local BDII	Conn. to local BDII			
1. Athens					-				-		-			-	Loca...	-	LFC_...		
2. Bari					-		ERRO...		Fail...			Fail...		-			LFC_...		
3. Birmingham					-									-			LFC_...		
4. Bologna					-		ERR					Fail...		-			LFC_...		
5. Cagliari					-		ERR					Fail...		-			LFC_...		
6. Catania					-		ERR					Fail...		-			LFC_...		
7. COMPASS					-									-			LFC_...		

ERROR: Couldn't find a valid proxy. Use -debug for further information. DN not specified and unable to get DN from user proxy; at /opt/lcg/bin/vobox-proxy line 48.

# CastorGrid scripts monitoring




- Migration
  - Amount, speed, errors
- Staging
  - Amount, speed, errors
- Nodes
  - Host parameters
- Xrootd resource usage
- File cache status
  - Used space, no. of files



# API Services monitoring

- API Service sessions
  - Established, active
- API Service users
  - Active, total
- Statistics
  - Executed commands



Established sessions				
Farm	Last value	Min	Avg	Max
Server 1	0000	0000	0000	0000
Server 2	9999	9999	9999	9999
Server 3	9999	9999	9999	9999
<b>Total</b>	<b>29997</b>	<b>29997</b>		

Active sessions				
Farm	Last value	Min	Avg	Max
Server 1	1	0	2.999	126
Server 2	59	44	289.8	771
Server 3	49	31	237.7	600
<b>Total</b>	<b>109</b>		<b>530.4</b>	

Valid Sessions				
Farm	Last value	Min	Avg	Max
Server 1	606	379	1369	2033
Server 2	6909	6858	7773	8358
Server 3	5458	5133	5474	6017
<b>Total</b>	<b>12872</b>		<b>14615</b>	

# VOBox monitoring

- Machine parameters, real-time & history
  - Load, memory & swap usage, processes, sockets

Global views:

Detailed history:




## VO Box machine status

Machine status																			
Site name	Last see online	CPU					Mem [% MB]		Swap [% MB]		Eth0 [KB/s]		Eth1 [KB/s]		Eth2 [KB/s]		System		
		Load5	Usage	User	System	Cnt	MHz	Usage	Total	Usage	Total	In	Out	In	Out	In	Out	Procs	Socks
1. Athens	2006-10-17 08h	0.288	14.6	11.82	2.781	1	2680	29	2010	0	1992	19.53	19.65	-	-	-	-	85	62
2. Bari	2006-10-17 08h	0.15	6.444	4.187	2.257	2	1800	55	1001	14.04	2047	0.213	0.07	16.13	8.643	-	-	69	44
3. Birmingham	2006-10-17 08h	0.359	17.59	14.2	3.389	2	800	53	1001	5.489	2047	15.9	7.91	-	-	-	-	92	49
4. BITP	2006-10-17 08h	0.09	1.819	1.018	0.801	4	2793	65	1001	8.805	2000	2.565	2.551	54.57	51.41	-	-	90	107
5. Bologna	2006-10-17 08h	0.061	1.75	1.16	0.59	4	3067	16	4005	0	2000	-	-	16.15	7.39	-	-	94	41
6. Cagliari	2006-10-17 08h	0.059	1.818	1.36	0.458	2	3199	16	2007	0	2000	8.594	7.159	-	-	-	-	63	32
7. Catania	2006-10-17 08h	0.075	1.912	1.271	0.642	4	2799	35	2006	0.206	4094	14.53	8.779	-	-	-	-	86	43
8. CCIN2P3	2006-10-17 07h	0.078	2.687	1.709	0.978	4	3000	38	2007	30.51	2047	1.278	0.079	10.63	5.16	-	-	119	45
9. CERN	2006-10-17 08h	1.113	23.13	18.78	4.354	4	2387	59	5768	1.316	8000	598.6	78	-	-	-	-	666	1340
10. CERN-L	2006-10-17 08h	0.317	15.87	12.17	3.7	2	2793	24	3991	0	4094	24.79	9.156	-	-	-	-	121	56
11. CERNMAC	2006-10-17 08h	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
12. Clermont	2006-10-17 08h	0.287	14.16	11.3	2.857	1	2007	23	3013	0	8189	30.96	23.89	-	-	-	-	111	92
13. CNAF	2006-10-17 07h	0.106	4.89	3.333	1.558	2	3067	45	4005	0.185	4000	17.17	10.37	-	-	-	-	147	94
14. Cytronet	2006-10-17 08h	0.066	3.249	2.963	0.285	2	1300	27	1982	0.578	1000	-	-	20.31	13.19	-	-	77	36
15. FZK	2006-10-17 07h	0.068	1.947	1.564	0.383	4	3000	56	2007	10.68	3827	10.79	4.595	0.204	0.186	-	-	167	115
16. GSI	2006-10-10 04h	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
17. Houston	2006-10-17 08h	0.244	9.331	7.436	1.895	1	1396	43	4014	2.529	4095	0.016	0.039	17.06	19.59	-	-	121	147
18. IJCL	2006-10-17 08h	1.388	2.555	0.973	1.003	4	2704	24	2006	4.255	4000	468.8	16.14	5.554	24.26	-	-	207	107

