GlobusWORLD 2011: Globus Toolkit Updates

Jim Basney
jbasney@ncsa.uiuc.edu

This material is based upon work supported by the National Science Foundation under grant numbers 0850557 and 0943633. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Outline

- CILogon Project
- Core Globus Toolkit Components:
  - MyProxy
  - GSI-OpenSSH
- GridShib
- CILogon Service and InCommon

CILogon Project Goals

- Foster secure, usable authentication for cyberinfrastructure (CI)
- Provide community-driven **software** support:
  - MyProxy (http://myproxy.ncsa.uiuc.edu)
  - GSI-OpenSSH (http://grid.ncsa.uiuc.edu/ssh)
  - GridShib (http://gridshib.globus.org)
- Provide a new **service** (https://cilogon.org/) that issues certificates to researchers for authentication to cyberinfrastructure

Core GT Components

- MyProxy
  - Certificate Authority and Credential Repository
  - Authentication methods: password, certificate, PAM, OTP, LDAP, SASL, Kerberos, VOMS
  - Client APIs: Java, Python, Perl, PHP, JAAS
  - http://myproxy.ncsa.uiuc.edu
- GSI-OpenSSH
  - GSI support for OpenSSH
  - Single sign-on remote login and file transfer
  - High Performance Networking (HPN) updates (PSC)
  - Java GSI-SSHTerm client
  - http://grid.ncsa.uiuc.edu/ssh

MyProxy News

- New releases (bug fixes, minor features)
  - GT 5.0.3 contains MyProxy 5.3 (latest release)
  - http://myproxy.ncsa.uiuc.edu/dl/VERSION
- Updates for jGlobus 2.0
  - MyProxy jGlobus API unchanged
  - http://myproxy.ncsa.uiuc.edu/jglobus
- Fedora/EPEL RPMs available
  - Thanks to Steve Traylen (CERN)
  - https://admin.fedoraproject.org/updates/myproxy
- GT 5.2 RPMs in preparation

GSI-OpenSSH News

- GSI-OpenSSH 5.2 (latest release)
  - Included in GT 5.0.3
  - OpenSSH 5.6p1 and HPN13v6
- GSI-OpenSSH 5.3 (coming soon)
  - OpenSSH 5.8p1 and HPN13v11
- Fedora/EPEL RPMs in preparation
  - Thanks to Mattias Ellert
    - https://bugzilla.redhat.com/show_bug.cgi?id=683587
- GT 5.2 RPMs in preparation

GridShib

- GridShib-CA
  - Online CA supporting SAML/Shibboleth and OpenID authentication
  - Used by CILogon Service (https://cilogon.org)
  - http://gridshibca.cilogon.org/
- GridShib SAML Tools
  - For using SAML assertions with certificates
  - Compatible with GRAM5
  - Used by TeraGrid Science Gateways
    - https://www.teragrid.org/web/science-gateways/security
  - Updated for jGlobus 2.0
  - http://gridshib.globus.org/

CILogon Service

- Facilitate campus logon to CI
  - Use researchers' existing security credentials at their home institution
  - Ease credential management for researchers and CI providers
- Bridge from:
  - Credentials issued by InCommon Federation members using SAML web browser single sign-on
- Bridge to:
  - Certificates that satisfy the requirements of CI projects

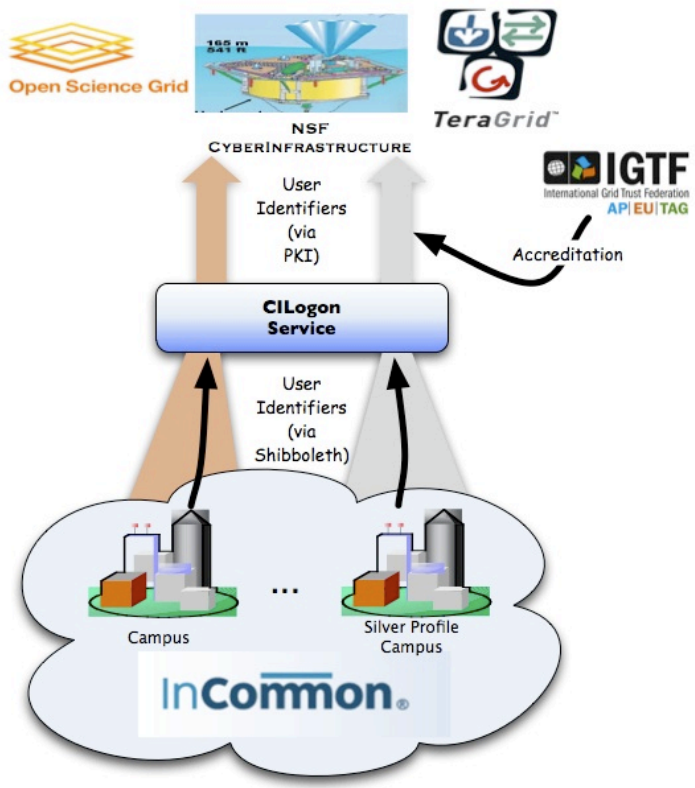

InCommon Federation

- Facilitates use of campus identity with external service providers
  - Supports adoption of standard mechanisms and policies
  - Distributes SAML metadata that identifies members
- Uses SAML Web Browser Single Sign-On protocols
  - Shibboleth implementation from Internet2
  - Works well for browser-based applications
- InCommon represents >200 institutions (>4m users)

