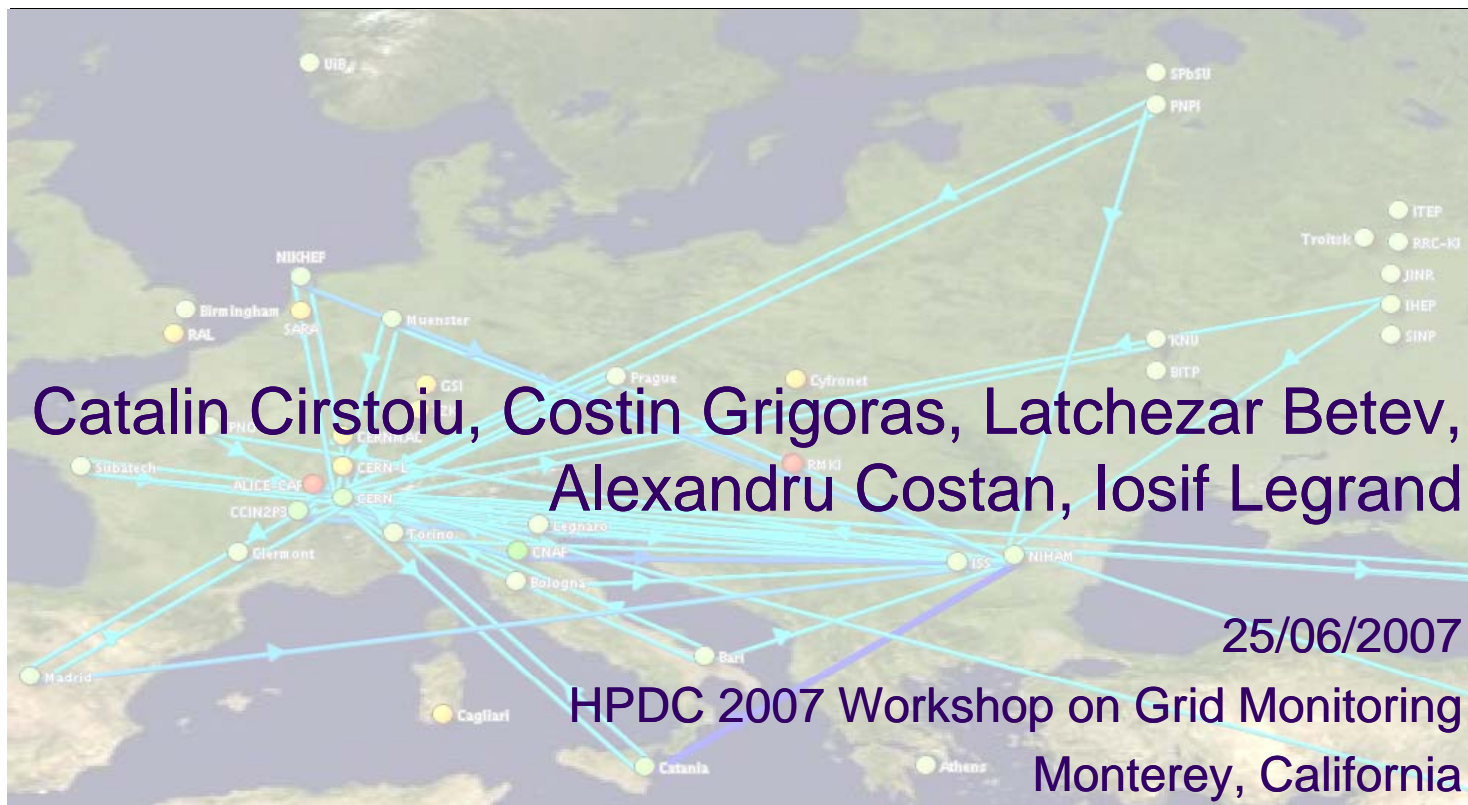


Monitoring, Accounting and Automated Decision Support for the ALICE Experiment Based on the MonALISA Framework



<http://monalisa.caltech.edu>



AliEn²
@GRID

<http://alien.cern.ch>





Contents

- Monitoring requirements
- MonALISA overview
- Application monitoring
- Monitoring architecture in AliEn
 - Jobs monitoring
 - Traffic monitoring
 - Services monitoring
 - Nodes monitoring
- Actions framework
- Feature snapshots

