Customizing, Extending, and Tweaking: A story of how OSG uses the Globus Toolkit

Alain Roy

Things do not change; we change
–Thoreau

Open Science Grid

Who Am I?

- Open Science Grid Software Coordinator
- Condor Team Member
- In an earlier life, a graduate student with the Globus Project
- An aspiring bread baker and cheese maker
- A wannabe amateur magician

What is Open Science Grid?

- OSG brings together computing and storage resources from campuses and research communities into a common, shared high-throughput computing infrastructure over research networks via a common set of software.
- High-Throughput? For 7-April:
  - ~570,000 jobs
  - ~1.2 million CPU hours
  - 96% success rate

Who is OSG?

- About 35 virtual organizations
  - Support US-LHC experiments (CMS, ATLAS)
  - LIGO (Gravitational waves)
  - SBGrid (Structural Biology Grid)
  - Other high-energy physics (CDF, Dzero…)
  - Many other smaller research groups
- Nearly 80 sites
  - Mostly in the US, but also in South America, China, and South Korea
- Funded by NSF & DOE

Open Science Grid

How do we use Globus?

- Primarily in three ways:
  - GSI for most authentication
  - Gatekeeper/Job Manager to accept jobs to sites
    - Often in combination with glide-in systems for distributing jobs
    - About 80 gatekeepers deployed
  - GridFTP to accept data at sites.
    - Often in combination with underlying storage management system such as Bestman or dCache
    - dCache uses GridFTP protocol but custom implementation
- Across about 80 sites

Customizing, Extending & Tweaking Globus

- Globus encourages customization in (at least) two ways:
  - GRAM job managers are written in Perl (easily examined and changed)
  - GridFTP allows plugins through well-defined interface (DSI)
- OSG has taken advantage of both of these
  - OSG software distribution includes several customizations
  - Individual sites have made customizations

OSG Customization #1: Accounting Integration

- Job managers call outs:
  - Job Manager saves user info (DN/FQAN)
  - Condor Job Manager has direct callout
  - For PBS/LSF/SGE, we watch from the sidelines (no Globus modification)
- Upload all information to Gratia accounting database
  - Central DB
  - Sites can have local DB for their own user
- Lots of results…