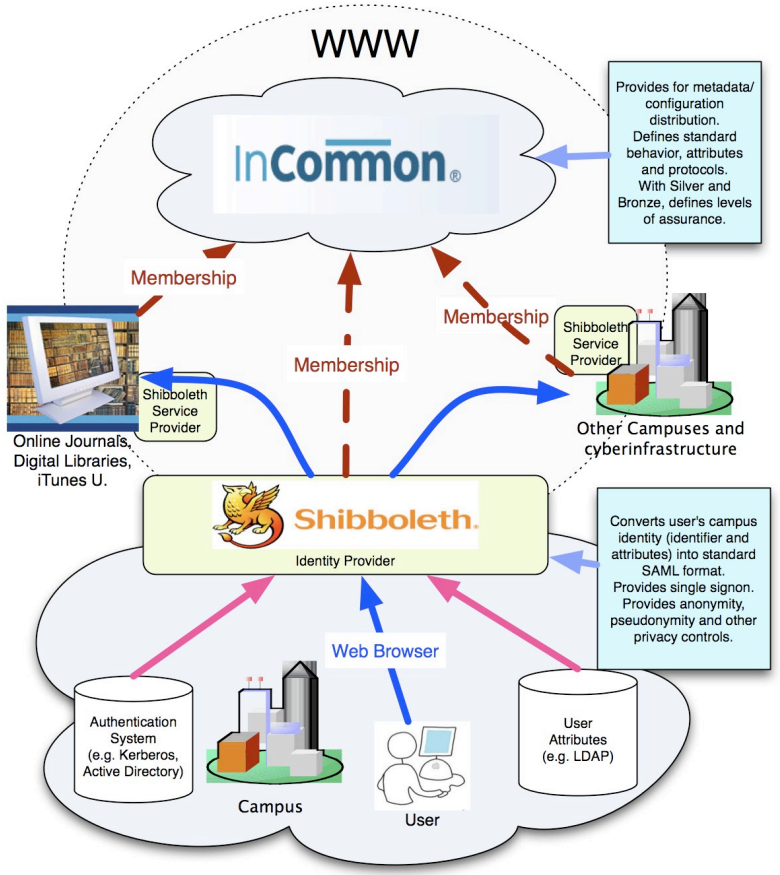

www.incommon.org

CILogon Service

- Delivers certificates to desktop, browser, & portals
- Available certificate lifetimes: from 1 hour to 13 months
- InCommon and OpenID authentication
- Custom appearance and behavior via “skins”
- Available now at https://cilogon.org
- More details at http://www.cilogon.org/faq

Thanks!

jGlobus 2.0

April 11, 2011

CILogon

www.cilogon.org

jGlobus Overview

- **jGlobus 2.0 Beta, released in March 2011, is a fork of the "COG-Jglobus" 1.8 project**
- **jGlobus 2.0 provides Java support for:**
  - GSI-enabling applications and services
  - Client APIs to GRAM
  - Client APIs to GridFTP
- **Developed by the Computation Institute with major contributions from NSCA**
- **Open-sourced under Apache 2 License**

jGlobus Highlights

- **Developed and tested against GT 5.0.2**
- **Packaged with Maven in 3 separate projects:**
  - **jGlobus Core** – GSI Security and Core classes
  - **jGlobus GRAM** – Client support for GRAM
  - **jGlobus GridFTP** – Client support for GridFTP
- **Built on Java Secure Sockets Extension (JSSE) instead of pureTLS (see 1.8.x)**
- **Includes support for SHA2 certificates**

Additional Changes from 1.8.X

- **The following packages were discontinued:**
  - **org.globus.gatekeeper** - jGlobus no longer allows developers to instantiate a GRAM gatekeeper in Java. All GRAM capabilities are provided by the Globus Toolkit.
  - **org.globus.mds** - Support for MDS has been removed as MDS is no longer packaged with the Globus Toolkit.
  - **org.globus.tools.ui** - Java Swing UI for creating proxies have not been carried over into jGlobus 2.0. Our plan is to review what tools are needed for improving the development experience, so ideas or feedback welcome!
- **Several unused classes and packages were deprecated**
- **Java 1.5 or higher now required**

Development and Releasing

- **Code now hosted in GitHub**
- **jGlobus encourages community contributions, developers can submit changes as Git pull requests**
- **JIRA now used to track change requests and bug reports**
- **Binaries releases are published to Maven Central**

Future Development

- **jGlobus 2.0 Beta undergoing joint security review by Computation Institute and NCSA**
- **jGlobus 2.0 production released expected in early Summer 2011**
- **Client support for Globus Online REST ful services will be added in Summer 2011**

References

- http://www.globus.org/toolkit/jglobus
- http://github.com/jglobus
- http://oss.sonatype.org/content/repositories/releases/org/jglobus/
- http://jira.globus.org/browse/CONTACTJGLOBUS
- http://www.globus.org/toolkit

GT 5.2 Native Packaging

Stuart Martin
Software Development Manager
University of Chicago / Argonne National Lab

What is Native Packaging?

- **Debian.org**
  - precompiled software bundled up in a nice format for easy installation on your machine
- **Ian Murdock has commented**
  - that package management is "the single biggest advancement Linux has brought to the industry", that it blurs the boundaries between operating system and applications, and that it makes it "easier to push new innovations [...] into the marketplace and [...] evolve the OS"

Goal of GT Native Packaging

- **Easily integrate Globus software into standard Linux distributions**
  - RedHat, RHEL, Fedora, CentOS, Scientific Linux, Debian, Ubuntu, Mint, Suse, …
- **Embrace native installation methods for other important systems**
  - OS X
  - Windows
- **Continue to provide a source installer for additional systems we don't (yet) support**

Target Packaging Formats

- **Linux**
  - RPM Based
  - Debian Based
- **Mac**
  - PackageMaker / .pkg
- **Windows**
  - Install shield / Windows installer

22

Package Management Systems

- **RPM Based Repos**
  - yum
- **Debian Based Repos**
  - apt

23

Get Accepted into Linux Repositories

- **RPM Based**
  - RedHat, RHEL, Fedora, CentOS, Scientific Linux
- **Debian Based**
  - Debian, Ubuntu

Relevance to other efforts?

