The Life-Science Grid Community

Tristan Glatard CNRS-CREATIS, Lyon

Current LSGC coordinator biomed VO manager

Cracow Grid Workshop 9-11-2011

LSGC overview

Federation of Virtual Organizations (~500 users)

- **biomed** (~20 countries): bioinformatics, medical imaging, drug discovery
- **Isgrid** (NL): bioinformatics
- medigrid (DE): German biomedical and LS community
- pneumogrid (DE): chronic obstructive pulmonary disease
- **vlemed** (NL): Bioinformatics and Medical Imaging

Supporting grid initiatives

- Official: Dutch, French, Italian, Spanish-Portuguese (Ibergrid) and Swiss
- + numerous resource providers (see on http://wiki.healthgrid.org/LSVRC:Index)

Organization

- Open community of grid users (not a project)
- Hosted by the HealthGrid association

Timeline

- 2004-2010: active life-science cluster in EGEE projects
- June 2010: open workshop at the HealthGrid conference
- Summer 2010: agreed on a statement of goals and missions
- June 2011: MoU signed with European Grid Initiative

Outline

• LSGC: general presentation

- Virtual Organizations
- Users
- Resources
- Applications

biomed technical teams

Current technical challenges

- Data management
- Job management
- VO management

Closing

- Feedback
- Community challenges
- What's in it for you?

LSGC: objectives

Represent life-science grid users

- Negotiate resources
- Promote requirements

Coordinate actions

- Serve as a contact point for new users
- Share expertise, resources, data and tools; avoid replication of efforts

Provide technical services

- Operate and support common VOs and services
- Provide targeted user support and application porting

Training and induction

- Organize community-specific training events
- Smooth the learning curve, lower the start-up cost

Dissemination

- Transfer knowledge among VRC partners
- Advertise actions

Life sciences in EGI

Consumed CPU time, 12/2010 – 11/2011

• Total: 1717 x 10⁶ hours



Virtual Research Community

• What is a VRC?

- Community of researchers with common technological interests
- Technical implementation: group of VOs

• Why being a VRC?

- Technical support from EGI through helpdesk
- Send requirements to middleware development

Recommendations to be a VRC

- Have supporting resources, i.e., supporting NGIs
- Involve significant VOs, i.e., represent users

VRCs (or VRC candidates) so far

 Astronomy and astrophysics, computational chemistry, earth sciences, e-Humanities, hydro-meteorology, life sciences, WeNMR, WLCG (High-Energy Physics)

Used software and services

- Authentication and security
 - Hydra file encryption
 - Myproxy servers
 - Robot certificate services
- VO management and operations
 - VO admin dashboard
 - VOMS
 - Monitoring services
 - Application database
- Pilot-job systems
 - DIANE, DIRAC, ToPoS

Workflow engines

- GWES
- MOTEUR
- P-Grade

File management

- LFC + LCG utils
- VBrowser and vlet API
- iRODS
- User and application support
 - GGUS
 - WS-GRASS, GASUC

VO?

Users (~500)

- Authenticated by Distinguished Name (X509 certificates)
- Authorized in VOMS
- Robots (a.k.a portals)

Supporting resources (EGI VOs only)

- Computing elements
 - biomed: 255
 - Isgrid: 37
 - vlemed: 37
- Meta schedulers (WMS)
 - biomed: 37
 - Isgrid: 4
 - vlemed: 4

- Storage elements (disk)
 - biomed: 122 (3.3 PB)
 - Isgrid: 21 (270 TB)
 - vlemed: 21 (280 TB)
- File catalogs
 - biomed: 1
 - Isgrid: 1
 - vlemed: 1

Source: BDII (cclcgtopbdii02.in2p3.fr)

Grid Activity (CPU) – do we need these resources?

CPU time consumed by life-science VOs on EGI



Source: http://accounting.egi.eu

Grid Activity (CPU) – is it significant?

