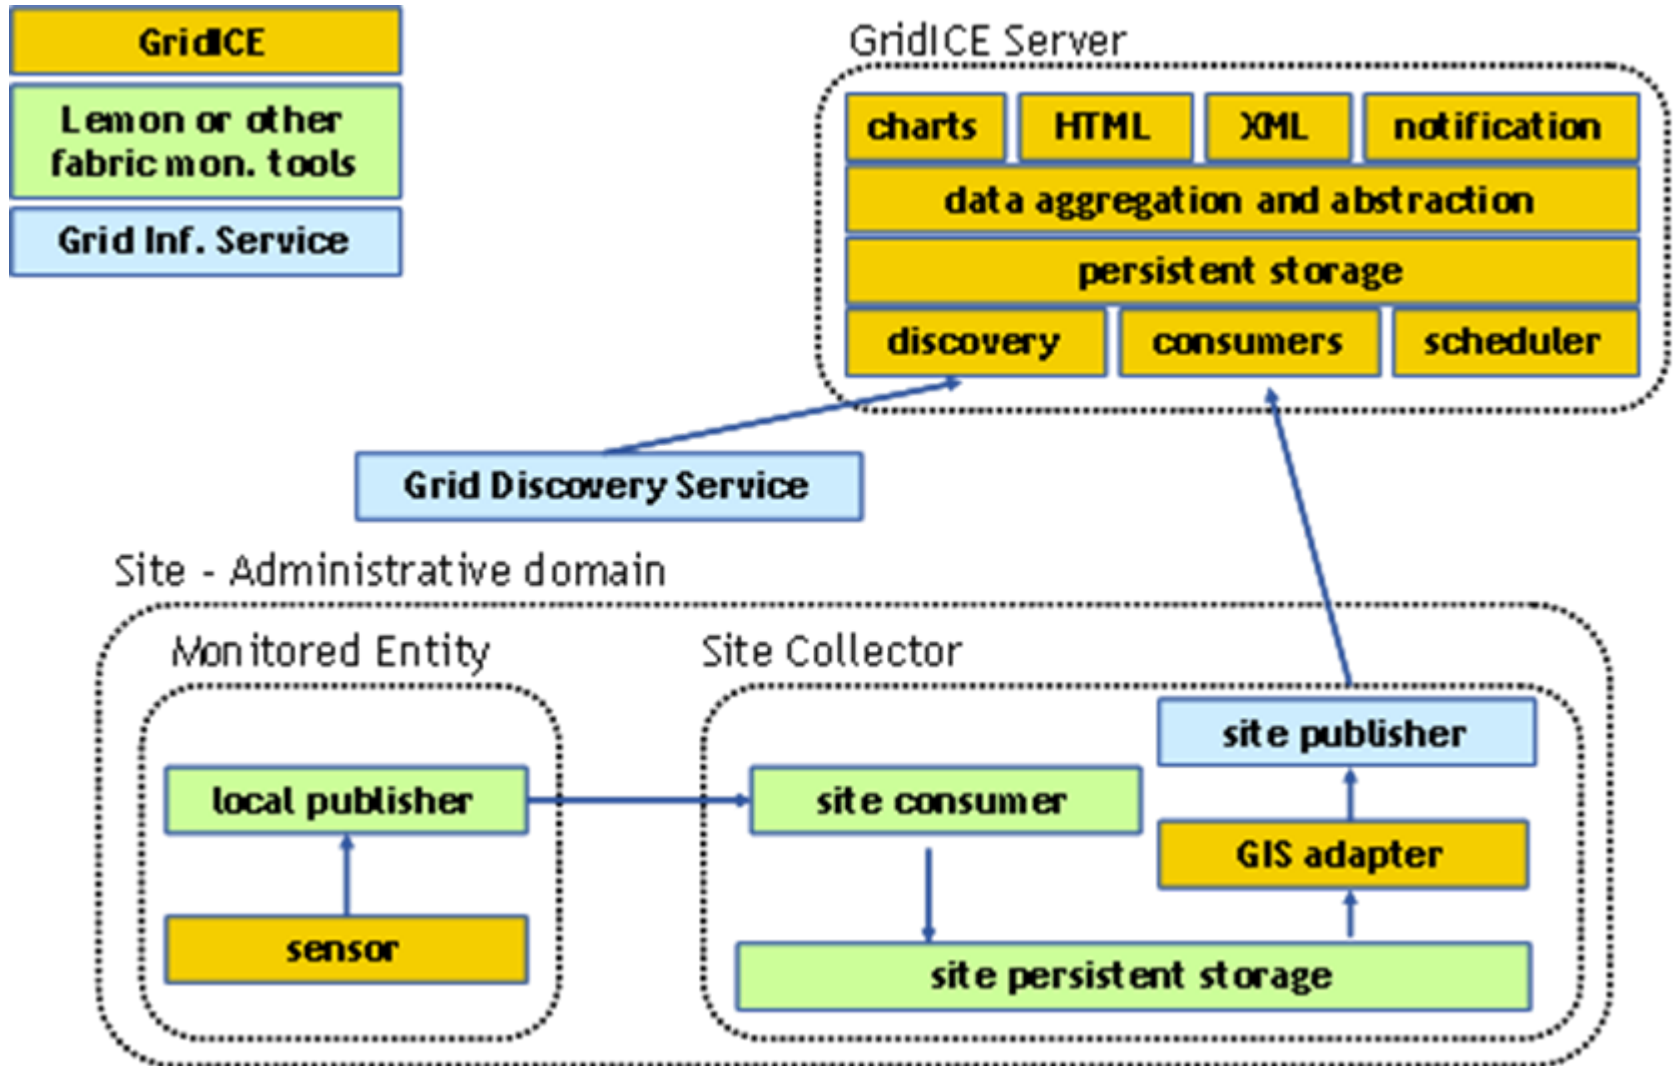


GridICE: Monitoring the user/application activities on the grid

Guido Cuscela
INFN – Bari

CHEP2007 – September 5 , 2007

- **General description**
 - Architecture
 - Main features (old and new ones)
- **What we are monitoring**
 - Fabric and application monitoring
- **Grid monitoring for different users/needs**
 - Details of the VOMS groups, roles and users.
 - Standard user/VO manager/Site Manager/Production Manager
 - Troubleshooting
- **Conclusions**



- **Transport of data (to the site collector) performed by LEMON (site level)**
- **A set of standard measurement for the fabric monitoring provided by LEMON sensors**
- **GridICE adds its own specific sensors (demon status check, job monitoring, LRMSinfo, etc.)**
- **Substantial effort spent to integrate the new LEMON version (v.2.15.x) in GridICE**
 - The new version is at the moment under test on selected sites
- **Grid monitoring integrated with local monitoring**
 - Farms, using LEMON for local monitoring, can add the GridICE Grid monitoring
 - Farms implementing GridICE can benefit from the LEMON farm monitoring
 - by enabling LRF (Lemon RRD framework)

- **complete web-based interface for data presentation**
- **the same data, shown in each view of the web-based interface, can be retrieved in XML format**
- **support for customized graph generation**
- **notification service**