- **Mattias Ellert www.grid.tsl.uu.se/repos/**
  - Debian, RHEL, OpenSUSE, Fedora, Ubuntu
- **VDT**
  - RedHat Compatible, Debian
- **GT will provide a development repo and push to Linux repos**
  - Other efforts should be unnecessary

GT 5.0 to 5.2 Changes

- **The concept of "flavors" of Globus libraries has been obsoleted**
  - Every library is built to be threadsafe
    - An application can choose at runtime
- **Compliant with the Filesystem Hierarchy Standard**
  - important prerequisite to being part of Linux distros
- **GLOBUS_LOCATION is no longer required to be set at runtime**
  - Non-standard for Linux distros
  - But, still honored if set
- **Each package has a .spec file (rpm)**

26

GT 5.2 Alpha 1

- **Released Dec 2010**
- **Includes only GridFTP packages**
- **Repos:**
  - Fedora 14, CentOS 5
- **Created a Globus yum repo**
  - www.mcs.anl.gov/~mlink/globusrepo
- **Command Examples:**

```
% yum groupinstall globus-gridftp
% chkconfig globus-gridftp-server on
% yum install globusonline-trust-roots
```

Steps to GT 5.2.0

- **Alpha2 – Mid April**
  - + MyProxy, GSISSH
  - + Fedora 13, Scientific Linux 5, RHEL 5
- **Alpha3 – ~May**
  - + Debian packaging for GT components
  - + Debian, Ubuntu repos
  - + Maybe include GRAM
- **Beta - ~June**
  - + GRAM if not already in Alpha3
- **GT 5.2.0 - Summer**

Thanks to the GT 5.2 Developers!!

- **Joe Bester - ANL**
- **Eric Blau – ANL**
- **Mike Link – UChicago**
- **Lukasz Lacinski - UChicago**
- **Jim Basney – NCSA**
- **Venkat Yekkirala – NCSA**

GRAM5 - Scalable and Reliable

Stuart Martin

Software Development Manager

University of Chicago / Argonne National Lab

What is GRAM?

- **GRAM is a Globus Toolkit component**
  - For Grid *job management*
- **GRAM is a unifying remote interface to Local Resource Managers**
  - Yet preserves local site security/control
- **GRAM provides stateful job control**
  - Reliable create operation
  - Asynchronous monitoring and control
  - Remote credential management
  - Remote file staging and file cleanup

Grid Job Management Goals

Provide a service to securely:

- **Create an environment for a job**
- **Stage files to/from environment**
- **Cause execution of job process(es)**
  - Via various local resource managers
- **Monitor execution**
- **Signal important state changes to client**

Traditional Interaction

- Satisfies many use cases
- TACC's Ranger (62976 cores!) the Costco of HTC ;-), one stop shopping, why do we need more?

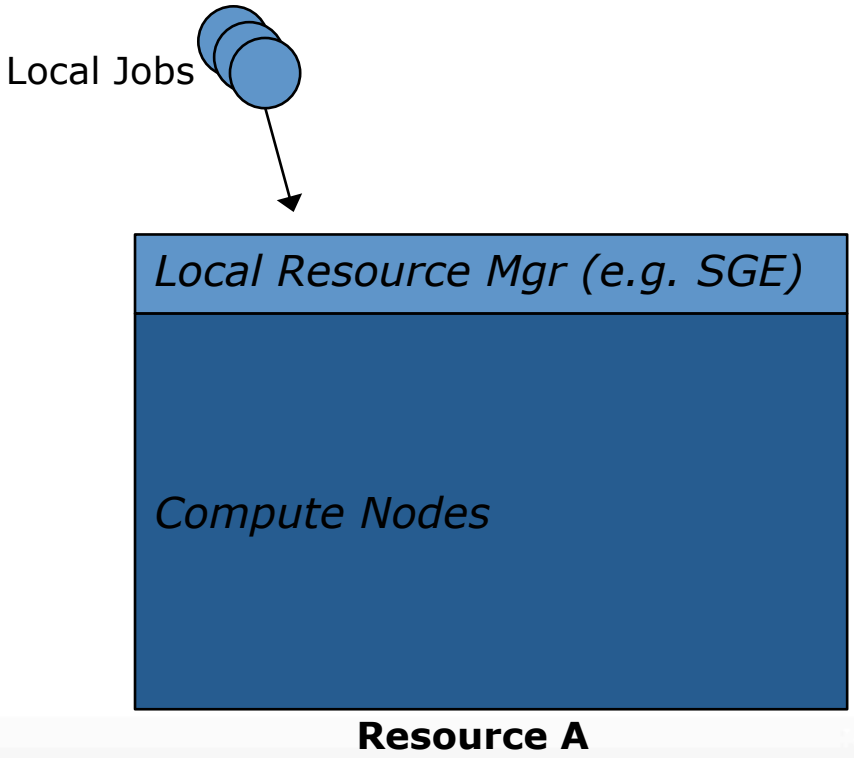

GRAM Benefit

- Add remote execution capability
- Enable clients/devices to manage jobs without logging into the cluster