# FTD Monitoring



- Status of the transfers
- Transfer rates
- Success/failures



# Annotations



**MonALISA Repository**  
**ALICE**

**MonALISA Client**  
Click on the button below to start the Monalisa Client.

**Client**

**ALICE Repository**

- Interactive Map
- Job Information
  - Site views
    - Summary plots
    - Job states
    - Jobs per site
    - Jobs per site
    - Job agents
    - Resource use
  - User views
    - Task queue
  - SE Information
  - Services
  - Network Traffic
  - AliEn Tests
  - FTD Transfers
  - CAF Monitoring
  - Original

close all

**ALICE Reports**

[Site Administration](#)

[Farm Colour Configuration](#)

[Bookmark this page \(url\)](#)

Job status:

Running jobs views:

Sites:  Athens  Bari  Birmingham  BITP  Bologna  Cagliari  Catania  CCIN2P3  CERN  CERN-L  CERNMAC  Clermont  CNAF  Cyfronet  FZK  GSI  Houston  IHEP  IPNO  ISS  ITEP  JINR  KNU  Kolkatta  LBL  Legnaro  Madrid  Muenster  NIHAM  NIKHEF  OSC  PNPI  Prague  RAL  RRC-KI  SARA  Sejong  SINP  SPbSU  Subatech  Torino  Troitsk  UIB

SUM

Interval selection: last month or 14-9-2006 22:00 - 15-10-2006 08:00 Image size: 800x550

Annotations:  enabled Show annotation text:  yes  [What is this about?](#)

**Running Jobs**

Running jobs [Number]