Gratia Example #1: Graphs

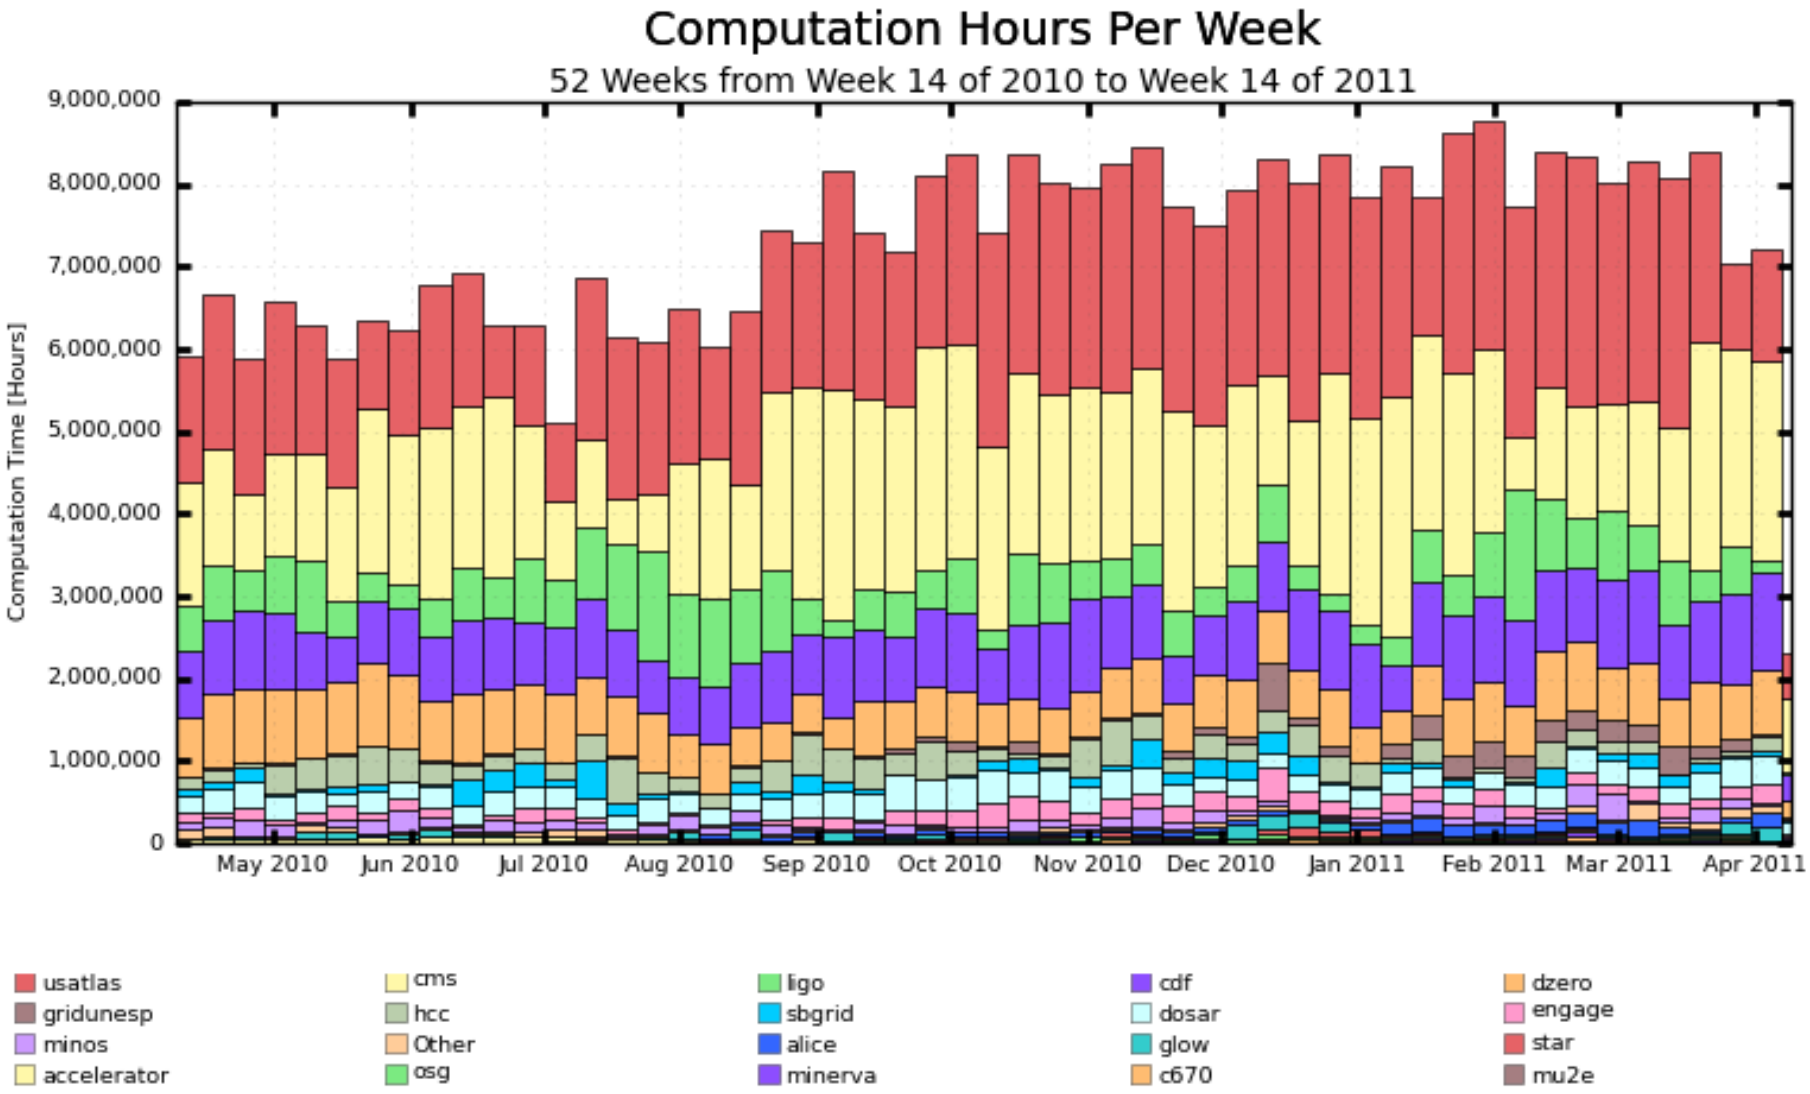

Gratia Example #2: Summary Emails

	VO	Success Rate	Wall Success	Wall Failed	Success	Failed
…						
16	cdf	99 %	174,591	1	20216	33
17	ligo	99 %	5,204	0	578	1
18	nova	99 %	11,629	22	2788	50
19	osg	99 %	2,567	6	14728	379
…						
23	minos	99 %	13,184	76	4116	244
24	alice	98 %	5,155	96	2124	27
25	usatlas	98 %	277,619	5,371	386887	1100
26	sbgrid	97 %	19,384	433	7920	711
28	cms	94 %	489,489	29,627	95680	5900
...						
	All VOs	94 %	1,101,712	69,349	562,470	18,161