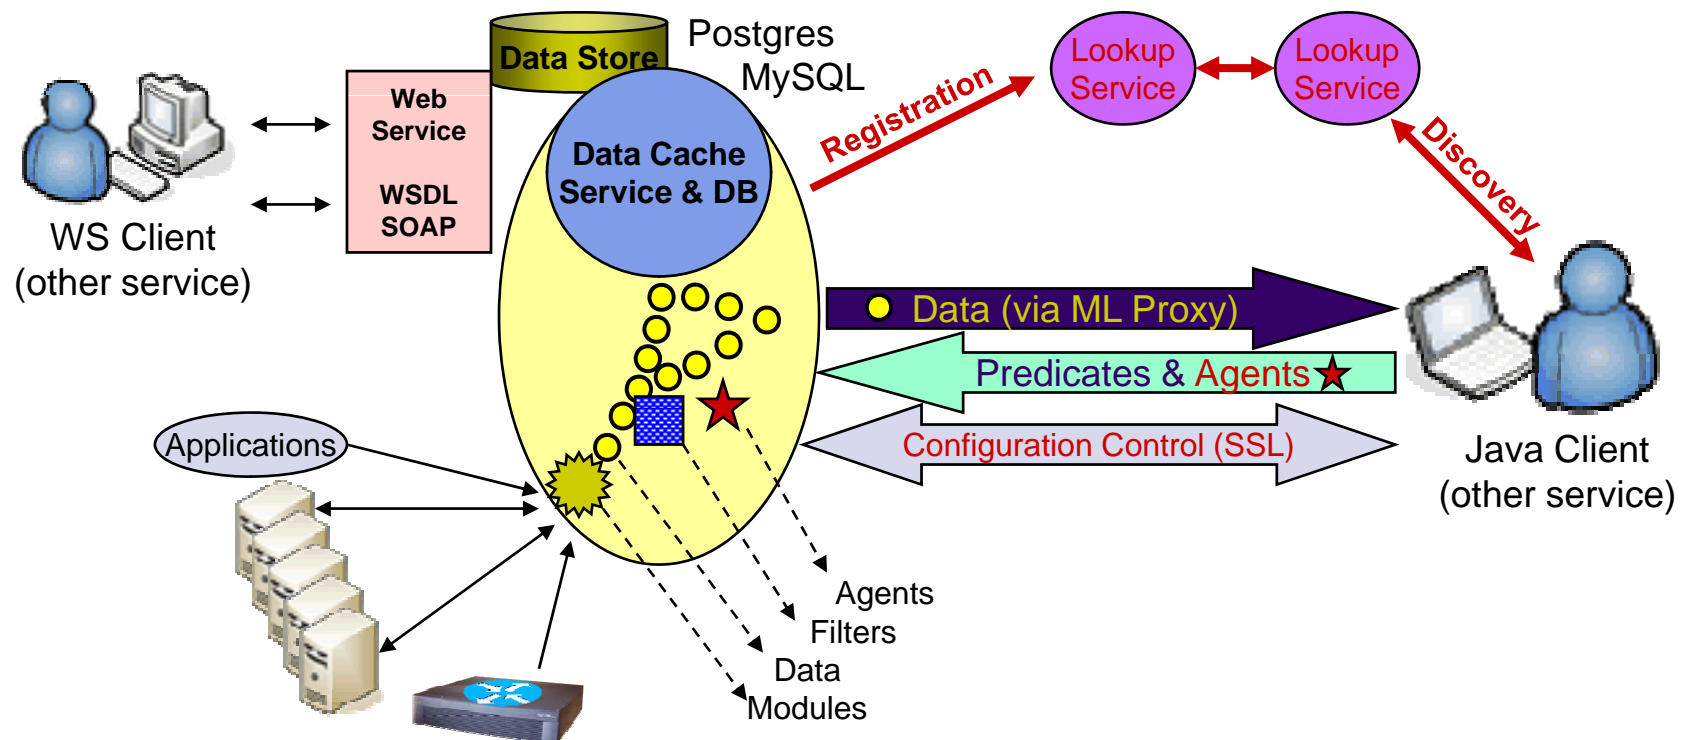


Monitoring Requirements

- Global view of the entire distributed system
 - Least-intrusive
 - As accurate as possible
 - Best-effort data transport
 - Minimizing the requirements for open ports
- 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 framework
- Collects any type of information from different systems
- Aggregates and analyzes it in near-real time
- Provides support for automated control decisions and global optimization of workflows in complex distributed systems.



ML Discovery System & Services

- Hierarchical structure of loosely coupled services
- Independent & autonomous entities able to
 - Publish their existence
 - Discover other available Jini-enabled services
 - Use a dynamic set of proxies to cooperate with them



Global Services or Clients

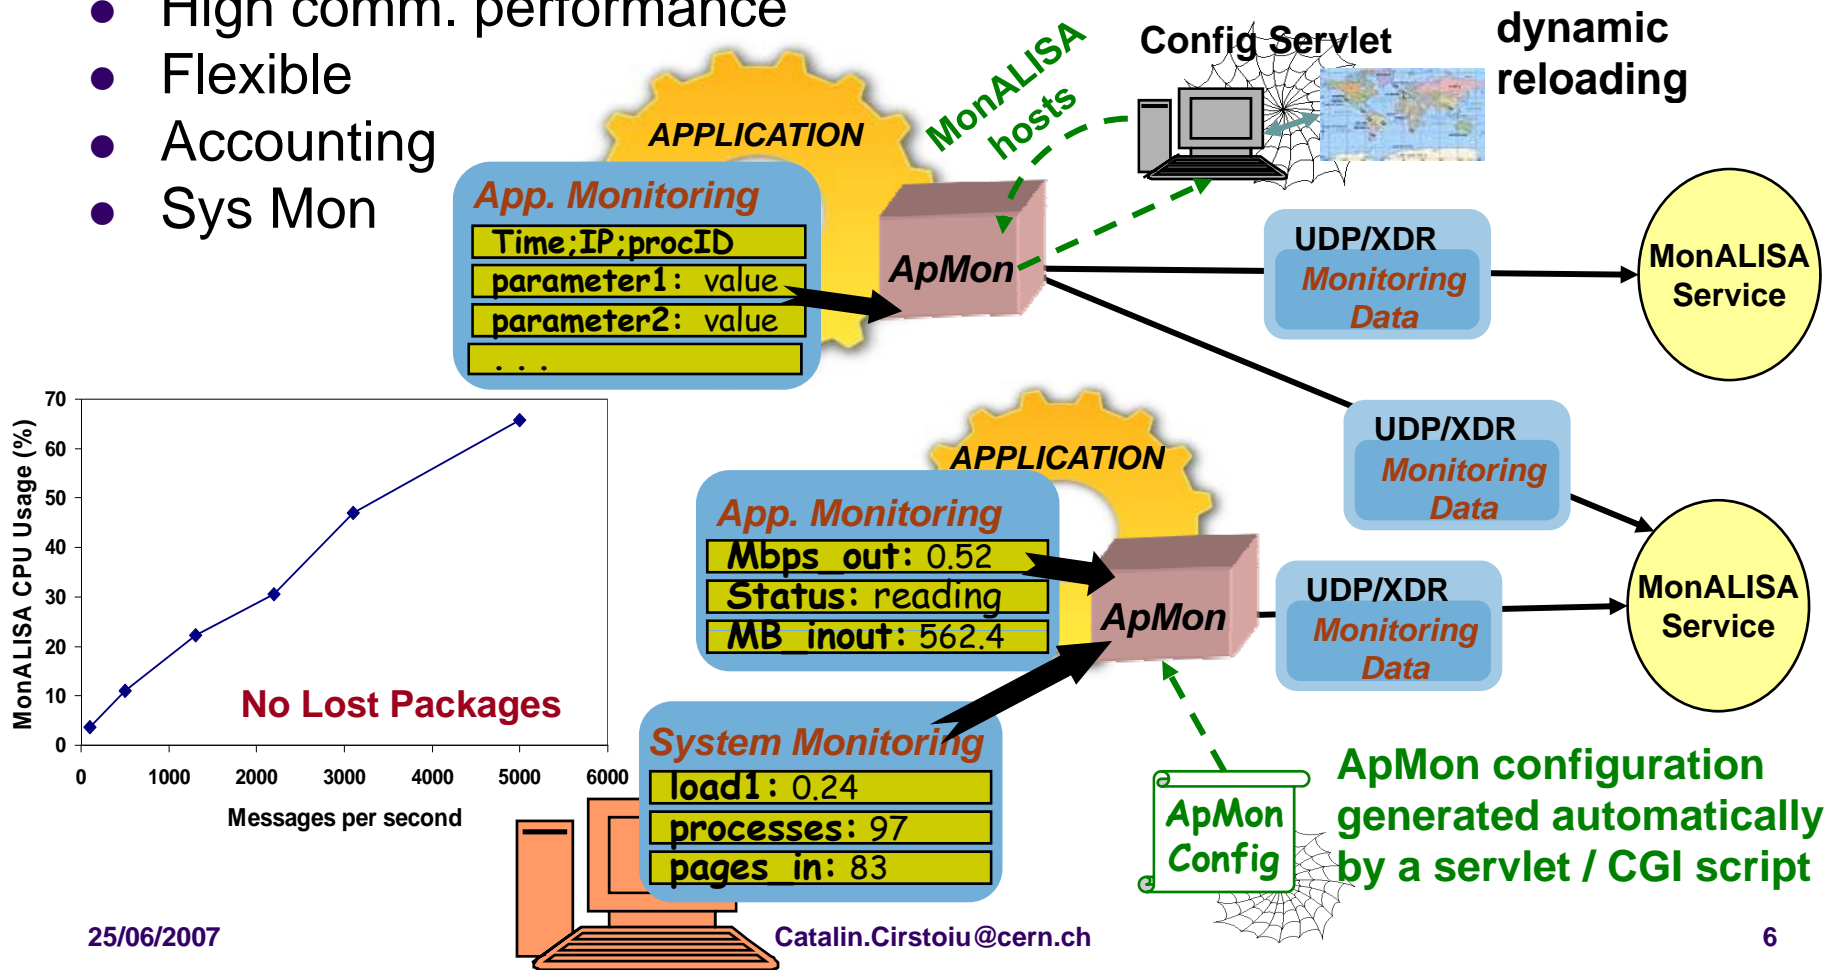
**Dynamic load balancing
Scalability & Replication
Security AAA for Clients**