CEST time

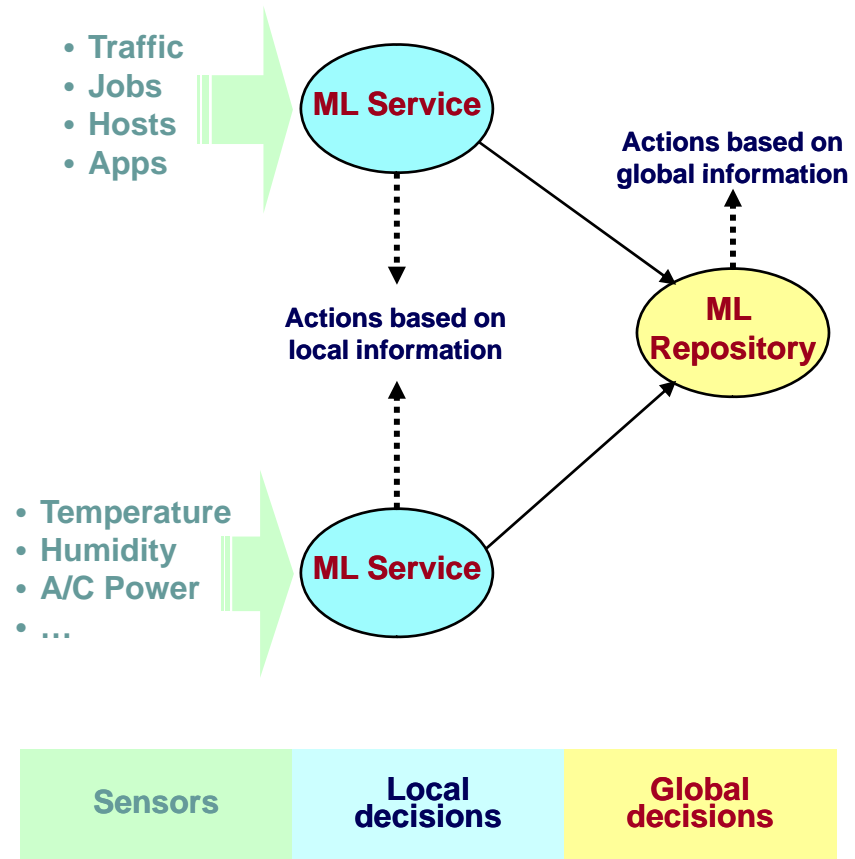
SUM  Athens  Bari  Birmingham  BITP  Bologna  Cagliari  Catania  CCIN2P3  CERN  CERN-L  CERNMAC  Clermont  CNAF  Cyfronet  FZK  GSI  Houston  IHEP  IPNO  ISS  ITEP  JINR  KNU  Kolkatta  LBL  Legnaro  Madrid  Muenster  NIHAM  NIKHEF  OSC  PNPI  Prague  RAL  RRC-KI  SARA  Sejong  SINP  SPbSU  Subatech  Torino  Troitsk  UIB

Running Jobs				
Farm	Last value	Min	Avg	Max
Athens	5	0	6.249	18
Bari	0	0	19.1	60
Birmingham	1	0	3.925	20
BITP	0	0	12.86	65
Bologna	0	0	4.432	9
Cagliari	0	0	17.61	27
Catania	1	0	74.62	157
CCIN2P3	3	0	78.54	249
CERN	0	0	8.323	201
CERN-L	0	0	171.4	598
CERNMAC	1	0	1.5	2
Clermont	0	0	34.02	55
CNAF	0	0	177.7	451
Cyfronet	0	0	0	0
FZK	0	0	71.87	850
GSI	(27.09.2006 14:00, 27.09.2006 22:00)	46	278	
Ho	Network intervention	86	99	
IHE	(27.09.2006 22:00, 04.10.2006 11:55)	31	64	
IPN	Site-wide Proxy problem	04	28	
ISS		3.4	116	
ITEP		0	11.78	22
JINR		2	6.07	12
KNU		1	2.913	50
Kolkatta		0	14.94	35
LBL		0	31.94	918
Legnaro		0	52.34	106
Madrid		0	50.21	76
Muenster		0	41.05	93
NIHAM		2	79.69	182
NIKHEF		1	53.19	191
OSC		0	28.32	104
PNPI		0	23.53	40
Prague		0	31.12	100
RAL		0	4.301	49
RRC-KI		0	20.84	40
SARA		0	1.723	20
Sejong		1	0.886	25
SINP		4	19.23	68
SPbSU		1	0.367	6




# Actions framework

- Based on monitoring information, actions can be taken in
  - ML Service
  - ML Repository
- Actions can be triggered by
  - Values above/below given thresholds
  - Absence/presence of values
  - Correlation between multiple values
- Possible actions types
  - Alerts
    - e-mail
    - Instant messaging
    - RSS Feeds
  - External commands
  - Event logging