Gratia Example #3: Web displays

http://display.grid.iu.edu

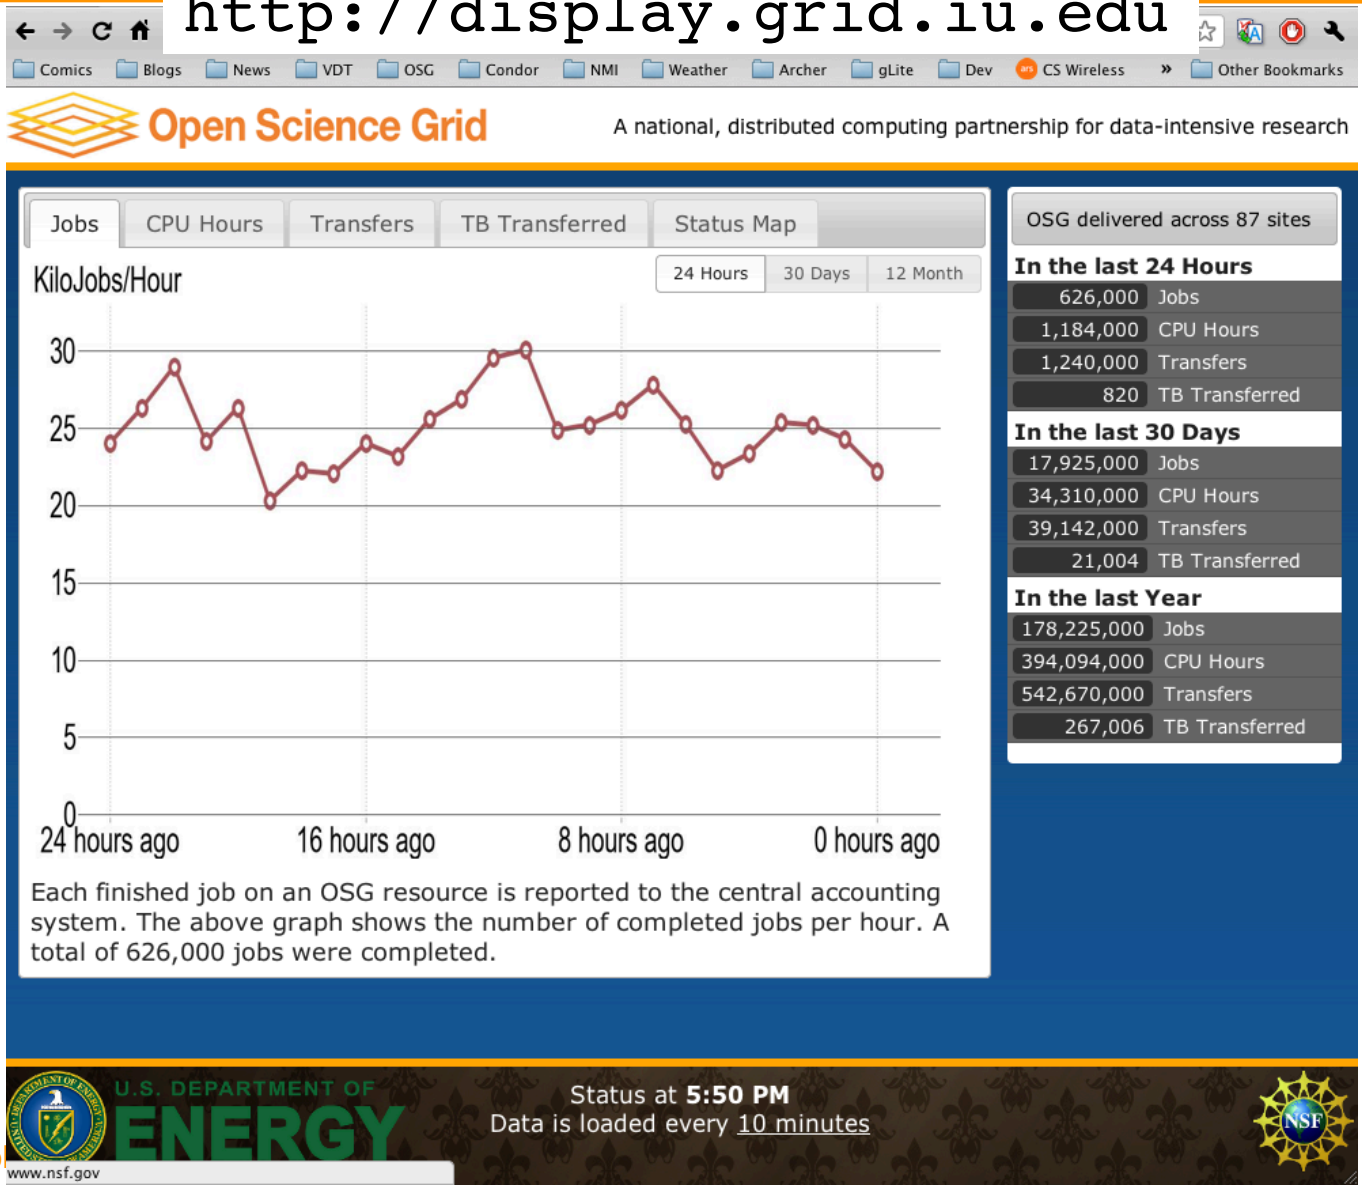

OSG Customization #2: Job Environment

- Extend Job Manager to extend job environment
  - Based on configuration files, automatically created but extensible by users
  - Provides "OSG environment"
    - Where to find installed applications
    - Where to find temporary disk space
    - Where to find local Squid server
    - …

OSG Customization #3: New Jobmanagers

- Condor "NFS-Lite" Jobmanager
  - Doesn't assume that input/output are in NFS
  - Uses Condor's file-transfer mechanism
- Managed Fork Job Manager
  - Allows site to control fork jobs and prevent them from overwhelming gatekeeper
  - Limit how many jobs run at a time based on number (e.g. 10) or policy (as long as load is small).
  - Uses Condor to manage the fork jobs

OSG Customization #4: GridFTP DSI Plugins

- Plugin to allow GridFTP to access files in Xrootd storage system via FUSE
  - Provided by Xrootd developers
- Plugin to allow GridFTP to access files in HadoopFS
  - Provided by CMS (High-Energy Physics experiment from LHC at CERN)

Individual Customizations #1
Access advanced batch system policies

- Assign Condor “accounting groups” or PBS queues based on job DN or VO
- Extend Condor job ClassAd with extra information (User, VO)
- Set max run time, max idle time, etc.
- Tweak allowed architectures for job to run (32- vs. 64-bit)
- Provide job wrapper script to create temporary directory (OSG_WN_TMP)
- Allow MPI jobs to work with local MPI implementation

Individual Customizations #2
New Jobmanagers

- Resubmit job to internal grid with Condor-G (Fermilab)
- Split “SAM” jobs into individual jobs and resubmit across OSG. (Fermilab)
- Condor NFS lite (mentioned earlier) was developed by a user and donated to OSG.

Questions? Comments?

- I'm here until the end of Tuesday: feel free to talk to me any time.
- Email me at: roy@cs.wisc.edu

Dr Charaka J. Palansuriya (EPCC)

Globus World 2011

Argonne National Laboratory

11 April 2011

OGSA-DAI STATUS REPORT

Outline

- What is OGSA-DAI?
- Users
- Current Status
- Future Plans
- Conclusions
- Further Information

What is OGSA-DAI?

- Framework for distributed data access and management
  - Access and update
  - Distributed Query Processing (DQP)
  - Workflow executor
  - Data streaming and concurrent execution
  - 100% Java Open Source
- Distributed heterogeneous data resources
  - Relational, XML, file, RDF
- Extensible framework
- WS built with different toolkits
  - GT, Axis, CXF