**Distributed System
for gathering and
Analyzing Information**

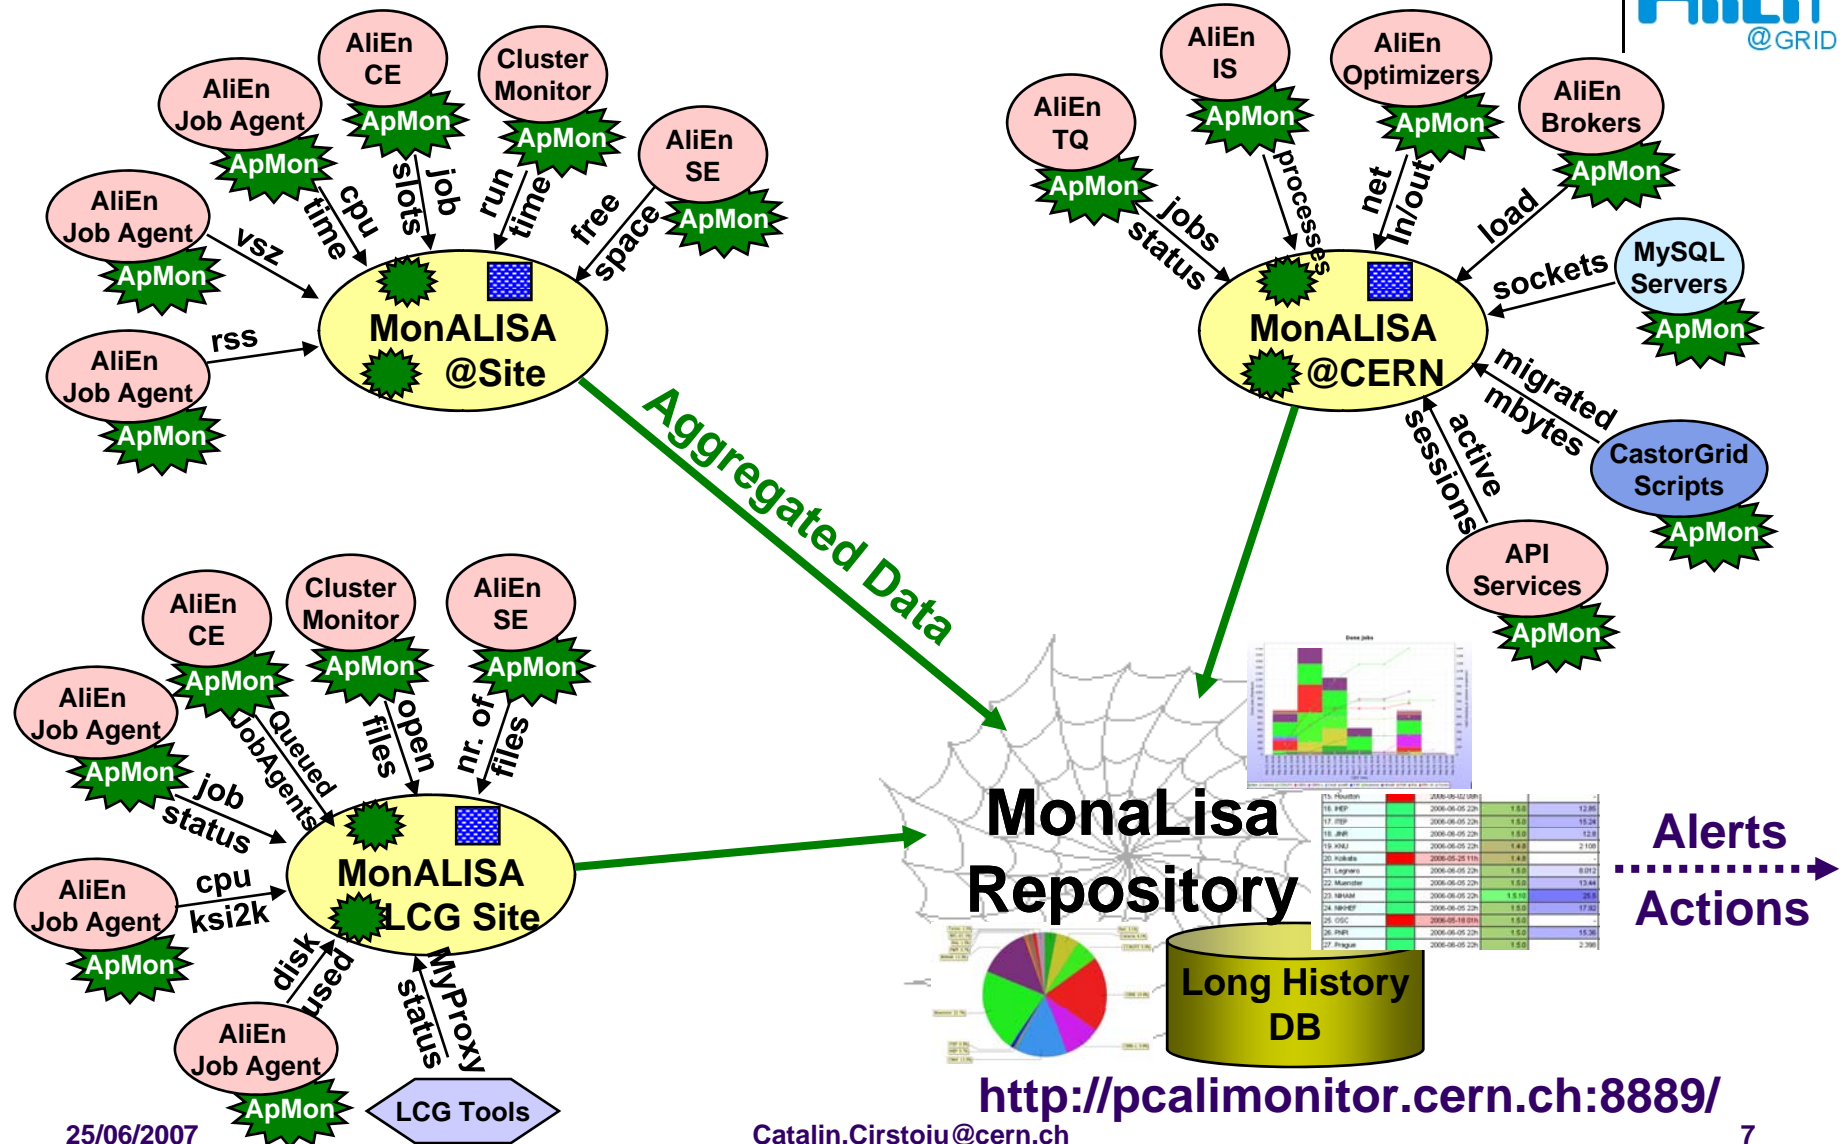
**Distributed Dynamic
Discovery-based on a lease
Mechanism and REN**