- **automatic discovery of new resources to be monitored through the Grid Information Service**
- **complete set of monitored metrics, from host-related to Grid services related info**
 - supports and extends the GLUE Schema
- **support for the following batch systems: OpenPBS, Torque, LSF**

- **Hardware monitoring:**
 - fabric level monitoring via LEMON sensors
 - new LEMON version integration (under test)
- **Services monitoring:**
 - For every grid node we check the related services (via standard GRIIS)
 - Monitoring of every process/daemon which has to run on nodes
- **Job monitoring:**
 - “lightweight” job monitoring sensors
 - no problems observed to the INFN-T1 with more than 12000 jobs R/Q (more then 5000 on a single CE)
 - jobs info retrieved both from batch system and grid logs
- **LRMS monitoring:**
 - LRMSinfo sensor provides preliminary SLA support and basic site CPU usage efficiency
 - No sensors on WNs (all needed information retrieved from the CE by querying the batch system)


- **Users are required to have a valid CA certificate**
- **Users identification is done through the digital certificate installed in their browsers (DN retrieved)**
- **https protocol used on the server side**
- **Two way to access data**
 1. “Standard Users”: only the info of user own jobs are provided
 2. “High level Users” can ask to be registered to the GridICE web site with a specific role
 - VO manager
 - Site manager
 - Production manager
 - ...