CPU time consumed by life-science VOs on EGI



Source: http://accounting.egi.eu

Grid Activity (jobs) — do we need more resources?

biomed



http://gstat-prod.cern.ch

LHCb



Do we need more resources? or to use them better...

• User has quit!



Grid Activity (data) - biomed

Overall: we don't need more...



• Local: we may need more... or better distribute



• What about data cleanup? best effort, a.k.a. inexistent



Jobs fail

vlemed applications

Medical Imaging

- Functional MRI
- DTI data analysis for disease markers
- Bone removal in CTA scans
- Image registration
- Next Generation Sequencing
- Proteomics
 - Mass spectrometry data analysis pipeline
- Metabolomics?









vlemed "e-bioinfra gateway"



http://orange.ebioscience.amc.nl/ebioinfragate/

vlemed e-BioInfra architecture



Applications in MediGRID-DE

- NeuroImaging
 - "Standard" tools like FSL, Freesurfer, SPM
- Clinical image processing and surgery support
 - Urology, radiology, liver surgery, cardiology, pneumology
- Biosignal processing
 - Sleep medicine and research
- Bioinformatics
 - Gene prediction
 - SNP selection

Hochschule und Wirtsch

University of A

MediGRID-DE: common Tools

- Gridportal
 - webbased access
 - application specific portlets
 - access via domain-specific portals

Gridworkflows

- application interoperability
- abstraction from underlying Grid infrastructure



Hochschule und Wirtsch

MediGRID-DE: common tools

Security enhancements

- Role-based access control
- GSI-enabled medical data transfer
- Userfriendly credential management
- Generic data protection concept
- De-identification components
- Selection of gridnodes depending on security level
- Audit-trails

cationsanfrage	X
verlangt, dass Si :harite.de:443	sich mit einem Zertifikat identifizieren:
arite - Universitaet : "Charite - Universi	edizin Berlin" etsmedizin Berlin"
Zostifikat das a	Identification uprocessing winds
s DFN-Verein ID #2	2:0C:7F:4E]
hiten Zertifikats:	
CN=Dagmar Kreftir Charite - Universita	
12:0C:7F:4E 6.2011 17:35:11 al (Signatur, Verschlüss nes Zertifikatsschlüs Daten-Verschlüssel krefting@charite.d ("N=DEN_Versin 0" eidung merken	Grand de + GRID Proxy Upload Tool by Fraunhofer Mit Interesting to (CALIFIC
,	Called Cadificate Course
	Select Certificate source Source: Certificate from FireFox Keystore
	Passphrase:
	Select certificate
	Certificate: Dagmar Krefting, Medizinische Info
	Issuer: DFN-Verein PCA Grid - G01, DFN-PKI, DFN Validity: 04.06.2010 - 04.07.2011
	Choose MyProxy Username / Password
	Username: Vse DN ?
	Password: ••••••
	Retype PW:
	Set MyProxy Options
	Proxy Lifetime (days): 7 ?
	🗉 Credential Manager 💿 🖨 🗘 🛇
	Credential Manager
	Credential successfully retrieved and stored.
	The following credential is in your credential store:
	Distinguished name Status Infetime
	/Co-DEO-GridOernany /OU=Charte - Universitatesmedizin Active 0 d, 23 h, Berlin/OU=Medizinische 59 m, 35 s Informatii/CN-Dagmar Kreiting
	To retrieve a credential, enter your credential's MyProxy password and press the button.
	M/Proxy password Retrieve Remove

Hochschule und Wirtsch

Distribution of biomed users in 2010



International catch-all VO for life sciences

Users

- 280 users from ~20 different countries
- Two SMEs (non commercial activities)
- Application fields: Bioinformatics, Drug discovery, Medical Imaging

Resources

biomed

- No formal agreement with sites
- Benefit from EGI-sites agreements

Tools used to access the infrastructure

- Heterogeneous tooling (portals, workflow engines, pilot-job systems): DIANE, DIRAC, WISDOM, OpenMole, Moteur, etc
- No central job submission control point

biomed: applications

Bioinformatics

- Protein clustering (CD-HIT)
- Network for flu observation (g-INFO)

Drug discovery

- Protein docking (MRPD)
- Docking on Malaria (WISDOM)

Medical imaging

- Simulators (ultrasound, MRI, PET, CT, radiotherapy)
- Image filtering

+ unknowns!