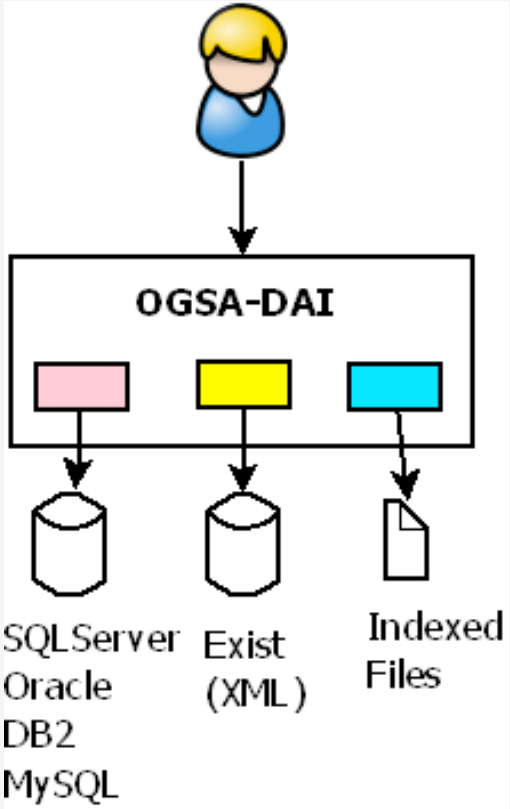

OGSA-DAI components

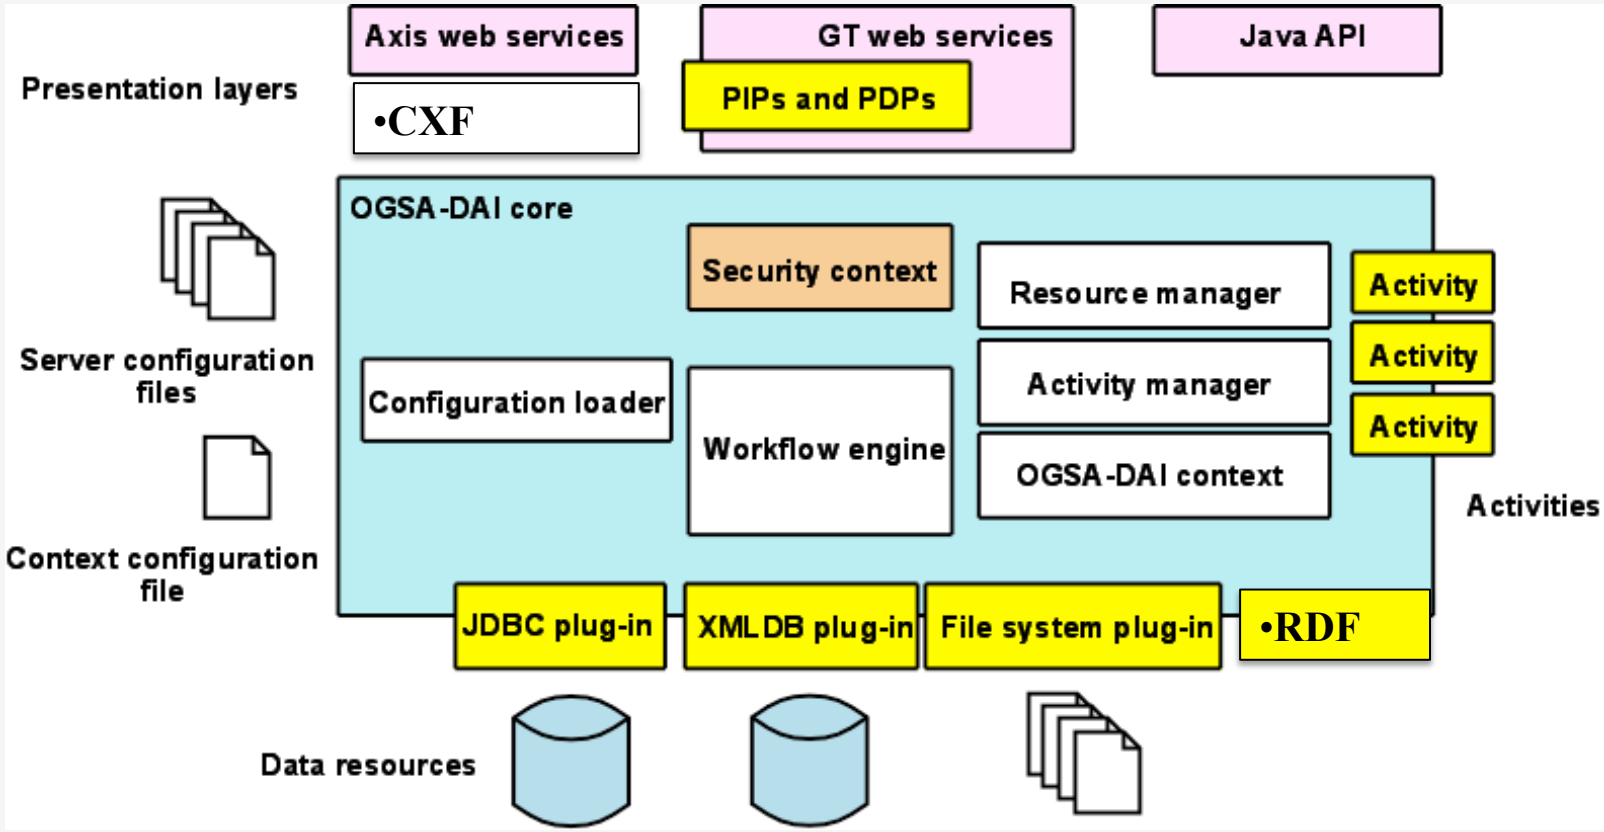

Users

LAQVAT

ADMIRE

omii europe
open middleware infrastructure institute

SEE GEO

GEOGrid

Astro Grid

D-GRID

ACGT
Advancing Clinico Genomic Trials on Cancer

GridMiner
Intelligent Grid Solutions

GIR

eDiaMoND

SIMDAT

BRIDGES

Database Grid

NGS
National Grid Service

Recent Users

•Many and varied!

•IS

•IGE

BIRN Biomedical Informatics Research Network

•2008 2009 2010 2011

OGSA-DAI in Astronomy

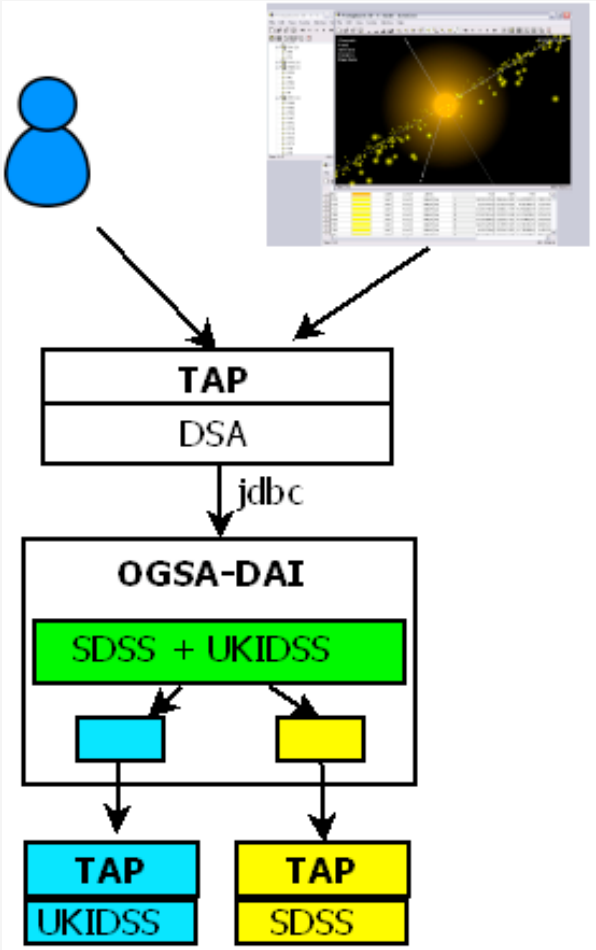

Current Status?