A new attribute was added for role identification

Field	Description
NAME	Job name
JOB_ID	Local LRMS job id
GRID_ID	Grid job unique id
USER	Local mapped username
VO	User VO name
QUEUE	Queue
QTIME	Job creation time
START	Job start time
END	Job end time
STATUS	Job status
CPUTIME	Job CPU usage time
WALLTIME	Job walltime
MEMORY	Job memory usage
VMEMORY	Job virtual memory usage
EXEC_HOST	Execution host (WN)
EXIT_STATUS	Batch system exit status
SUBJECT	User DN
VOMS_ROLES	User VOMS roles

- "localUser=2455" "~~userDN=/C=IT/O=INFN/OU=Personal Certificate/L=Bari/CN=Guido Cuscela~~"
~~"userFQAN=/infngrid/Role=NULL/Capability=NULL"~~ "jobID=https://egee-rb-01.cnaf.infn.it:9000/M61LksOXimoLyEM9yH6pCQ" "ceID=gridba2.ba.infn.it:2119/jobmanager-lcgpbs-cert" "lrmsID=537640.gridba2.ba.infn.it" "timestamp=2007-04-20 13:26:52"
- "localUser=42000" "~~userDN=/C=IT/O=INFN/OU=Personal Certificate/...~~"
~~"userFQAN=/cms/Role=production/Capability=NULL"~~ "jobID=https://rb122.cern.ch:9000/..."
 "ceID=gridba2.ba.infn.it:2119/jobmanager-lcgpbs-short" "lrmsID=536594.gridba2.ba.infn.it"
 "timestamp=2007-04-20 00:00:06"
- "localUser=7002" "~~userDN=/C=IT/O=INFN/OU=Personal ...~~"
~~"userFQAN=/bio/Role=SoftwareManager/Capability=NULL"~~
~~"userFQAN=/bio/Role=NULL/Capability=NULL"~~ "userFQAN=/bio/libi/Role=NULL/Capability=NULL"
 "jobID=https://gridit-rb-01.cnaf.infn.it:9000/..." "ceID=gridba2.ba.infn.it:2119/jobmanager-lcgpbs-short" "lrmsID=461894.gridba2.ba.infn.it" "timestamp=2007-01-15 12:52:30"