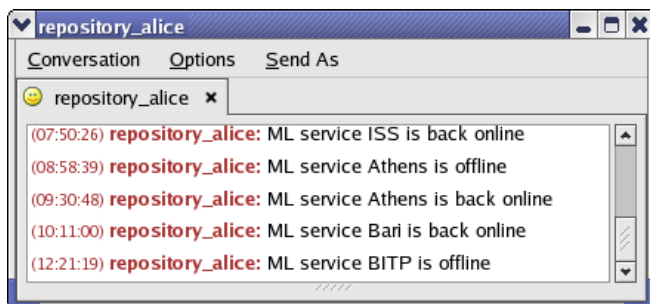
# Alerts and actions



**MySQL daemon is automatically restarted when it runs out of memory**  
**Trigger: threshold on VSZ memory usage**



**ALICE Production jobs queue is automatically kept full by the automatic resubmission**  
**Trigger: threshold on the number of *aliproduct* waiting jobs**



**Administrators are kept up-to-date on the services' status**  
**Trigger: presence/absence of monitored information**



# Summary

- All aspects of the system are monitored
  - Job execution, job data transfers
  - Central and site services
  - Machines
- LCG services monitoring is done through custom scripts
  - This should be improved
- MonALISA is a very flexible tool
  - Provides a top-down monitoring solution for large wide distributed systems
  - It allows using the gathered data for intelligent control and decisions taking