- Version 4.1 was released (March 2011)
  - Efficient binary transfers
  - RDF data resources
  - Other improvements and bug fixes
- Adding CXF presentation layer
  - Initial work GoSC 2010
  - Developing as part of IGE WP4
    - Plan to merge this with new Globus WS Core when available
  - A technology preview is available from IGE
- ADMIRE:
  - Ongoing improvements to DQP
  - Further support for RDF data resources

Future Plans

- Fully functioning CXF web services layer
  - Provide all functionality available via Axis
  - This will eventually replace the Axis 1.4 layer
- Integration with new Globus Java WS Core
  - Globus Crux Toolkit?
- Further support for RDF data resources
  - W3C SPARQL 1.1 Federation Extensions
- DQP improvements
- REST based presentation layer

Conclusions

- OGSA-DAI continues to evolve
  - Widely used internationally and has a high profile
  - Provides a proven technology for managing and integrating distributed heterogeneous data sources
    - Since 2002
- OGSA-DAI is an Open Source project
  - Available from SourceForge
  - Contributors are welcome
- Edinburgh will continue to evolve OGSA-DAI
  - Continuing relationship with Globus
- IGE and other projects will provide necessary support

Further Information & Acknowledgement

- Project web site:
  - http://www.ogsadai.org.uk
- Product SourceForge location:
  - Top page:
  - http://sourceforge.net/projects/ogsa-dai/
  - TRAC wiki:
  - http://sourceforge.net/apps/trac/ogsa-dai/
- Further information:
  - info@ogsadai.org.uk
- OGSA-DAI current and recent funders
  - Initiative for Globus in Europe (IGE)
    - Currently funding CXF work and integration with Globus WS Core
  - ADMIRE
    - Currently funding DQP and RDF extensions
  - The University of Edinburgh
  - Euro-VO AIDA (Royal Observatory of Edinburgh)
  - Google Summer of Code

Alexander Papaspyrou, Dissemination Manager, IGE

GlobusWORLD 2011, April 11-13, 2011, Argonne (IL), United States

GLOBUS IN EUROPE

ADOPTION IN THE OVERSEAS DCIS

Why Globus in Europe?

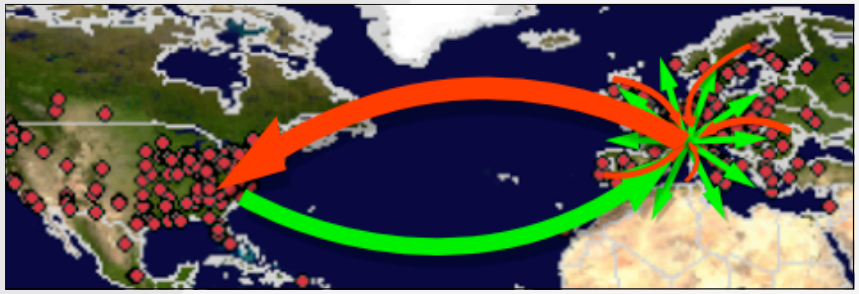

- interfaces
- support
- certification
- privacy
- packaging
- security
- training

A Brief Introduction to the European DCI Ecosystem...