Source: http://appdb.egi.eu

biomed: virtual imaging platform

http://vip.creatis.insa-lyon.fr

Application repository



• Examples



Paraster Paraster Using FIE [L. Grevillot, D. Sarrut, 2011]

CPU = 8 hours



Parasternal short-axis echocardiography simulation using FIELD-II [O. Bernard, 2011] VRC:Index Cierro de Investigaciones Energéticas, Medioambientales y Tecnológicas



• jModelTest / ProtTest3

- likelihood scores to establish the BEST-FIT model of DNA/protein evolution (88/120 models)
- widely employed (25.000/5.000 active users)
- Reference: http://darwin.uvigo.es/
- Codes developed by Prof. David Posada (U. Vigo)
 - HPC & Grid versions developed by U. Vigo, U. Coruña and CIEMAT
- Both of them based on PhyML
- Grid versions making use of GridWay & DRMAA
 - EGI & GISELA applications
 - Reference: http://www.ciemat.es/portal.do?IDR=1481&TR=C

(contact: Rafael Mayo)

Outline

LSGC: general presentation

- VOs
- Users
- Resources
- Applications

biomed technical teams

Current technical challenges

- Data management
- Job management
- VO management

Closing

- Feedback
- Community challenges
- What's in it for you?

Technical teams: goals

- Monitor resources supporting the VO
 - Proxies can be generated, files are accessible, jobs can be run
 - Check issues and report to sites & NGIs

Be a technical interface between sites and users

- React to site & NGI requests to users
- e.g. "user **x** was banned to to jobs consuming 10GB of RAM"
- Help users finding solutions

Handle operational issues

- Full storage
- Storage decommissioning

Investigate technical issues

- e.g. "how comes that information system reports wrong figures?"
- Improve support efficiency, QoS

Technical teams: organization

Participants

- 10 organisations
- 9 teams, 14 people
- Volunteer effort of most active VO users

Support organized in shifts

- 1 week shift for one-member team
- 2 weeks shift for 2+-members teams
- Phone conference at each beginning of shift

• Wiki: http://wiki.healthgrid.org/Biomed-Shifts:Index

- Minutes of meetings
- Daily tasks and procedures
- Operation tools

Tools

- VO dashboard
 - Centralizes monitoring and operational tools

https://vodashboard.lip.pt











Specific tools

- File management tools (LFCBrowseSE)
- List DNs, files migration
- Report generator for SE



Nagios: monitoring probe results and downtimes

- Provided by EGI
- Tests for Storage Elements, Computing Elements, Workload Management Systems
- Issues of transient alarms: when should they be reported?

Incident reporting

Global Grid User Support (GGUS)

- Submit incidents to sites, react to VO incidents
- "Team" tickets, shared by shifters

GGUS
Did you know 🎌
Documentation
Training
Registration
💼 🛐 🖂
Search ticket
Submit ticket
Support staff
Navigation on top