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



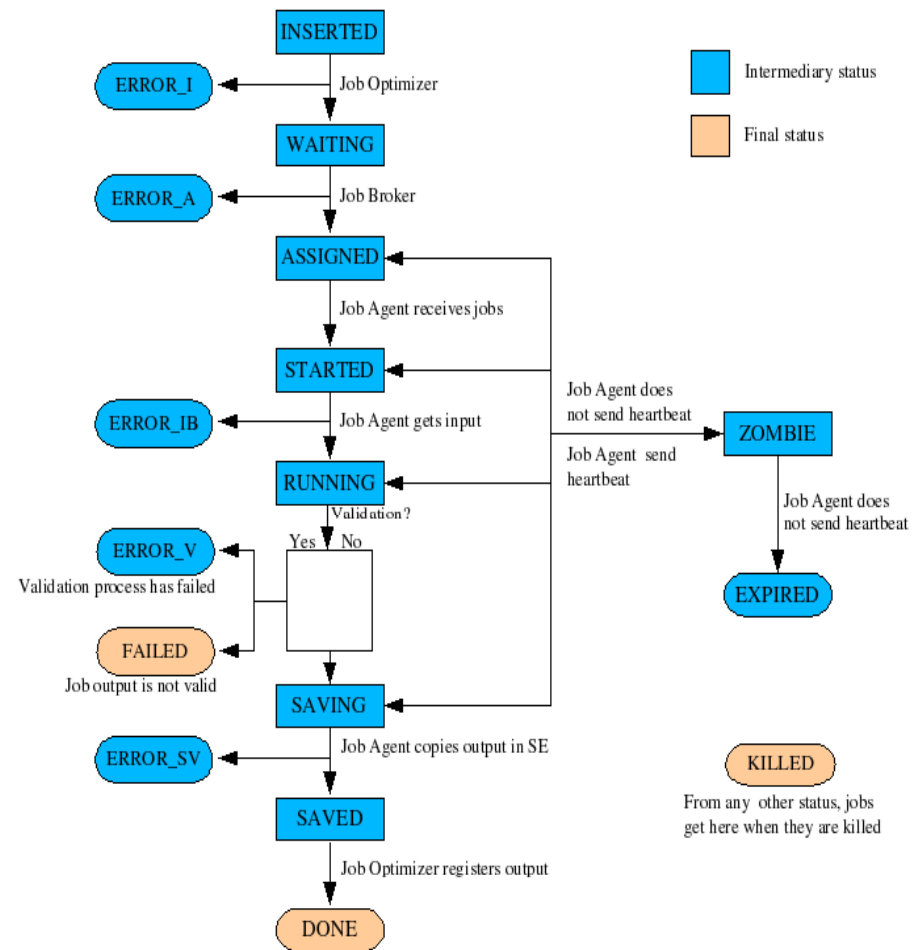
Monitoring architecture in AliEn



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



Real-time Map



MonALISA Repository for ALICE



[Repository Home](#) |
 [Administration Section](#) |
 [ALICE Reports](#) |
 [Events XML Feed](#) |
 [Firefox Toolbar](#) |
 [MonaLisa GUI](#)

ALICE Repository

- ALICE Repository
- Google Map
- MonALISA Map
- Running trend
- Job Information
- SE Information
- Services
- Network Traffic
- FTD Transfers
- CAF Monitoring
- SHUTTLE
- LCG exp. monitoring
- Build system

close all

This page: [bookmark](#), [URL](#)