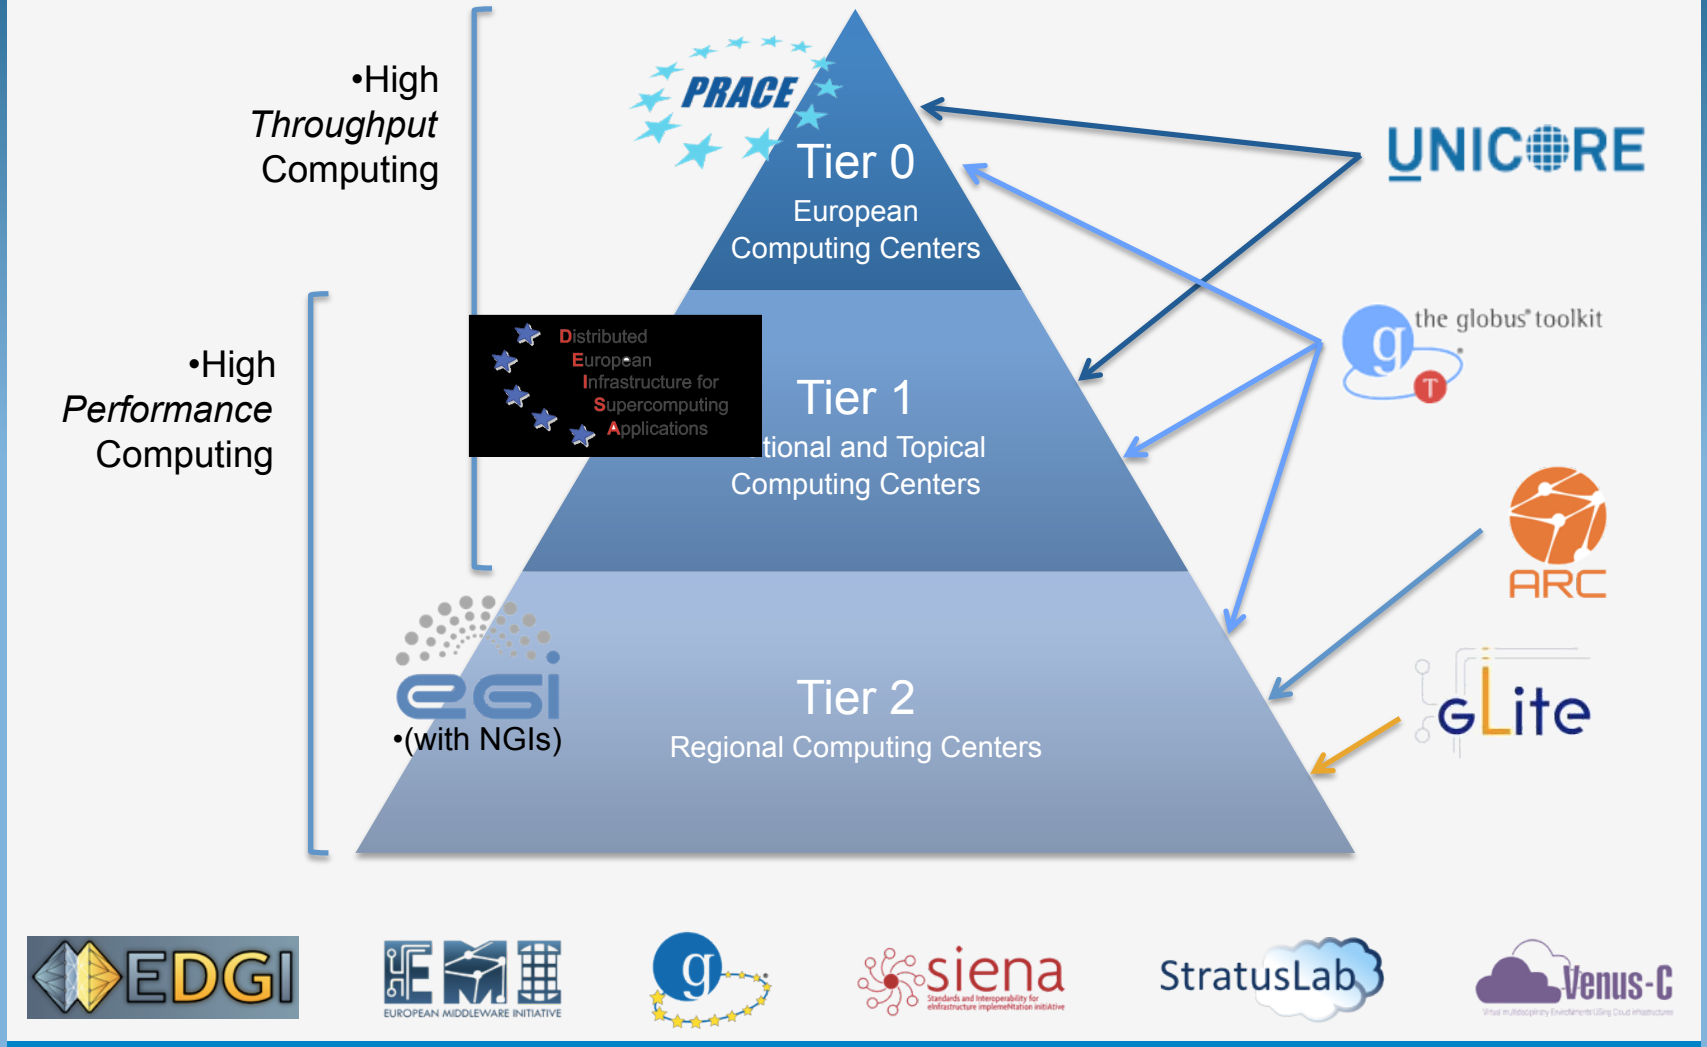

So, what is the role of IGE in here?

•Help the scientists in the European Research Area!

egcf

•*Coordinate*

Bring together
•European Users,
•Developers, and Operators

egh

•*Streamline*

•Provide a
•Single Point of
•Information for Europe

eglo

•*Voice*

•Add the
European Perspective
•to future developments

What we offer as of today.

- European counterparts of Globus resources
  - s and integration testbed
  - cy!)
  - structures
- Development
  - Patches for L
  - Support for additional interfa
  - Integration of requested sec