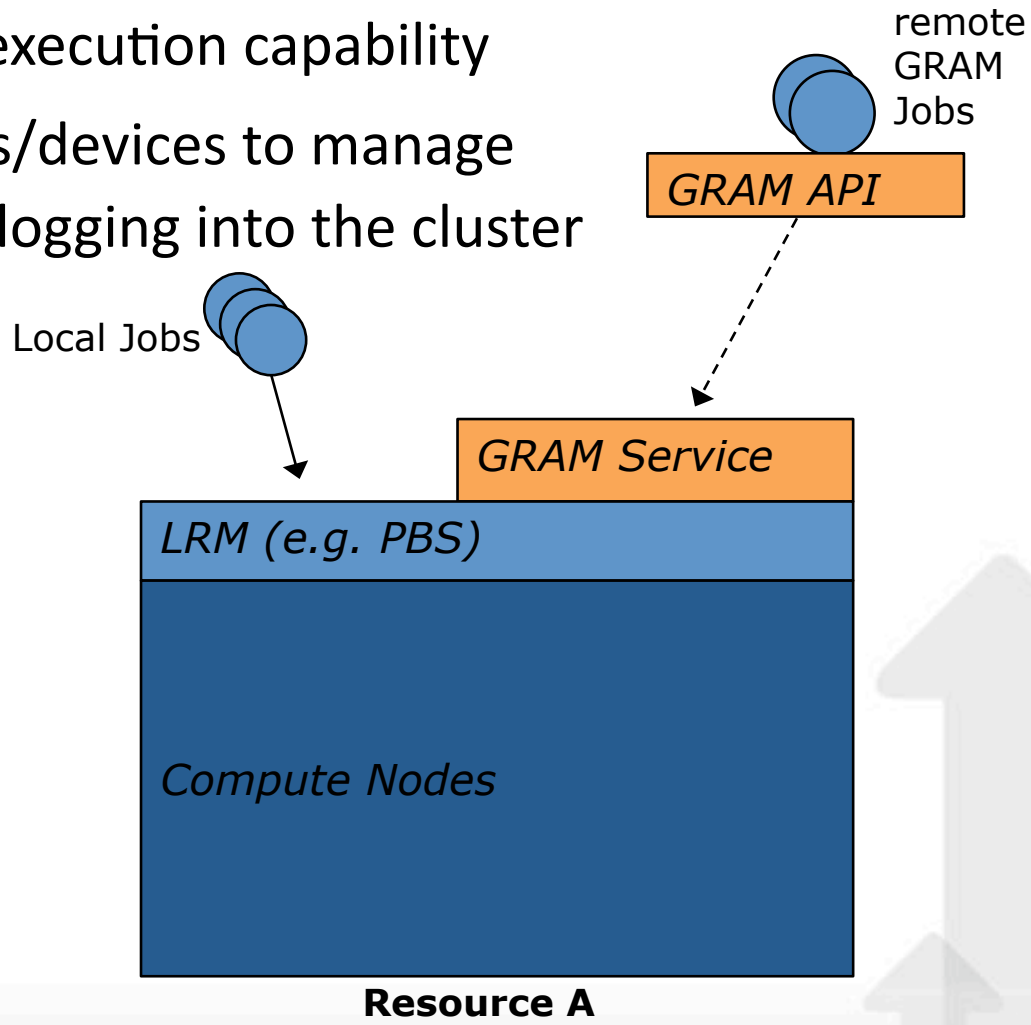

GRAM Benefit

- Provides LRM abstraction

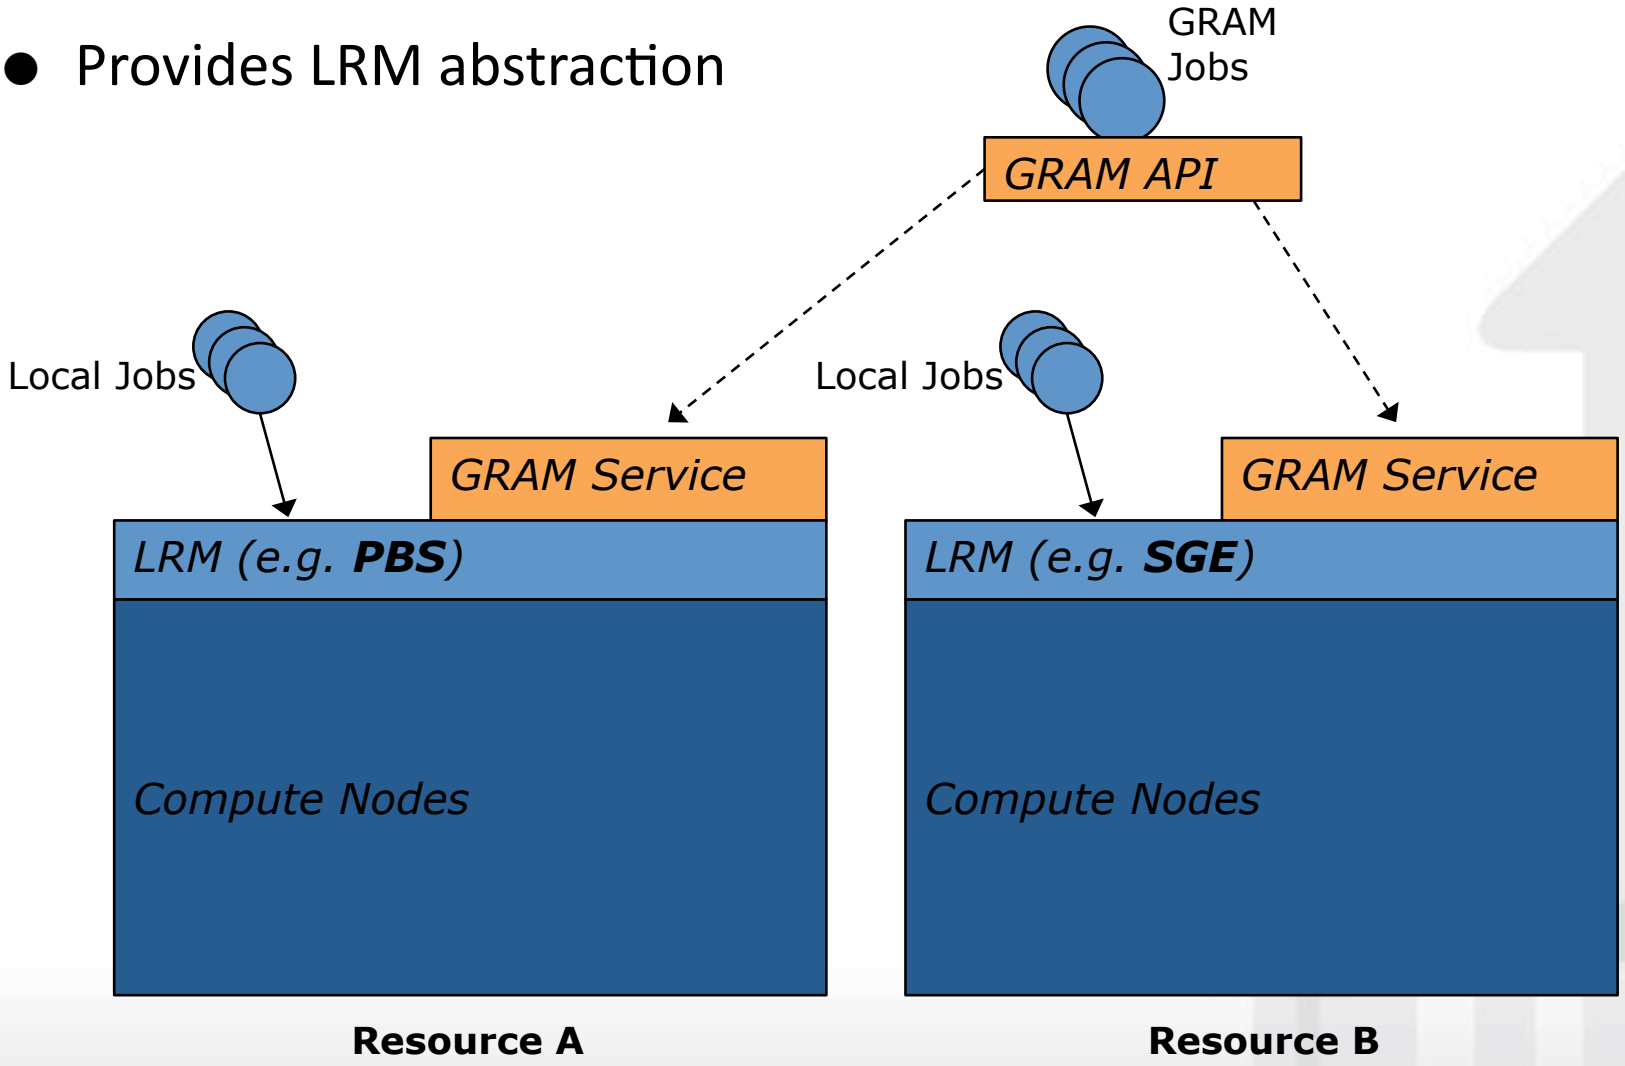

GRAM Benefit

- Scalable job mgmt
- Interoperability

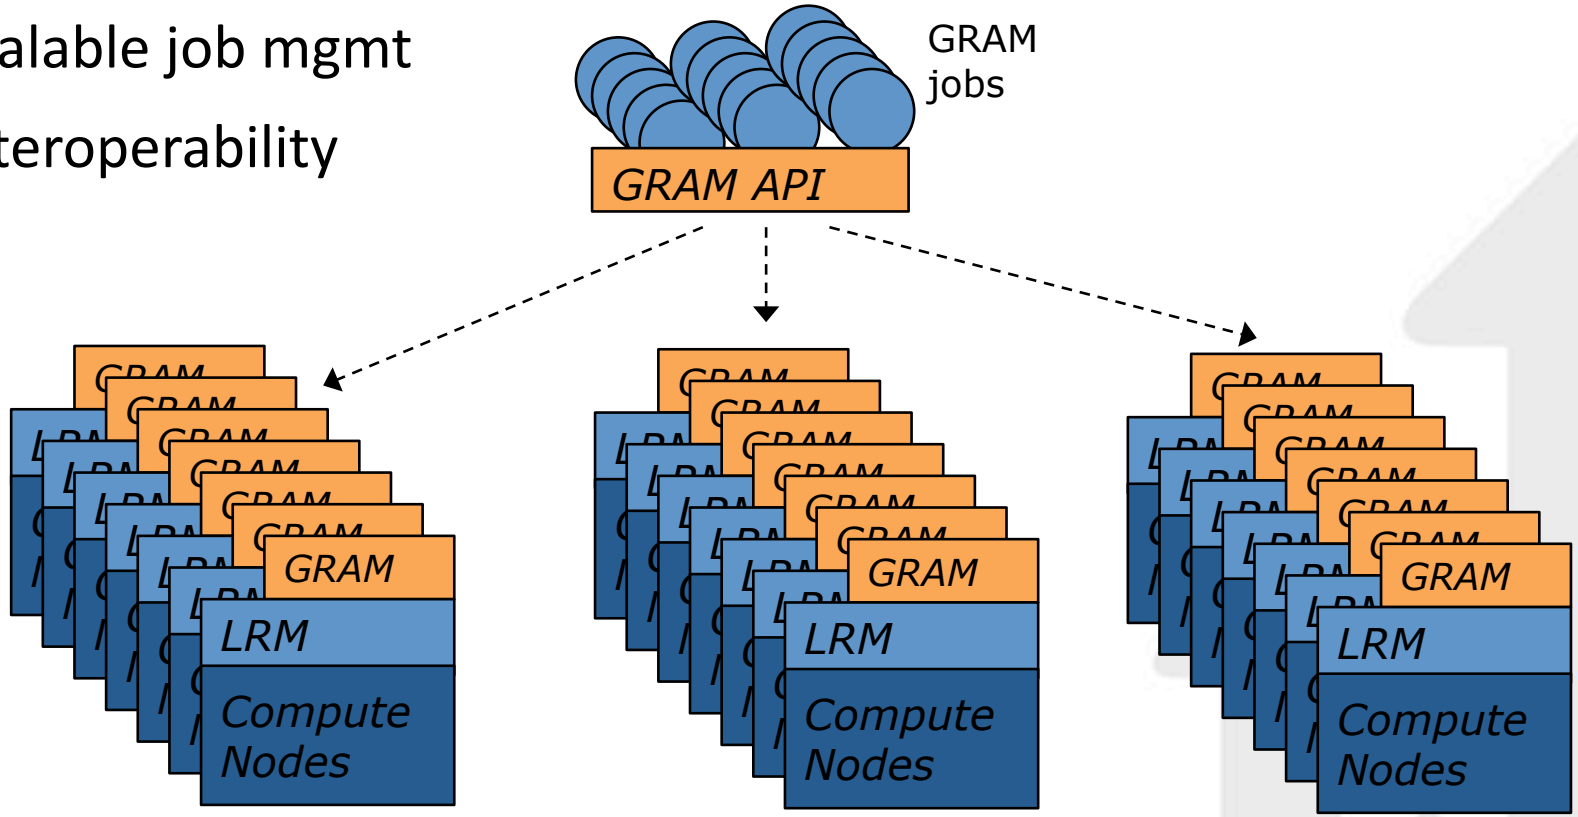

Open Science Grid

Grid Applications
Grid Middleware
Networking
NAREGI

TeraGrid

Protocol Hourglass

Users/Applications:
Science Gateways, Portals, CLI scripts,
App Specific Web Service, etc.