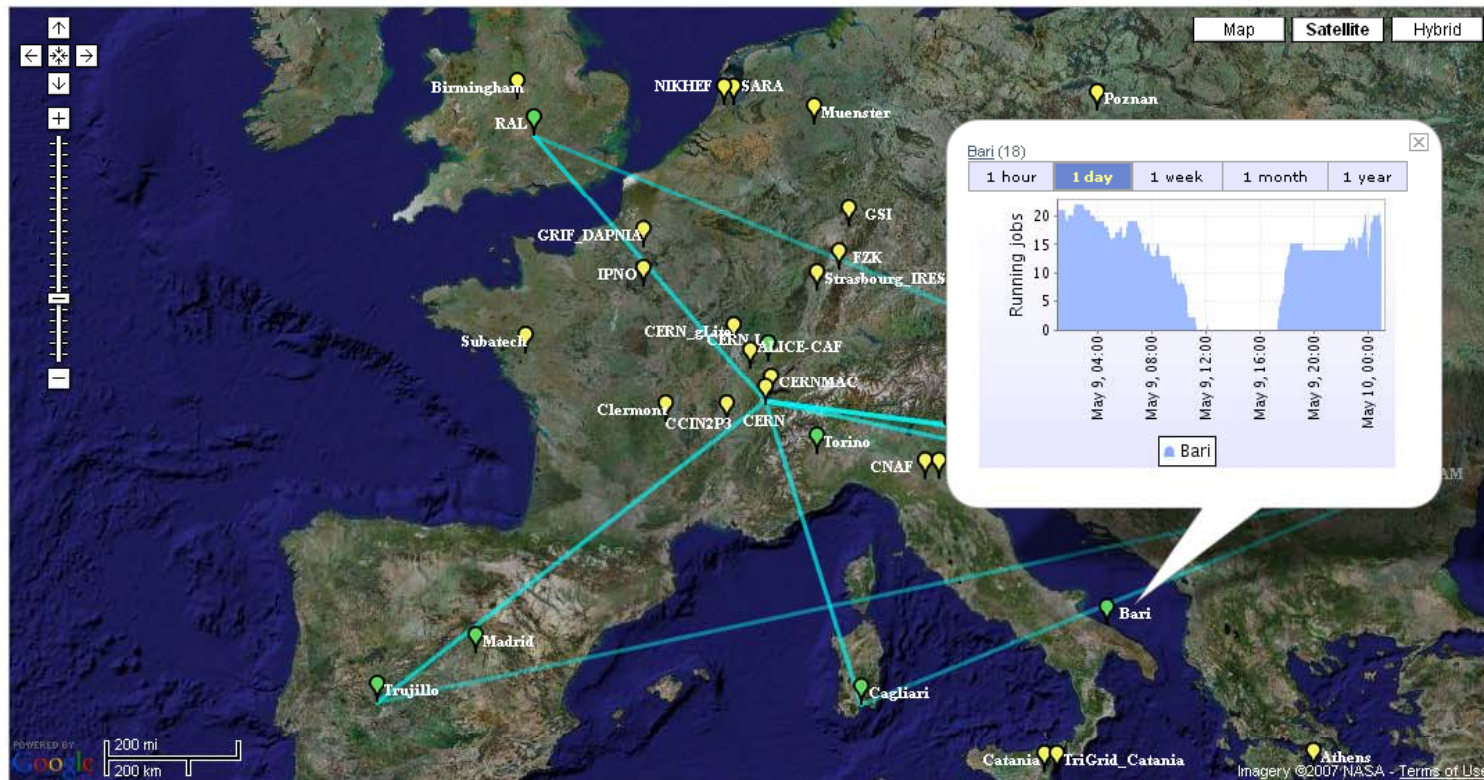
Running jobs trend

Jobs: 1292

Running jobs trend

[24h](#)
[12h](#)
[6h](#)
[1h](#)

(click arrows for detailed view)



● Running Jobs
 ● ML Service Down
 ● No Active Jobs
 ● ML Service Down & no running jobs

Catalin.Cirstoiu@cern.ch

25/06/2007

Integrated Pie Charts



Series ▾ Options Options Alternative Views ▾

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

Done jobs views: History | Histogram

(check all | uncheck all)


Aalborg Athens Bari Birmingham BITP Bologna Cagliari Catania CCIN2P3 CERN CERN_gLite CERN-L
 CERNMAC Clermont CNAF Cyfronet FZK GRIF_DAPNIA GSI Houston IHEP IPNO ISS ITEP JINR
 Jyvaszkyla KFKI KISTI KNU Kolkata Kolkata Kosice LBL LCG Legnaro Madrid Muenster NBI
 NIHAM NIKHEF OSC PNPI Prague RAL RAL RRC-KI SARA Sejong SINP SPbSU Strasbourg_IRES
 Subatech Torino Trigrad TriGrid_Catania Troitsk Trujillo UIB UNAM UPB

Function: Integrate ▾

Interval selection: last month ▾ or « 2007-04-07 21:00 - 2007-05-08 07:00 » Plot

What is this about?

Done jobs statistics



History Plots, Annotations




Series ▾ Options ▾ Alternative Views ▾ (check all | uncheck all)

Aalborg Athens Bari Birmingham BITP Bologna Cagliari Catania CCIN2P3 CERN CERN-L CERN_gLite
 CERNMAC Clermont CNAF Cyfronet FZK GRIF_DAPNIA GSI Houston IHEP IPNO ISS ITEP JINR
 Jyväskylä KFKI KISTI KNU Kolkata Kolkatta Kosice LBL LCG Legnaro Madrid Muenster NBI
 NIHAM NIKHEF OSC PNPI Prague RAL RRC-KI SARA Sejong SINP SPbSU Strasbourg_IRES
 Subatech Torino Trigrad TriGrid_Catania Troitsk Trujillo UIB UNAM UPB
 SUM

Interval selection: 3 months or « 2007-02-06 01:00 - 2007-05-08 08:00 » Plot

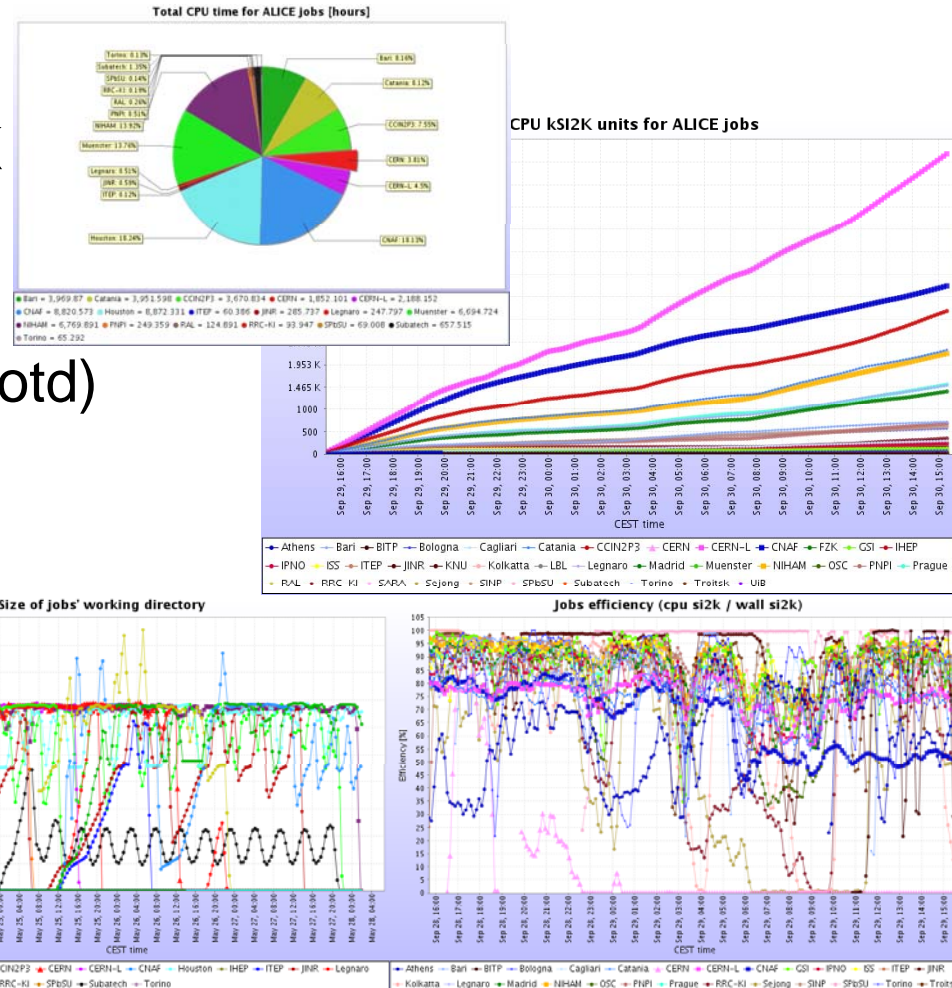
Annotations What is this about?