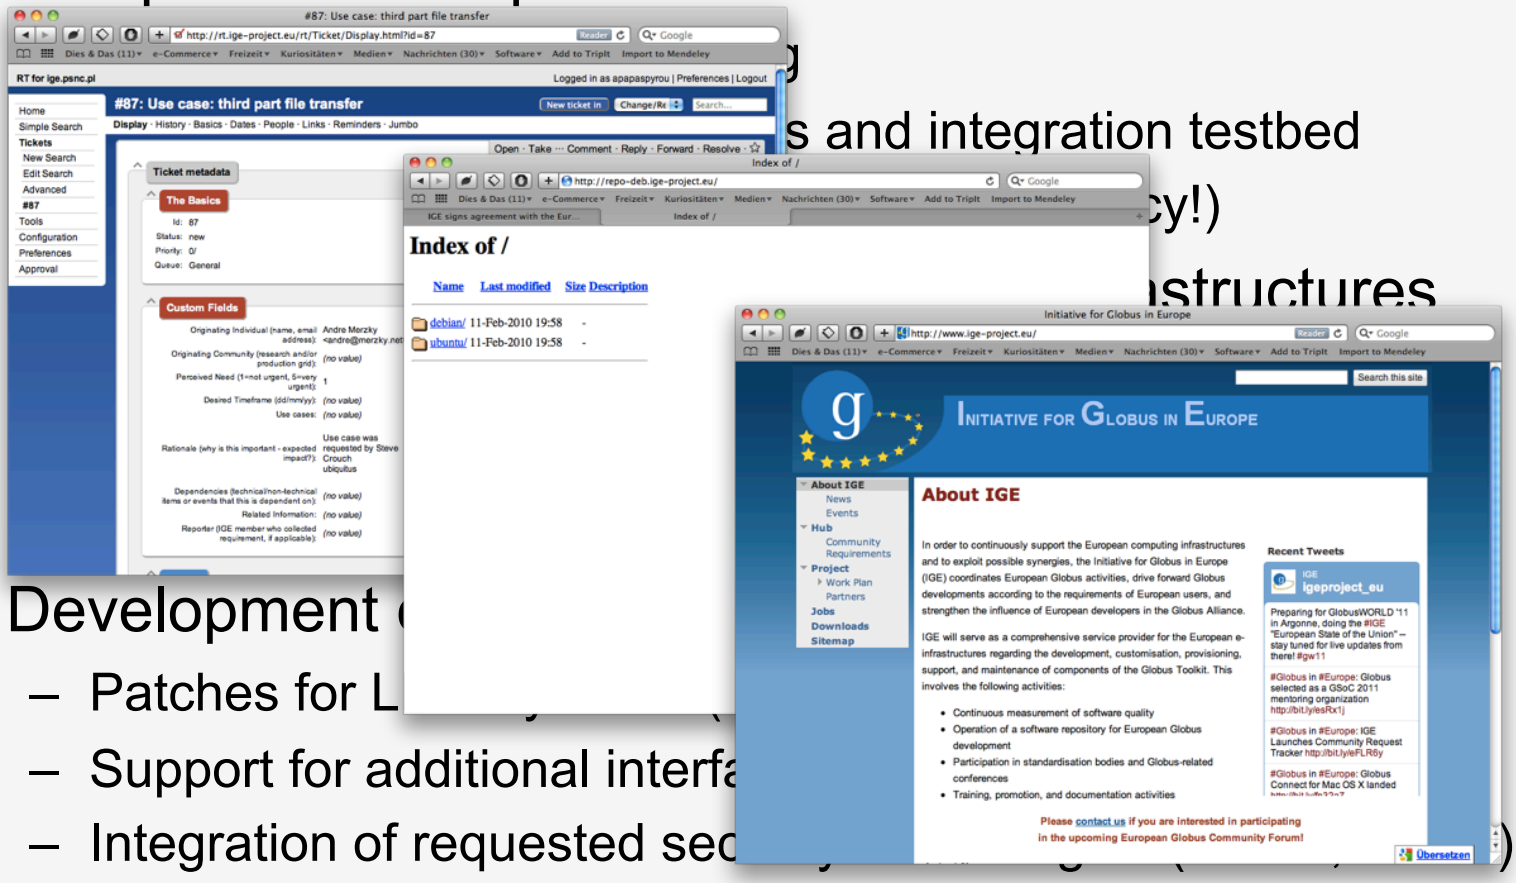

Success stories so far.

- The DEISA / VPH problem
  - Virtual Physiological Human Community
    - Uses SAGA based software
    - Expects GT5 on the Resources
  - Distributed European Infrastructure for Supercomputing Applications
    - Mostly IBM brand machines (running LoadLeveler)
    - IBM does not support GT5
- IGE's success
  - Connect people
    - Rob van der Waal (SARA, Netherlands) contacted us
    - Eduardo Huedo (UCM, Spain) offered expertise with GT and LL
  - Provide solution
    - Joint venture in porting (SARA: testbed, UCM: adapter)
    - IGE takes care of maintenance, packaging, and licensing with IBM

- *and...* This was also useful for TeraGrid, resulting in a use case for interoperation in the PGI-WG within Open Grid Forum: http://go.ige-project.eu/pgi-usecase

Success stories so far.

- The GridKa Summer School 2010 problem
  - On
  - Wa
- IGE's
  - Kn
    -
    -
  - De
    -
    -

Evaluation talks and tutorials/workshops

- Talks:
  - Grid and Clouds, a Look ahead (Fabrizio Gagliardi): 8.9
  - Introduction to Grid and Cloud Computing (Tony Cass): 8.8
  - Grid Computing for LHC (Johannes Elmsheuser): 8.5
  - Grid Security (John White): 8.5
  - Virtualization (Ulrich Schwickerath): 8.4
- Best workshops/tutorials:
  - Globus (Laitinen, Zrenner): 9.4
  - Cloud Computing (Mauch and team): 9.3
  - gLite Administration (Freitag, Feldhaus): 9.3
  - GAT (Beck-Ratzka): 9.3
  - Unicore (Breu, Demuth): 9.3
- All courses had a rating of 8.0 or better
- All talks but one had a rating of 7.8 or better

4 Dr. Christopher Jung - GridKa School 2010

•And… „best of show“!

Success stories so far.

- The Production Infrastructure Adoption problem
  - Formal pr
  - Need for
  - ... But UC
- IGE's solu
  - Liaise
    - IGE a
    - EGI w
  - Take res
    - IGE represents Globus in all relevant bodies
      - Within EGI (Software Vulnerability, Support, Packaging & Release, etc.)
      - Within SEE (Standardization and Interoperability)
    - EGI has a formally established MoU and SLA to Globus towards the Production Infrastructure via EGI.

- *And I hasten to add:*
- **Globus is the first middleware with a formal MoU and SLA**

Next steps.

- Statistics collector,
  - Currently running internally, for testing
  - Integration into IGE packages
  - Will arrive within Q2/11
- Training Portal
  - To feature slides, hands-on material, videos, FAQs, ...
  - Part of the European Globus Hub effort
  - Will arrive end of 2011
- Contiuing development cycles
  - PDCA style, every six months
  - Catering of standardisation needs (with SIENA)
  - Globus “appliances” (with StratusLab)

Thank you for listening.