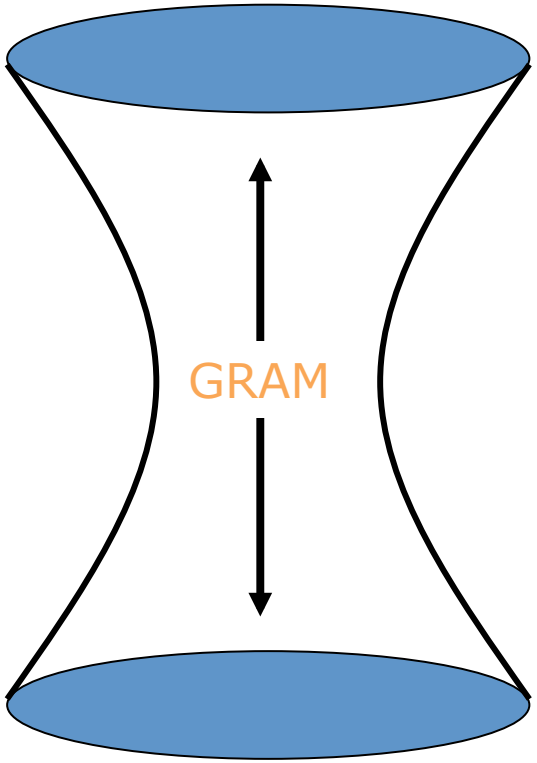

Local Resource Managers:
PBS, Condor, LSF, SGE, Fork

Higher-Level Client Examples