- User that has no jobs submitted and no role registered



An authenticated user sees only his/her own jobs

GridICE >> Site::ALL >> Site::ALL

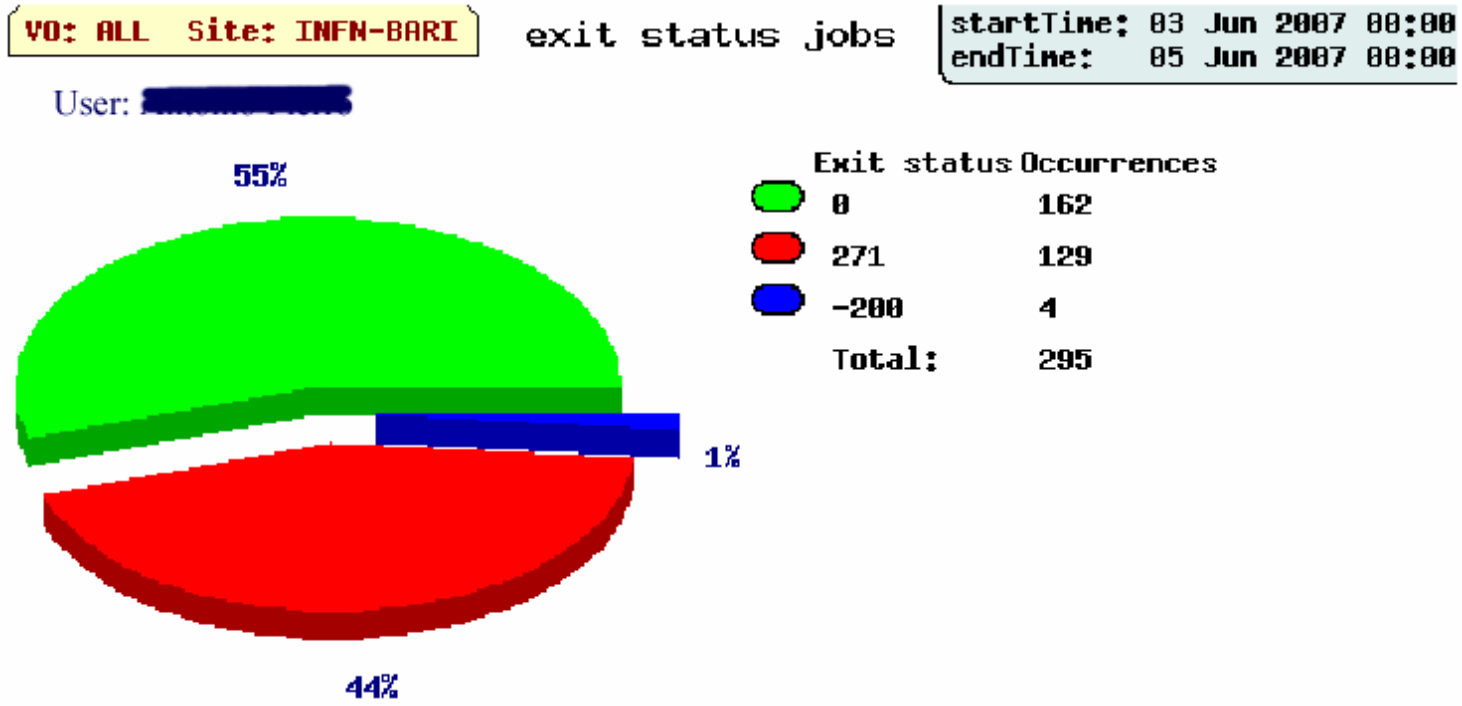
Site: VO: Global ID: Exec-host:
 From: Status: LocalID: Local-user:
 To: Jobs x page: (max 200) Personal Jobs: (Request a browser certificate)
 Time interval:

578 Entries in 25 Pages

1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | Next | Last

#	LocalID	VO	Site	Status	Creation	GlobalID
1	547976	biomed	INFN-BARI	W	2007-05-04 13:57	https://grid09.lal.in2p3.fr:9000/oLr7XHskVhu6x2d7rgx9Sg
2	547975	biomed	INFN-BARI	W	2007-05-04 13:57	https://node04.datagrid.cea.fr:9000/RgehwL_rwf7deuzz9_3Atw
3	547974	biomed	INFN-BARI	W	2007-05-04 13:57	https://grid09.lal.in2p3.fr:9000/cv5HqtWwqzOtn9gyW9vyPg
4	547973	biomed	INFN-BARI	W	2007-05-04 13:57	https://grid09.lal.in2p3.fr:9000/yq4H6U_XT9cfJoGzcRj_Cg
5	547972	biomed	INFN-BARI	W	2007-05-04 13:56	https://grid09.lal.in2p3.fr:9000/TO9_I08r5h3nT-ZCk-U2Aw
6	547971	biomed	INFN-BARI	W	2007-05-04 13:56	https://grid09.lal.in2p3.fr:9000/tHaM9074jcULwRzoH-iAGg
7	547970	biomed	INFN-BARI	W	2007-05-04 13:55	https://node04.datagrid.cea.fr:9000/6hGEkpWPwTD4dG9J33GSpQ
8	547968	biomed	INFN-BARI	W	2007-05-04 13:55	https://grid09.lal.in2p3.fr:9000/Qn3-oXUv6iKU2taqqW4mmg
9	547969	biomed	INFN-BARI	W	2007-05-04 13:55	https://grid09.lal.in2p3.fr:9000/qauNKCLRlqtk-KQ1k6RDJQ
10	547966	biomed	INFN-BARI	W	2007-05-04 13:54	https://grid09.lal.in2p3.fr:9000/GtVgLVu3vOD2LR45PHP6oA
11	547965	biomed	INFN-BARI	W	2007-05-04 13:54	https://grid09.lal.in2p3.fr:9000/QC0C_aeV9ldFXVrXb5DtxQ
12	68842	biomed	INFN-ROMA3	R	2007-05-04 13:53	https://grid09.lal.in2p3.fr:9000/q_icUbjqrhwLcB9wG6ZQkQ
	CPU/Wall: -	Exit: -	RAM: 0	VM: 0	RB: grid09.lal.in2p3.fr	LocalUser: biomed002 Queue: biomed
13	68841	biomed	INFN-ROMA3	R	2007-05-04 13:53	https://grid09.lal.in2p3.fr:9000/taRtF5QHmQcND6bJHCFlgA