GGUS Team tickets

All GGUS TEAM tickets related to VO:biomed All TEAM members of VO biomed • here ► ID Status Submitter Date Info ▶ 76119 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel 2011-11-08 (almost) No space left for biomed users on SE bohr3226.tier2.hep.manch.. in progress ▶ 76117 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel 2011-11-08 No space left for biomed users on SE se.polgrid.pl assigned 76113 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel solved 2011-11-08 SE srm2.grid.sinica.edu.tw is not working for the biomed VO 76112 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel in progress 2011-11-08 SE se.cat.cbpf.br is not working for the biomed VO 76111 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel in progress 2011-11-08 SE se03.grid.acad.bg is not working for the biomed VO 76110 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel waiting for reply 2011-11-08 SE se03.esc.gmul.ac.uk is not working for the biomed VO 76109 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel solved 2011-11-08 SE fal-pygrid-30.lancs.ac.uk is not working for the biomed VO • 75898 /O=GRID-FR/C=FR/O=CNRS/OU=CREATIS/CN=Tristan Glatard 2011-11-02 VOMS web interface of VOs biomed and shiwa doesn't work solved 75722 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel 2011-10-26 Biomed VOMS admin web interface unresponsive verified 75714 /O=GRID-FR/C=FR/O=CNRS/OU=LPC/CN=The Quang Bui 2011-10-26 CE ce01.eela.if.ufrj.br is not working for the biomed VO solved ▶ 75713 /O=GRID-FR/C=FR/O=CNRS/OU=LPC/CN=The Ouang Bui 2011-10-26 SE gridsrm.pi.infn.it is not working for the biomed VO solved 75376 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel 2011-10-17 Biomed VOMS Admin web interface unreachable verified 75358 /O=GRID-FR/C=FR/O=CNRS/OU=CREATIS/CN=Tristan Glatard solved 2011-10-16 VOMS web interface is down • 75321 /O=GRID-FR/C=FR/O=CNRS/OU=CREATIS/CN=Tristan Glatard 2011-10-14 biomed VOMS web interface is down solved 75310 /O=GRID-FR/C=FR/O=CNRS/OU=CREATIS/CN=Tristan Glatard verified 2011-10-14 LFC is getting slow ▶ 75177 /C=IT/O=INFN/OU=Personal Certificate/L=Bari/CN=Guido Cuscela in progress 2011-10-10 CEs ce05-lcg.cr.cnaf.infn.it ce06-lcg.cr.cnaf.infn.it are not working 75157 /C=IT/O=INFN/OU=Personal Certificate/L=Bari/CN=Guido Cuscela verified 2011-10-10 CE axon-g01.ieeta.pt is not working for the biomed VO 74958 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel verified 2011-10-05 Biomed VOMS certificate will expire on Nov. 3rd 2011 74887 /C=IT/O=INFN/OU=Personal Certificate/L=Bari/CN=Guido Cuscela reopened 2011-10-03 SE serv02.hep.phy.cam.ac.uk is not working for the biomed VO 74885 /O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Franck Michel verified 2011-10-03 Biomed VOMS admin web interface unresponsive ▶ 74882 /C=IT/O=INFN/OU=Personal Certificate/L=Bari/CN=Guido Cuscela verified 2011-10-03 SE tbn18.nikhef.nl is not working for the biomed VO ▶ 74880 /C=IT/O=INFN/OU=Personal Certificate/L=Bari/CN=Guido Cuscela 2011-10-03 Problem with nagios box at grid04.lal.in2p3.fr verified ▶ 74878 /C=IT/O=INFN/OU=Personal Certificate/L=Bari/CN=Guido Cuscela verified 2011-10-03 SE moboro.uniandes.edu.co is not working for the biomed VO 74681 /O=GRID-FR/C=FR/O=CNRS/OU=LPC/CN=Paul De Vlieger 2011-09-26 grid-lab-ce.ii.edu.mk errors for biomed users assigned

Ticketing activity

Identification and follow-up of issues (~1 ticket a day)

- SEs (81 tickets in 2011)
- CEs (31 tickets, started 07/2011)
- VOMS, LFC, Nagios... (13 tickets in 2011)

Follow up of tickets from NGIs or users

- "VO Support" tickets (18 tickets in 2011)
- VO users management
- => Significant effort (~1 FTE)
- => Most incident reports could be automated

Decommissioned storage

Problem detection

- Lucky day: get information from broadcast emails
- Otherwise: when the SE is decommissioned

Problem resolution

- Open a GGUS team ticket including the expected decommissioning date
- Ask the site admin to forbid lcg-cr writes from now on
- while (decommissioning deadline is not passed)
 - Ask list of LFNs/DNs having SURLs on SE
 - Notify the concerned users
- end while
- Send decommissioning green light to SE.