Condor-G Architecture

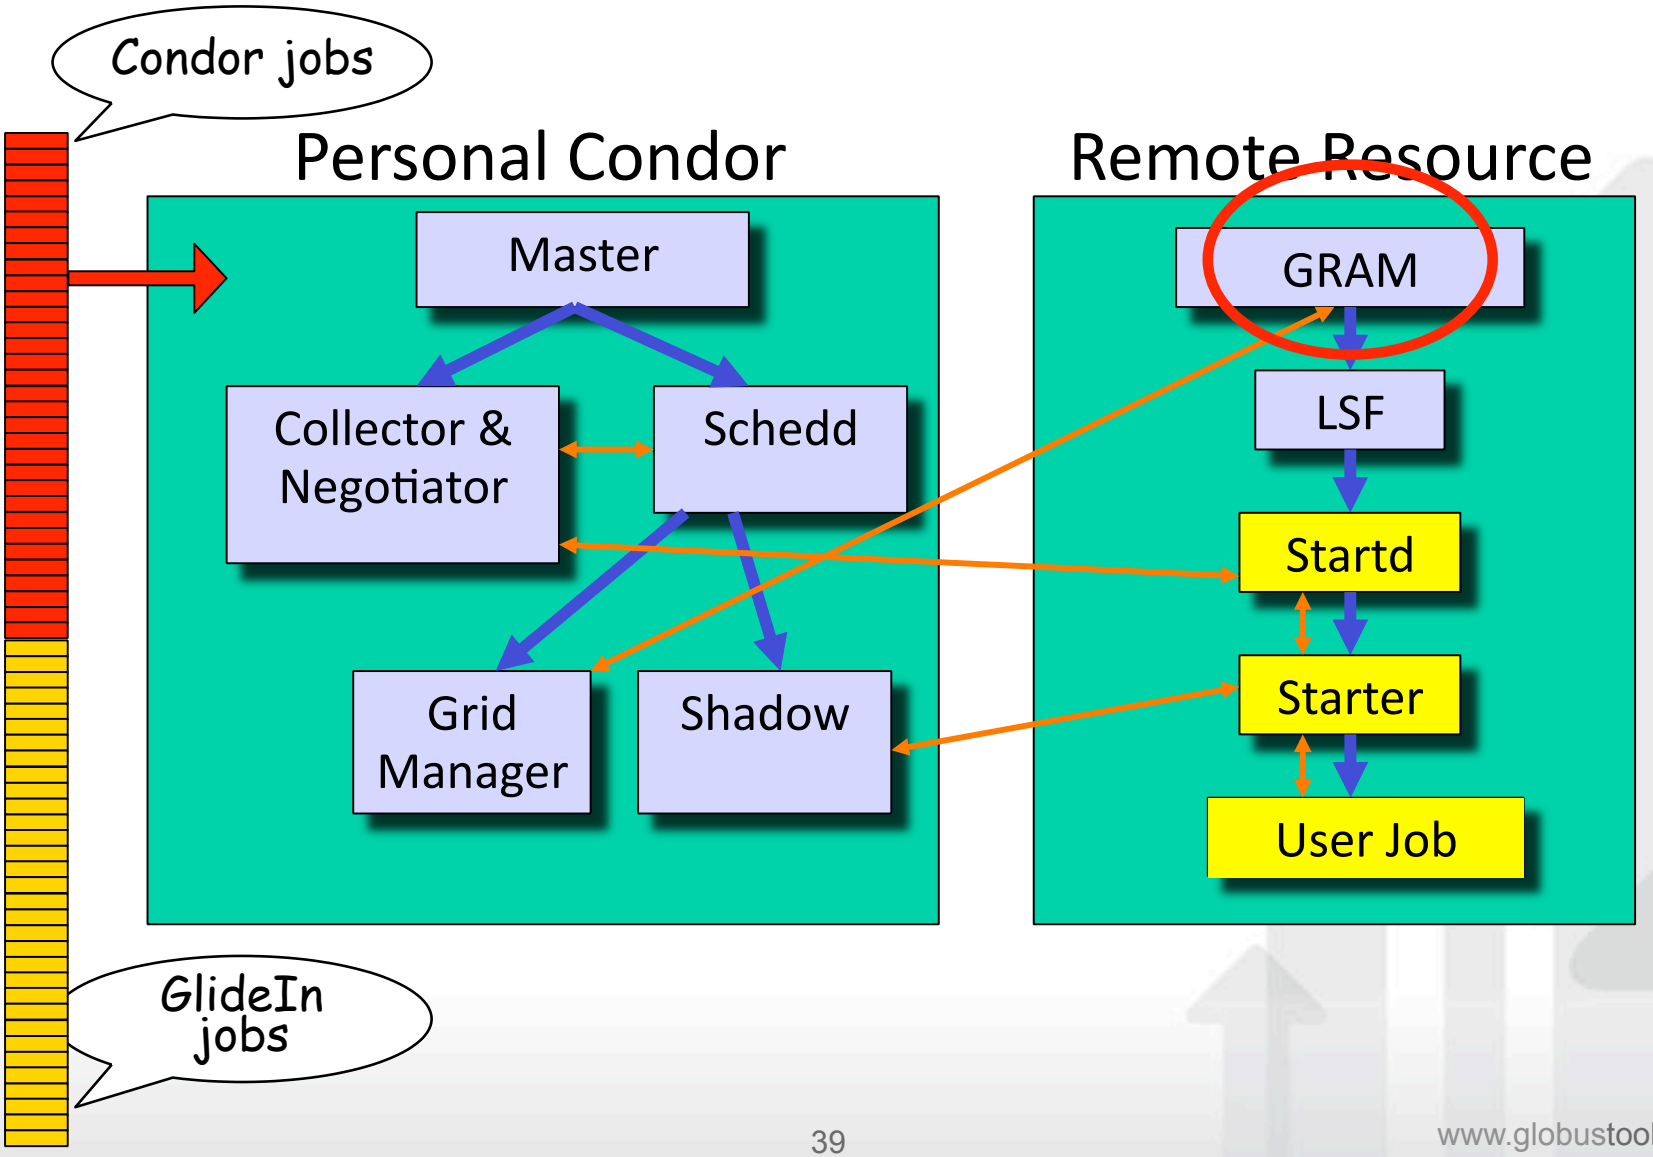

GridWay Architecture