An authenticated user sees only his/her own jobs



- **Different types of users require different views (aggregation) of monitoring information:**
 - **Virtual Organization managers** require the ability of observing and analyzing the **performance** of the “actual ” **system they are using** (which can dynamically change over time)
 - Both **site administrators** and **grid operation centre managers** require performance analysis and fault detection of the resources they have to manage
 - **Grid Service developers** require the ability of analyzing the **behaviour of their middleware components** (e.g. how does a resource broker dispatch jobs over a set of available resources)
 - A **production manager** could be interested to see the status of all the jobs of the production for which is responsible


GridICE >> Charts:Group/voManager

CE SE Gris Host Job Charts







Single-Charts | Choose your preferred Site/VO |
 | Local Owner (optional)

Chart for User | Site Manager | **VO Manager** | time intervals

VO: biomed Site: ALL | Farm Usage | startTime: 29 Apr 2007 00:00 | View Chart
 endTime: 02 May 2007 16:43



Name	VO	VOMS group	ROLE	CPUTime (min)	CPUTime (%)	WallTime (min)	WallTime (%)	CPU/WallTime(%)	E (#)	E (%)	R	Q
Total	All	/	/	2659	0,00	13712.68	0,00	19,39	560	100	15	
Jaiphet	biomed			20	0,75	481	3,51	4,16	201	35,89	13	
Roman	biomed			3	0,11	55	0,40	5,45	10	1,79	2	
Schreiner	biomed	/biomed/lcg1	NULL	180	6,77	7060	51,48	2,55	245	43,75		
Jisakul	biomed	/biomed/lcg1	NULL	3	0,11	3205	23,37	0,09	19	3,39		
Royent	biomed			2357	88,64	2640	19,25	89,28	27	4,82		
Marelli	biomed			3	0,11	105	0,77	2,86	23	4,11		
Estro	biomed			68	2,56	83	0,61	81,93	1	0,18		
Hung	biomed	Not used	Not used	28	1,05	56	0,41	50,00	10	1,79		
Gerard	biomed			2	0,08	31	0,23	6,45	23	4,11		
german	biomed	/biomed/lcg1	NULL	1	0,04	1	0,01	100,00	1	0,18		

	CNR-ILC-PISA	-	7	33	-	-	-	40
	INFN-TORINO	-	1	24	-	-	-	25
	ITB-BARI	-	-	79	-	-	-	79
	INFN-ROMA2	-	-	24	-	-	-	24
	INFN-PADOVA	-	-	23	-	-	-	23
	SPACI-NAPOLI	-	-	10	-	-	-	10

GridICE >> Charts:Group/voManager

CE SE Gris Host Job **Charts**

Single-Charts | Choose your preferred Site/VO

Chart for User: **Site Manager** | VO: **biomed** | Local Owner (optional)

VO Manager | time intervals

VO: ALL Site: ITB-BARI **VO Usage** startTime: 29 Apr 2007 00:00 endTime: 02 May 2007 18:33

VO: ALL Site: ITB-BARI **Jobs status** startTime: 29 Apr 2007 00:00 endTime: 02 May 2007 18:33