Running Jobs				
Farm	Last value	Min	Avg	Max
Aalborg	0	0	2.855	35
Athens	1	0	5.021	15
Bari	3	0	3.604	49
Birmingham	0	0	9.695	24
BITP	0	0	9.951	31
Bologna	0	0	6.928	14
Cagliari	13	0	10.62	26
Catania	0	0	67.39	150
CCIN2P3	0	0	78.27	246
CERN	0	-34	21.36	298
CERN-L	169	0	361.4	1342
CERNMAC	0	0	1.083	3
Clermont	0	0	72.32	157
CNAF	279	0	241	450
Cyfronet	0	0	1.789	12
FZK	7	0	176.4	697
GRIF_DAPNIA	0	0	28.52	99
GSI	(16.02.2007 02:55, 20.02.2007 01:05)		5.824	186
IHEP	LCG CE maintenance		20.32	62
IPNO		56.43		226
ISS	36	0	55.44	105
ITEP	3	0	18.46	47
JINR	0	0	7.638	16
Jyväskylä	0	0	1.642	5
KFKI	0	0	19.47	30
KIU	0	0	0.442	5
Kolkata	2	0	14.98	46
Kosice	4	0	7.662	15
LBL	21	0	28.88	47
Legnaro	47	0	48.25	129
Madrid	38	0	42.24	100
Muenster	0	0	50.31	80
NBI	0	0	0.5	4
NIHAM	159	0	97.93	269
NIKHEF	1	0	30.64	53
OSC	9	0	18.16	70
PNPI	0	0	26.16	70

Job Resource Usage


- 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



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



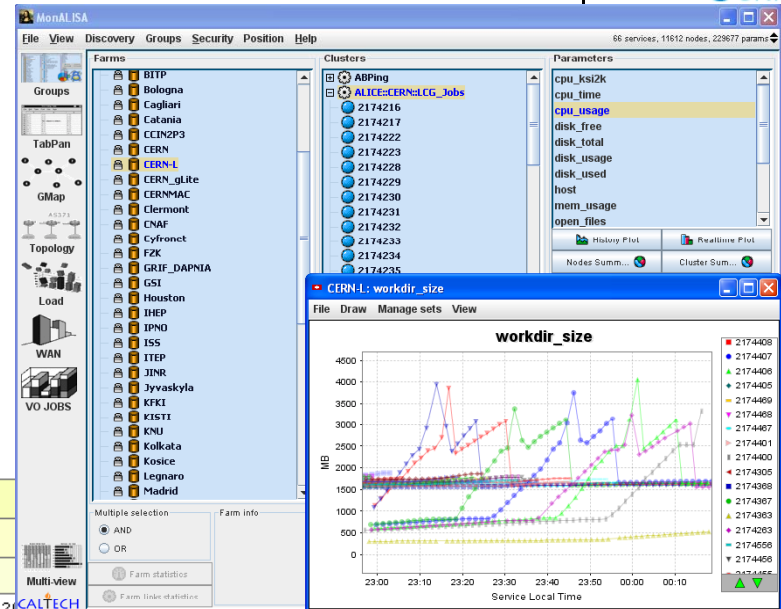
Farm	Last value	Min	Avg	Max	Total
CERN - Castor2	5.827 MB/s	0 B/s	25.4 MB/s	156.3 MB/s	25.6 TB
CERN - se	9.587 KB/s	0 B/s	0.83 KB/s	0.164 MB/s	761.7 MB
IIHAM - File	0 B/s	0 B/s	5.633 KB/s	0.155 MB/s	5.676 GB
Total	5.836 MB/s		25.41 MB/s		25.6 TB

Farm	Last value	Min	Avg	Max	Total
CERII - Castor2	0 B/s	0 B/s	0.861 MB/s	200.9 MB/s	889 GB
CERII - se	10.92 KB/s	0 B/s	5.703 KB/s	0.441 MB/s	5.747 GB
Jyvaskyla - File	0 B/s	0 B/s	16.12 KB/s	1.218 MB/s	16.24 GB
Total	10.92 KB/s		0.883 MB/s		910.9 GB

Individual job tracking



- Based on AliEn shell cmds.
 - top, ps, spy, jobinfo, masterjob
- 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/2006/pp_minbias/5188/
5187	2471497	aliproduct		v5-14-00	v4-04-Rev-10	v1-6-2	21.03.2007 21:07	/alice/sim/2006/pp_minbias/5187/
5186	2470231	aliproduct	8500	v5-14-00	v4-04-Rev-10	v1-6-2	21.03.2007 19:08	/alice/sim/2006/pp_minbias/5186/

pid	owner	first seen	last seen	subjobs	Job states													
					SPLIT	WAITING	STARTED	RUNNING	SAVING	DONE	ERRORS	ERROR_V	ERROR_S					
	aliproduct (1023)	- All -	Active															
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%
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%
2467498	aliproduct	21.03.2007 13:06	22.03.2007 00:07	1000	1					38%	380	0%	9	59%	592	1%	18	0%
2463641	aliproduct	21.03.2007 10:06	22.03.2007 00:07	1000	1					15%	150	0%	4	83%	833	1%	13	0%
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%

AliEn & LCG Services monitoring



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

Global view
 AliEn Services: History of sites' status (efficiency) | Detailed history for each site
 LCG Services: History of sites' status (efficiency) | Detailed history for each site

VOBOX AliEn and LCG services status

What is this about?

AliEn services													
Service	VOBox		AliEn services							AliEn tests			
	Address	AliEn version	Monitoring script	CE	SE	PackMan	Monitor	FTD	add	get	rm	whereis	
1. Aalborg	francis.grid.aau.dk		Failed wit...	-	-	-	-	-	-	-	-	-	-
2. Athens	xg010.inp.demokritos.gr	v2-13.58		DEAD...	DEAD...	DEAD...	DEAD...	-	Fail...	Fail...	Fail...	Fail...	
3. Bari	alicegrid6.ba.infn.it	v2-13.58			May ...			-	Fail...				
4. Birmingham	epbf008.ph.bham.ac.uk	v2-13.58						-	May ...				
5. BITP	alice9.bitp.kiev.ua	v2-13.58						-	May ...				
6. Bologna	boalice6.bo.infn.it	v2-13.58						-	May ...				

- LCG environment and tools
 - Integrating the VoBOX tests previously run by ML within the SAM framework
 - Proxy lifetime, gsiscp, LCG CE/SE, Job submission, BDII, Local catalog, software area etc.
 - Error messages in case of failure
 - Efficiency
 - ML Alerts are used for problems notification

What is this about?

LCG services																		
Service	Site		Proxy							VOBOX					BDII			
	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	Local catalog	Software area
1. Athens		Fail...							-					Fail...			LFC_...	
2. Bari		Fail...							-	Fail...			Fail...	Fail...			LFC_...	perm...
3. Birmingham									-	Fail...							LFC_...	
4. Bologna		Fail...							-	Fail...				Fail...			LFC_...	
5. Cagliari		Fail...		Prox...					-	Fail...				Fail...			LFC_...	
6. Catania		Fail...			Click for more details				-	Fail...				Fail...			Cann...	
7. CCIN2P3					Proxy renewal service not running				-	Fail...					Loca...			