Long-term fixes

- Improve decommissioning notifications
- Data migration service

Full storage

Tested functionality

- Icg-cr does not produce "No space left on disk" message
- Critical issue when jobs write on site's default SEs

Problem detection

• Same as before (i.e. once the problem happens...)

Problem resolution

- Open team ticket ; assign it to VOSupport
- Check the **consistency of lcg-infosites** w.r.t available space
- Ask list of LFNs/DNs having SURLs on SE to the LFC admin
- Ask the users to move or delete their data
- Send reminders every week until at least 50% (or at least 500 Go for big SEs) of the SE space for biomed is free

Long-term fixes

- Raise alarms when x% of the SE is full
- Provide data migration service (can crash LFC)

Outline

LSGC: general presentation

- VOs
- Users
- Resources
- Applications
- biomed technical teams

• Current technical challenges (see https://rt.egi.eu/rt/index.html)

- Data management
- Job management
- VO management

Closing

- Feedback
- Community challenges
- What's in it for you?

Data management issues

Storage space management

- Data cleanup is manual
- Data distribution is manual
- Scalability
 - Workload management on data servers
 - Management of many small files

Administration tools (see operations)

- File migration
- Storage decommissioning
- Replica access control
 - Consistency among sites and catalog
- File update
 - Currently, update = delete + write

Data management: replica selection

• Example: file transfer from

List replicas

[glatard@kingkong ~]\$ lcg-lr lfn:/grid/biomed/creatis/tristan/hello.txt srm://ccsrm.ihep.ac.cn/dpm/ihep.ac.cn/home/biomed/generated/2011-11-08/file849bd... srm://ccsrm02.in2p3.fr/pnfs/in2p3.fr/data/biomed//generated/2011-11-08/file78e8601...

Transfer file

[glatard@kingkong ~]\$ lcg-cp -v lfn:/grid/biomed/creatis/tristan/hello.txt file:hello.txt [...]

Trying SURL srm://ccsrm.ihep.ac.cn/dpm/ihep.ac.cn/home/biomed/generated/2011-11-08/file84...
[...]

Transfer took 5230 ms

	Top 10 Users ordered by Normalised CPU time (kSI2K)													
User		Jobs		CPU time		Norm. CPU time		WCT		Norm. WCT		CPU Efficiency	Avg. CPU time	Avg. WCT
#	ID	#	%	Hrs	%	Hrs	%	Hrs	%	Hrs	%	%	Hrs	Hrs
L	Top User 1	327,732	8.2%	2,342,353	35.5%	4,762,752	36.0%	2,456,514	24.4%	5,017,070	29.9%	95.4	7.15	7.50
2	Top User 2	478,234	12.0%	594,550	9.0%	1,234,794	9.3%	672,749	6.7%	1,384,461	8.3%	88.4	1.24	1.41
3	Top User 3	407,681	10.2%	540,699	8.2%	1,093,894	8.3%	688,314	6.8%	1,405,670	8.4%	78.6	1.33	1.69
1	Top User 4	41,601	1.0%	249,201	3.8%	496,747	3.8%	254,693	2.5%	505,659	3.0%	97.8	5.99	6.12
5	Top User 5	181,059	4.5%	142,234	2.2%	300,632	2.3%	298,186	3.0%	607,159	3.6%	47.7	0.79	1.65
ò	Top User 6	71,111	1.8%	89,270	1.4%	200,855	1.5%	97,433	1.0%	219,224	1.3%	91.6	1.26	1.37
7	Top User 7	365,744	9.1%	77,915	1.2%	154,955	1.2%	169,651	1.7%	319,660	1.9%	45.9	0.21	0.46
3	Top User 8	22,286	0.6%	71,267	1.1%	151,988	1.1%	73,885	0.7%	157,384	0.9%	96.5	3.20	3.32
)	Top User 9	27,448	0.7%	53,385	0.8%	111,614	0.8%	56,730	0.6%	119,418	0.7%	94.1	1.94	2.07
0	Top User 10	28,701	0.7%	46,024	0.7%	99,882	0.8%	141,378	1.4%	330,751	2.0%	32.6	1.60	4.93
	Others (DN known)	665,840	16.7%	237,087	3.6%	465,364	3.5%	407,474	4.0%	838,530	5.0%	58.2	0.36	0.61
	Others (DN unknown)	1,381,599	34.5%	2,160,419	32.7%	4,167,670	31.5%	4,766,508	47.3%	5,854,812	34.9%	45.3	1.56	3.45
	Total	3,999,036		6,604,404		13,241,147		10,083,515		16,759,798		65.5	1.65	2.52