•Follow us on

•http://www.ige-project.eu/news

twitter •@igeproject_eu

LinkedIn •http://go.ige-project.eu/linkedin

1st EGCF Meeting

May 18-19, 2011

Leibniz Supercomputing Center

Garching near Munich, Germany

GlobusEUROPE

September 19, 2011

EGI Technical Forum

Lyon, France

National Cancer Institute

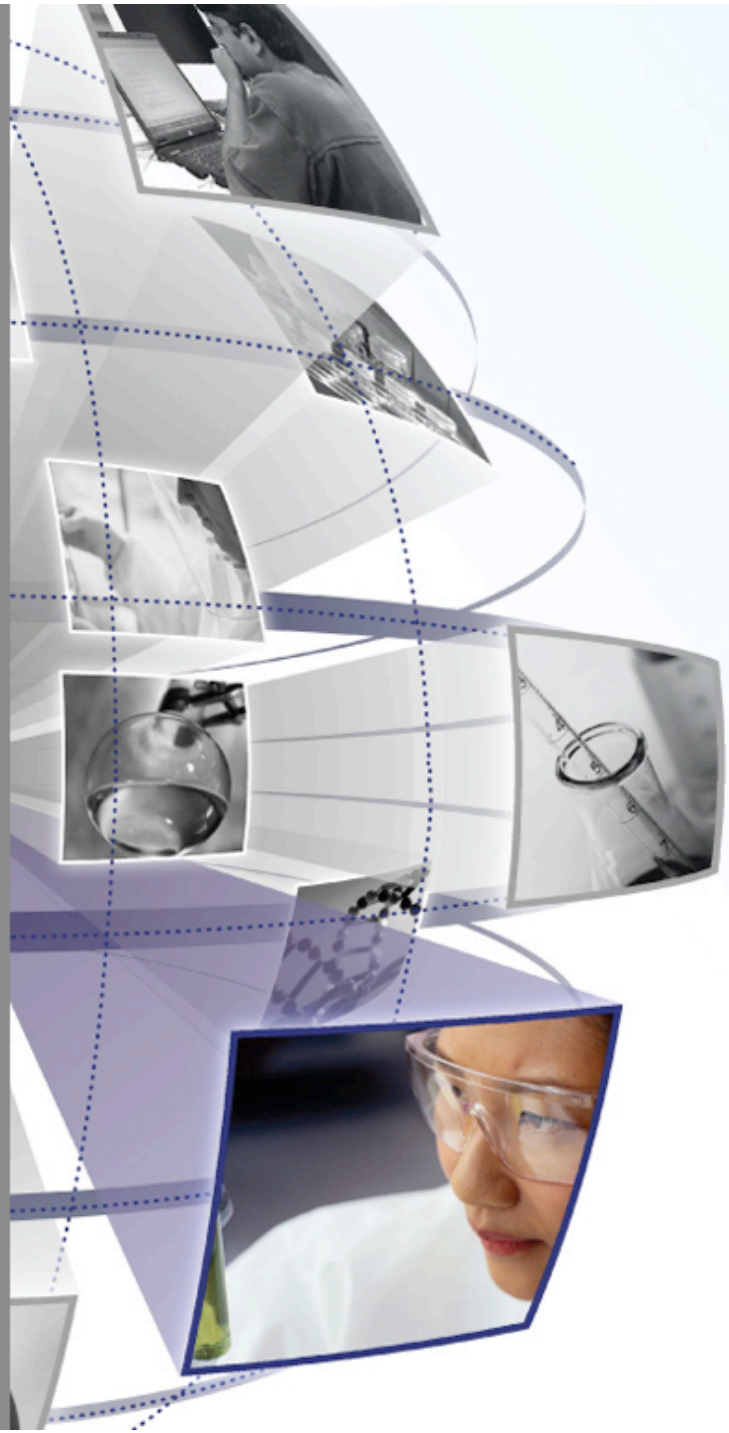

caBIG™
cancer Biomedical Informatics Grid ™

caBIG Update

Ravi K Madduri

University of Chicago and Argonne National Laboratory

Outline

- caBIG Workflow update
- caGrid 2.0 Roadmap and prototyping
- BSA Report
- Future Directions
- Q & A

Technical Accomplishments

- **Continued work on enhancing Taverna Workbench**
- **Integrated Taverna Credential Manager with caBIG GAARDS infrastructure**
  - Added ability to invoke https services
  - Integrated user login to Dorian from Taverna
  - Added ability to delegate users' credential using CDS
- **Enhancements to CQL Builder**

Technical Accomplishments

- **Integrated caGrid Transfer functionality for Bulk data transfer**
- **Workflow Portlet**
  - Reduce barrier-to-entry for end users
  - Ability to discover workflows from myExperiment and execute them from the portal

Accomplishments

- **Lymphoma workflow – Among the top 20 most viewed/downloaded Workflows in myExperiment**
  - This is more impressive given that this workflow was uploaded much later than the other workflows
- **Our BMC-Bioinformatics Article on "caGrid Workflow Toolkit: A Taverna based workflow tool for cancer Grid" achieved "Highly Accessed" relative to its age**
- **We are part of the CVRG Project that recently got renewed**

caGrid 2.0 Prototyping

- **Semantic Workflows**
  - Worked on creating a prototype of Semantic Workflow Engine using SADI plugin in Taverna
  - User can start with a data set or a service and the engine “recommends” next steps in the workflow using Semantic Metadata
  - Joint work with colleagues from NIU (Special Thanks to Jia Zhang and her students)
  - Results in providing guidance on how capabilities can be annotated so they can be discovered automatically
  - Leverages the power of Semantic Web and Linked Data Concepts

caGrid 2.0 Security Prototyping

- **Platform Security and Tooling Team**
  - Goals
    - Ensure interoperability across programming models
    - Ensure interoperability across platforms, Java and .NET
    - Assess tools available for use and development efforts for production system
  - Deliverables
    - Prototyping Authentication Model for REST and SOAP
    - Prototype Authorization Model
    - Prototype Delegation Model
- **Enterprise Security Program**
  - Goals
    - Investigate support for LOA3 and LOA4 credentials.
  - Deliverables
    - Prototype a WebSSO model for supporting LOA3 and LOA4 credentials

NCI-BSA Report on caBIG

- Institute an immediate moratorium on all ongoing internal and commercial contractor- based software development projects while initiating a mitigation plan to lessen the impact of this moratorium on the cancer research community.
- Institute a one-year moratorium on initiation of new projects, contracts and subcontracts by caBIG®.
- Provide a one-year extension on caBIG®-supported academic efforts for development, dissemination, and maintenance of new and existing community-developed software tools.
- Establish an independent oversight committee, representing academic, industrial, and government (NCI, NIH) perspectives to review ongoing and planned initiatives for scientific merit and to recommend effective transition options for current users of caBIG® tools
- Conduct a thorough audit of all aspects of the caBIG® budget and expeditures.

Lessons Learned

- Service oriented infrastructures should realize that Return on Investment is in creating reusable workflows
- Workflows are only as good as the services that are available
- Lowering the barriers to entry for sharing data and analytics
- Software is surprisingly hard to use for end users – more so if the benefit is not all too clear

Future Directions

- **The future of caBIG program is uncertain – Lot of good work has been done that can be leveraged**
- **We are focusing on working with end users to get requirements**
- **Planning to investigate hosted Workflow Solution– SaaS workflow tools**
  - Galaxy Prototype
  - Globus Online
  - Taverna Server Prototype
  - HTC Integration

Acknowledgments

- **UChicago**
  - Ian Foster, Dina Sulakhe, Wei Tan (Now at IBM), Bo Liu, Mark Kedzierski, Mike Russell
- **University of Manchester**
  - Stian Soiland Reyes, Alex Nenadic, Carole Goble


MADDURI, RAVI K. 3/23/1978 33 E1853286 4/10/2011 XR ELBOW (MIN 3 VIEWS) LEFT XRAY; SCALETTA THOMAS A.
PLAINFIELD, By: Alan Nazeran, MD