Waiting Jobs
50.0%

Executed Jobs
52.0%

Name	VO	VOMS group	ROLE	CPUTime (min)	CPUTime (%)	WallTime (min)	WallTime (%)	CPU/WallTime(%)	E (#)	E (%)	R	Q
Total	All	/	/	18149.75	0,11	59436.63	0,37	30,54	628	100	46	1248
Wolke	cms	/cms	NULL	19	0,10	26417	44,45	0,07	133	21,18	16	95
Flippa	cms	/cms/Commissioning/Tracker/cmsTAC	production	9308	51,28	12518	21,06	74,36	109	17,36	20	566
Patrizio	alice	Not used	Not used	5566	30,67	10458	17,60	53,22	26	4,14	9	22
Andreas	cms	/cms	NULL	3063	16,88	4247	7,15	72,12	68	10,83		99
elias	bio	/biomed/lcg1	NULL	1	0,01	1183	1,99	0,08	3	0,48	1	
Andreas	dteam	/dteam/cern	NULL	28	0,15	1105	1,86	2,53	66	10,51		2
Quirijn	atlas	/atlas/lcg1	production	2	0,01	783	1,32	0,26	3	0,48		142
Dierker	atlas	/atlas/lcg1	NULL	6	0,03	472	0,79	1,27	44	7,01		77
Dimitry	atlas	/atlas/lcg1	production	2	0,01	389	0,65	0,51	3	0,48		229
Dimitry	atlas	/ops	NULL	2	0,01	105	0,18	1,90	5	0,80		4
Thomas	atlas	Not used	Not used	1	0,01	45	0,08	2,22	7	1,11		6
Jaiphe	biomed	Not used	Not used	1	0,01	35	0,06	2,86	2	0,32		1
stefano	cms	/cms	NULL	86	0,47	942	1,58	9,13	105	16,72		
Dimitry	cms	/cms/dcms	NULL	59	0,33	306	0,51	19,28	22	3,50		

	-	-	-	1	-	-	-	1				
	cms	-	-	1	-	-	-	1				
Total:		0	2	196	0	0	0	198				

- **GridICE tool can show you only the info you are interested in**
- **With the authentication based on personal certificate, the data privacy is always guaranteed**
- **Work is going on the server side and on the web interface to improve performance and features**
- **GridICE is always open to fit everything your monitoring needs**

GridICE dissemination: <http://gridice.forge.cnaf.infn.it>

GridICE Publications:

- [1] S. Androzzzi, N. De Bortoli, S. Fantinel, A. Ghiselli, G. L. Rubini, G. Tortone, M. C. Vistoli **GridICE: a monitoring service for Grid systems**, Future Generation Computer System 21 (2005) 559–571
- [2] C. Aiftimiei, S. Androzzzi, G. Cuscela, N. De Bortoli, G. Donvito, S. Fantinel, E. Fattibene, G. Misurelli, A. Pierro, G.L. Rubini, G.Tortone. **GridICE: Requirements, Architecture and Experience of a Monitoring Tool for Grid Systems**. In Proceedings of the International Conference on Computing in High Energy and Nuclear Physics (CHEP2006), Mumbai, India. 13-17 February 2006.
- [3] C. Aiftimiei, S. Androzzzi, G. Cuscela, N. De Bortoli, G. Donvito, S. Fantinel, E. Fattibene, G. Misurelli, A. Pierro, G.L. Rubini, G.Tortone. **Flexible notification service for Grid monitoring events**. In Proceedings of the International Conference on Computing in High Energy and Nuclear Physics (CHEP2006), Mumbai, India. 13-17 February 2006.
- [4] S. Androzzzi, A. Ciuffoletti, A. Ghiselli, C. Vistoli. **Monitoring the Connectivity of a Grid**. In Proceedings of the 2nd International Workshop on Middleware for Grid Computing (MGC 2004) in conjunction with the 5th ACM/IFIP/USENIX International Middleware Conference, Toronto, Canada, October 2004.
- [5] C. Aiftimiei, S. Androzzzi, G. Cuscela, G. Donvito, S. Fantinel, E. Fattibene, G. Misurelli, A. Pierro, Vihang Dudhalkar, Giorgio Maggi. **Recent Evolutions of GridICE: a Monitoring Tool for Grid Systems**. In Proceedings of the IEEE International Symposium on High Performance Distributed computing (HPDC 2007), Monterey Bay (California), USA, 25-29 June 2007.