• **Consequence**... and/or people implement their own solution

Job management: pilot jobs

٢

٩



Job management: parallel jobs (MPI)

Tested 32-core MPI "Hello, world!" in biomed

- Sites supporting MPI: 30
- Reliability: 20 jobs were successful
- Average queuing time: 90 minutes [49 seconds 5 hours]

Reproduced in production

Possible solutions

- Report errors to sites, use mpi start wrapper (OK)
- Investigate scheduling of MPI jobs
- MPI support in pilot-job systems (e.g. DIRAC)

Outline

LSGC: general presentation

- VOs
- Users
- Resources
- Applications
- biomed technical teams

Current technical challenges

- Data management
- Job management
- VO management

Closing

- Feedback
- Community challenges
- What's in it for you?

LSGC: feedback on the VRC model

VOs are intensively used on a daily basis

Many (small) local VOs increase technical burden

- Services are hardly shared
- Scaling up to larger VOs is not trivial

Mutualize operations, infrastructure effort

- International VO
- Volunteer teams of shifters, from most active users

Structure, organization

- Lightweight, open structure
- Main beneficiaries are natural contributors

Volunteer funding leads to sustainability

LSGC: community challenges

Develop domain-specific activities

- Current effort mostly aims at technical issues
- Support application services, databases, etc
- Foster scientific collaborations in / across VOs

Liaise with scattered user communities

- Collect requirements, issues, publications, usage scenarios
- Low response rate... except when technical issues occur

Few dedicated resources

- Foster contribution of (even small) local clusters
- Clouds? desktop grids? non-gLite VOs?

Ensure sustainability of the VRC model

No formal agreements with (manpower, resource) providers

LSGC: the next day

Grid mostly used for computing (+temp storage)

- Application "challenges" or portal services
- => application sharing?
- => data sharing?
- => scientific collaborations?

Infrastructure is mostly EGI, gLite

- => non EGI VOs?
- Iocal resources, fine-grained access?
- => cloud resources?

VO operations mostly aimed at technical issues

... but shift teams balance load

What's in it for you, life scientists?

Established VOs

- Discuss technical issues with other VOs
- Submit requirements to EGI
- Advertise your services ; use others'

New communities

- Join existing VOs
- Use existing, monitored resources
- Get information on related applications and porting efforts
- Get application porting support

=> Open policy to resource access, support and development

What's in it for you, computer scientists?

Practical assumptions still hardly studied

- Non-clairvoyant job scheduling (applications, resources)
- Online job scheduling
- Semi-automatic file placement and replication

Assistance for issue reporting

- Transient alarms
- Diagnosis methods and tools

Evaluation and validation

- In production conditions
- At a large scale

=> LSGC will provide logs, describe issues, give feedback

Thank you!

Contact & more information

- Wiki page: http://wiki.healthgrid.org/LSVRC:Index
- Mailing list: lsvrc@healthgrid.org

Credits

- vlemed: Silvia D. Olabarriaga
- MediGRID-DE: Dagmar Krefting
- Tools and support: EGI.eu
- Resource providers