FTD/FTS Monitoring



- Status of the transfers
- Transfer rates
- Success/failures
- Efficiency via ARDA Experiment Dashboard



FTS EFFICIENCY

Click on any Site, and you will have a breakdown according to the errors transferring files to that site

This table presents the transfers that have been done from CERN to the ALICE T1

Transfers done on: Sat 20 Jan 2007

Site (click on any site)	Successful transfers	Failed transfers	Efficiency
ALICE::CCIN2P3::LCG	0	2079	0.00 %
Error message			Counter
The FTS transfer _transferid_ failed (Transfer failed. ERROR the server sent an error response: 425 425 Cannot open port: java.lang.Exception: Pool manager error: Best pool too high : 2.0E8)			2075
The FTS transfer _transferid_ failed (Failed on SRM put: Failed SRM put on httpg://ccsm.in2p3.fr:8443/srm/managerv1			3
executing /usr/local/grid/glite/3.0.4/UI/glite/bin/glite-transfer-submit			1
ALICE::HOUSTON::FILE	152	287	34.62 %
ALICE::CCIN2P3::TAPE	1723	1	99.94 %
ALICE::FZK::LCG	5168	1	99.98 %



VOBox/Head node monitoring



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

Global views: General status
 Detailed history: Load | Cpu usage | Memory usage | Swap usage | eth0 IN | eth0 OUT | eth1 IN | eth1 OUT | Processes | Sockets

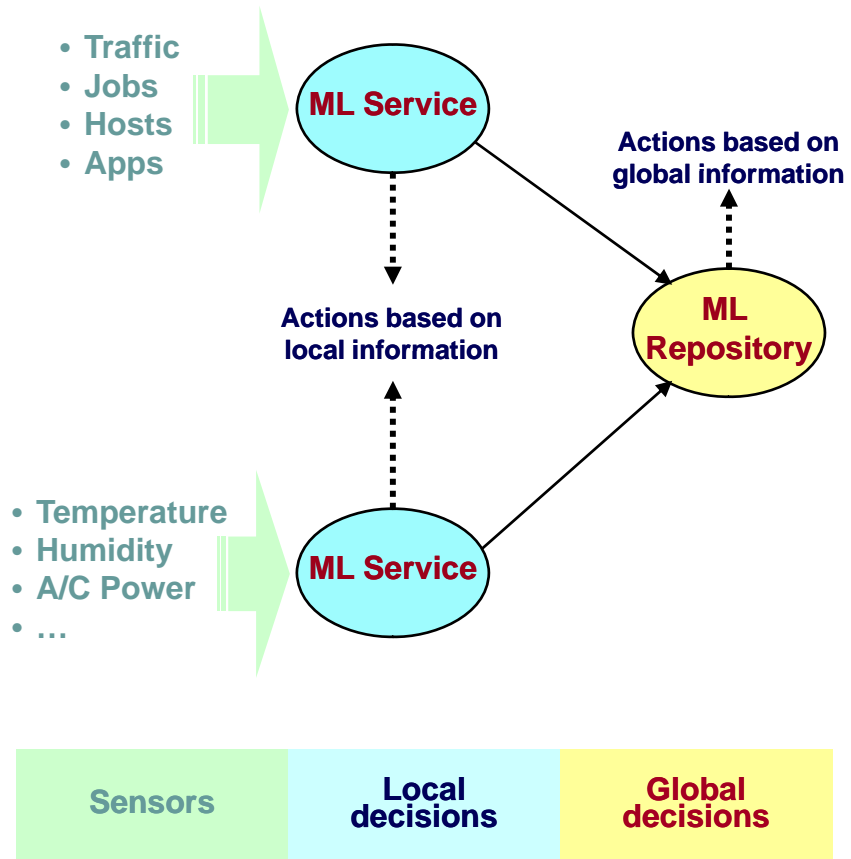
VO Box machine status

What is this about?