# Questions?



# Thank you!



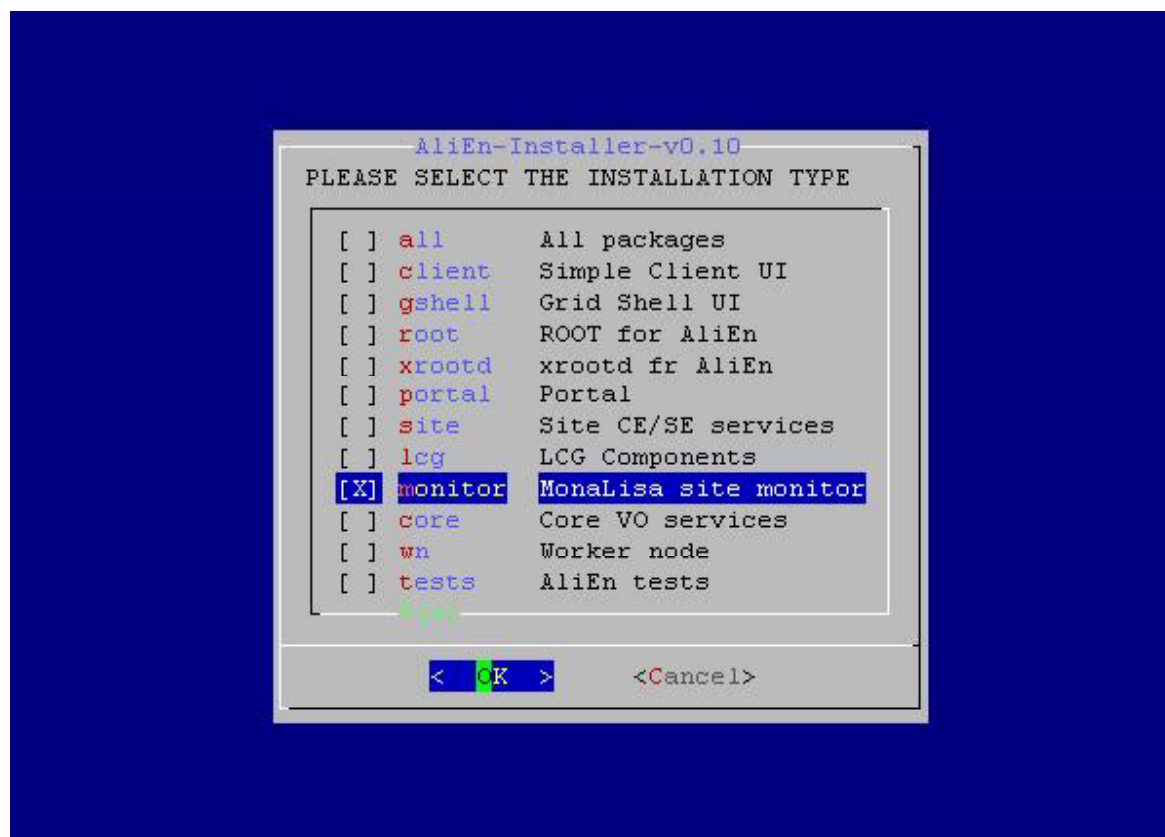
<http://alien.cern.ch>



<http://monalisa.caltech.edu>

# ML Service deployment

- MonALISA is packaged and prepared for installation by the AliEn Build System (BITS)






# ML Service configuration

- From site administrators point of view, it is just like any other AliEn service
  - You start it with ``alien StartMonaLisa``
  - You stop it with ``alien StopMonaLisa``
  - Check status with ``alien StatusMonaLisa``
- Configuration files for ML are generated automatically from AliEn LDAP

# PROOF CAF Monitoring



- Each host reports
  - CPU, memory, swap, network
- Each slave reports
  - Summaries per query type
  - CPU, memory
  - Event rate
  - File rate
  - I/O vs. network rate



Traffic between the cluster machines (MB/sec) (last 0.5h average)

Machine	6047	6048	6049	6050	6052	6053	6054	6055	6056	6057	6058	6059	6060	6061	6062	6063	6064	6065	6066	6067	6068	6069	607
1. 6047	0	-	-	-	-	-	-	2.927	2.018	-	-	1.094	-	-	-	1.908	4.112	-	-	0.974	0.614	0	0
2. 6048	-	9.406	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
3. 6049	-	-	8.678	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
4. 6050	-	-	-	6.692	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
5. 6052	-	-	-	-	3.913	-	1.454	-	-	-	-	3.084	-	0.317	0	0	0	0	-	-	0.985	4.447	-
6. 6053	-	-	-	-	-	6.803	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
7. 6054	0	-	-	1.363	-	6.195	-	-	-	-	0	-	-	-	0	-	-	-	-	-	0	-	1.58
8. 6055	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
9. 6056	-	-	-	-	-	-	-	4.962	-	2.442	0.525	-	-	-	-	-	-	-	-	-	-	-	-
10. 6057	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
11. 6058	1.164	-	-	-	0	-	-	2.531	-	0	0	-	-	-	-	1.103	-	0	-	-	-	-	-
12. 6059	3.755	-	0.622	-	-	-	-	-	-	-	-	11.76	1.955	0	0.677	1.848	0	-	-	-	2.812	-	0.78
13. 6060	-	-	-	-	-	-	-	2.068	-	-	-	-	11.59	-	-	1.06	-	-	-	-	-	-	2.08
14. 6061	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
15. 6062	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
16. 6063	-	-	-	-	1.655	0.27	2.416	-	-	-	-	-	-	0	-	6.38	-	0	-	0	-	-	-
17. 6064	-	-	-	-	-	1.123	-	2.822	-	-	-	-	1.621	-	0	-	3.117	-	0	0	-	-	0.58
18. 6065	0	-	-	-	3.52	3.165	-	0	-	0	-	-	-	-	-	3.034	0	1.579	-	0	-	-	-
19. 6066	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-



# Summary

- Monitoring is vital in a large distributed system
  - Distributed systems should have built-in monitoring capabilities
- We have to deal with a lot of data
  - We need decentralization
  - And aggregation
- MonALISA is very flexible and powerful
  - Monitoring is just the beginning
  - Gathered data can be used for intelligent control and decisions taking