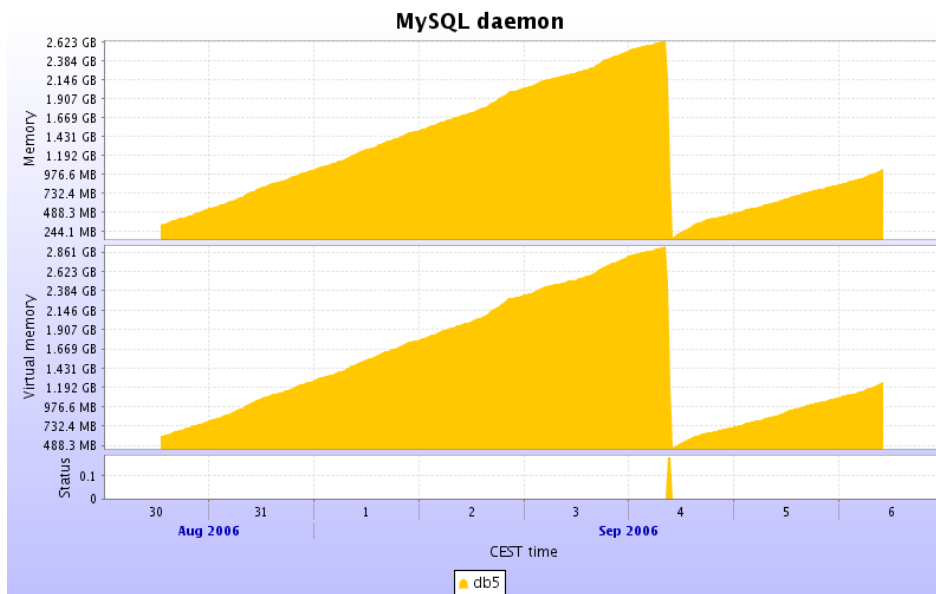
Machine status (last hour average values)																									
Site name	Last seen online	CPU										Mem [% MB]		Swap [% MB]		Eth0 [KB/s]		Eth1 [KB/s]		Eth2 [KB/s]		System			
		Load5	User	System	IOwait	Int	SoftInt	Nice	Steal	Idle	Cnt	MHz	Usage	Total	Usage	Total	In	Out	In	Out	In	Out	Procs	Socks	
1. Aalborg	2007-05-06 08h	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
2. Athens	2007-05-08 06h	0.042	3.449	0.792	0.016	0.067	0.053	0	-	95.62	1	2680	17	2010	0	1992	12.58	5.263	-	-	-	-	57	26	
3. Bari	2007-05-08 06h	0.106	4.206	1.212	0.134	0.023	0.397	0	-	94.03	2	1800	46	1001	6.787	2047	0.196	0.072	27.64	9.307	-	-	74	52	
4. Birmingham	2007-05-08 06h	0.223	8.699	1.612	0.201	0.01	0.373	0	-	89.1	2	800	54	1001	4.093	2047	11.08	2.335	-	-	-	-	91	58	
5. BITP	2007-05-08 06h	0.035	0.737	0.204	0.679	0	0.115	0	-	98.27	4	3192	21	3999	0	8181	1.098	0.25	8.8	1.26	-	-	91	56	
6. Bologna	2007-05-08 06h	0.145	2.904	0.495	0.099	0.006	0.112	0.002	-	96.38	4	3067	17	4005	0	2000	-	-	23.55	11.56	-	-	99	52	
7. Cagliari	2007-05-08 06h	0.044	1.95	0.68	0.021	0.007	0.256	0	-	97.09	2	3199	31	2007	0	2000	16.86	7.491	-	-	-	-	92	63	
8. Catania	2007-05-08 06h	0.323	6.155	1.162	0.744	0.022	0.176	0	-	91.74	4	2799	36	2006	0.193	4094	45.36	26.09	-	-	-	-	98	58	
9. CCIN2P3	2007-05-08 06h	1.012	3.159	10.32	14.8	0.042	0.15	0	-	71.54	4	3000	40	2007	3.791	2047	1.709	0.14	14.97	4.482	-	-	105	61	
10. CERN	2007-05-08 06h	0.66	13.15	2.425	0.167	0.016	1.979	0	-	82.26	4	2387	28	5768	0.1	8000	28.89	13.9	-	-	-	-	305	604	
11. CERN_gLite	2007-05-08 06h	0.106	3.337	1.426	0.436	0.005	0.158	0	-	94.64	2	3000	30	3995	0	4094	9.933	1.555	-	-	-	-	100	37	
12. CERN-L	2007-05-08 06h	1.098	25.04	5.07	5.342	0.136	1.841	0	-	62.58	2	2793	29	3991	0.901	4094	1957	69.1	-	-	-	-	146	107	
13. CERNMAC	2007-05-08 06h	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
14. Clermont	2007-05-08 06h	0.102	5.721	1.142	0.016	0.192	0.173	0	-	92.76	1	2007	17	3013	0	8189	18.16	7.845	-	-	-	-	110	66	
15. CNAF	2007-05-08 06h	0.524	16.14	1.775	0.438	0.048	0.813	0	-	80.79	2	3067	28	4005	0	4000	95.66	43.46	-	-	-	-	208	112	
16. Cyfronet	2007-05-08 06h	0.499	3.537	0.535	12.72	0.027	0.113	0	-	83.07	2	1300	58	1982	40.48	1000	-	-	20.93	15.02	-	-	102	146	
17. FZK	2007-05-08 06h	0.099	0.021	0.594	2.591	0.002	0.362	1.988	-	94.44	4	3000	72	2007	10	3827	9.564	1.506	36.23	1.602	-	-	164	162	
18. GRIF_DAPNIA	2007-05-08 06h	0.026	3.493	0.688	0.042	0.032	0.039	0	-	95.71	1	2793	50	2001	0.044	2047	14.04	6.082	-	-	-	-	99	54	
19. GSI	2007-05-08 06h	1.075	44.73	5.522	0.131	5.664	0.224	0	0	43.73	1	2667	58	820.9	0.636	512	28.11	9.568	-	-	-	-	84	72	
20. Houston	2007-05-08 06h	0.047	3.06	0.581	0.02	0.015	0.386	0	-	95.94	1	1396	18	4014	2.861	4095	0.018	0.036	34.1	14.91	-	-	106	81	
21. IHEP	2007-05-08 06h	0.172	2.942	0.638	0.559	0.011	0.155	0.353	-	95.34	4	2791	55	2006	11.69	1999	21.34	7.266	17.02	94.06	-	-	209	129	
22. IPNO	2007-05-08 06h	0.039	2.081	0.676	0.096	0.003	0.132	0	-	97.01	2	2394	28	3994	0	1992	17.2	8.175	-	-	-	-	98	57	

Actions framework

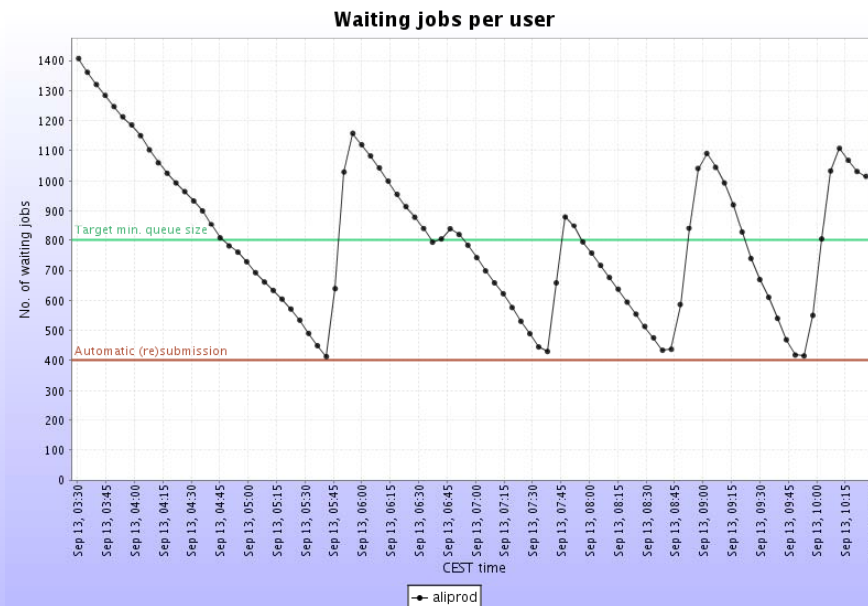
- Based on monitoring information, actions can be taken in
 - ML Service
 - ML Repository
- Actions can be triggered by
 - Values above/below given thresholds
 - Absence/presence of values
 - Correlation between multiple values
- Possible actions types
 - Alerts
 - e-mail
 - Instant messaging
 - RSS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



Fact figures

- Raw parameters when running 4K Jobs
 - Unique data series: 300K with frequency 1-15 minutes
 - Message rate: 16K / minute
- Site aggregated parameters
 - Message rate: 2K / minute
 - Bandwidth rate: 300Kbps
- Repository
 - DB Size: 70 GB ~ 700M records
 - Data Reduction Schema
 - 2 Months with 2 minutes bins
 - 1 Year with 30 minutes bins
 - ~Forever with 2 hours bins
- Response time
 - App -> ML Service – network speed
 - ML Service -> ML Clients
 - Subscribed parameters – network speed
 - One shot requests (history requests) – ~5 seconds
 - Repository dynamic history requests – ~300ms / page
- No incoming ports are required



Summary

- The MonALISA framework is used as a primary monitoring tool for the ALICE Grid since 2004
- Presently the system is used for monitoring of all (identified) services, jobs and network parameters necessary for the Grid operation and debugging
- The add-on tools for automatic events notification allow for more efficient reaction to problems
- The framework design and flexibility answers all requirements for a monitoring system
- The accumulated information allows to construct and implement automated decision making algorithms, thus increasing further the efficiency of the Grid operations

Thank you!



Questions?



<http://alien.cern.ch>



<http://monalisa.caltech.edu>