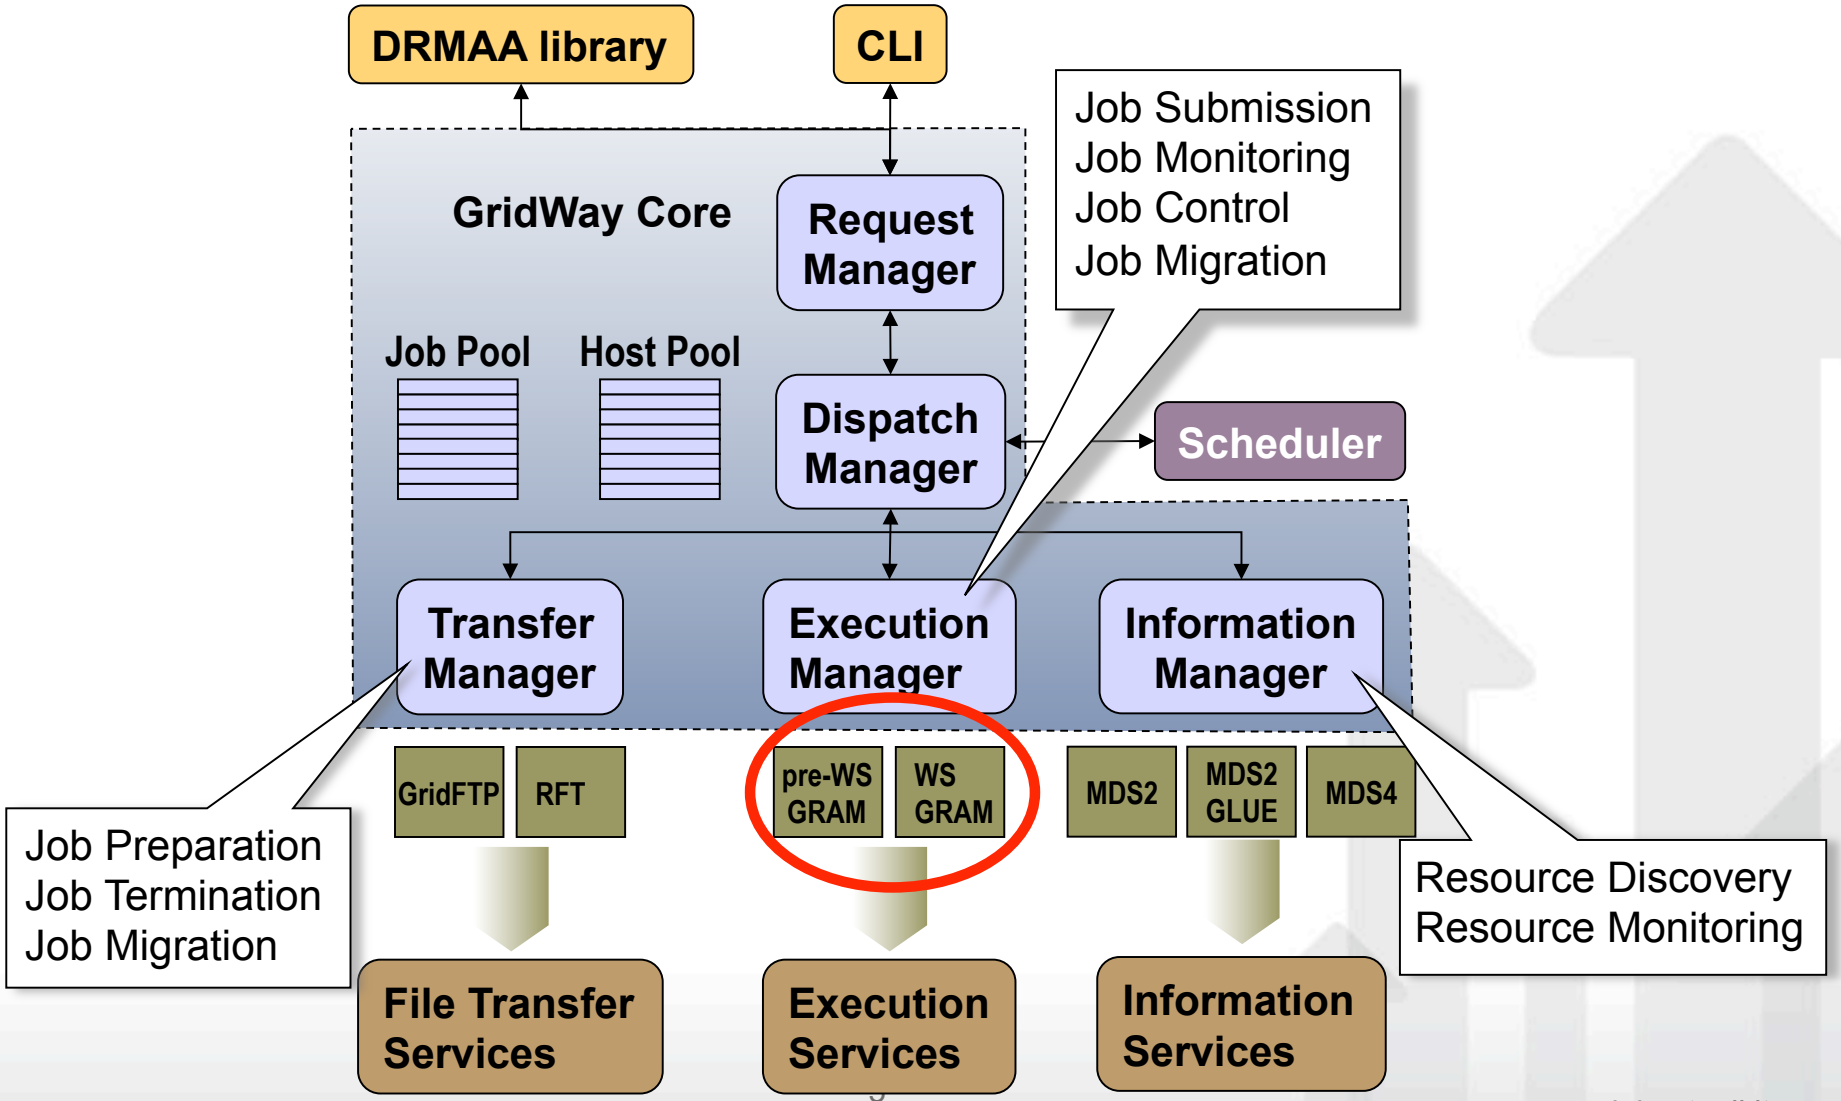

GridWay / Condor-G Benefit

- Scalable job management
- Throttling
- Metascheduling

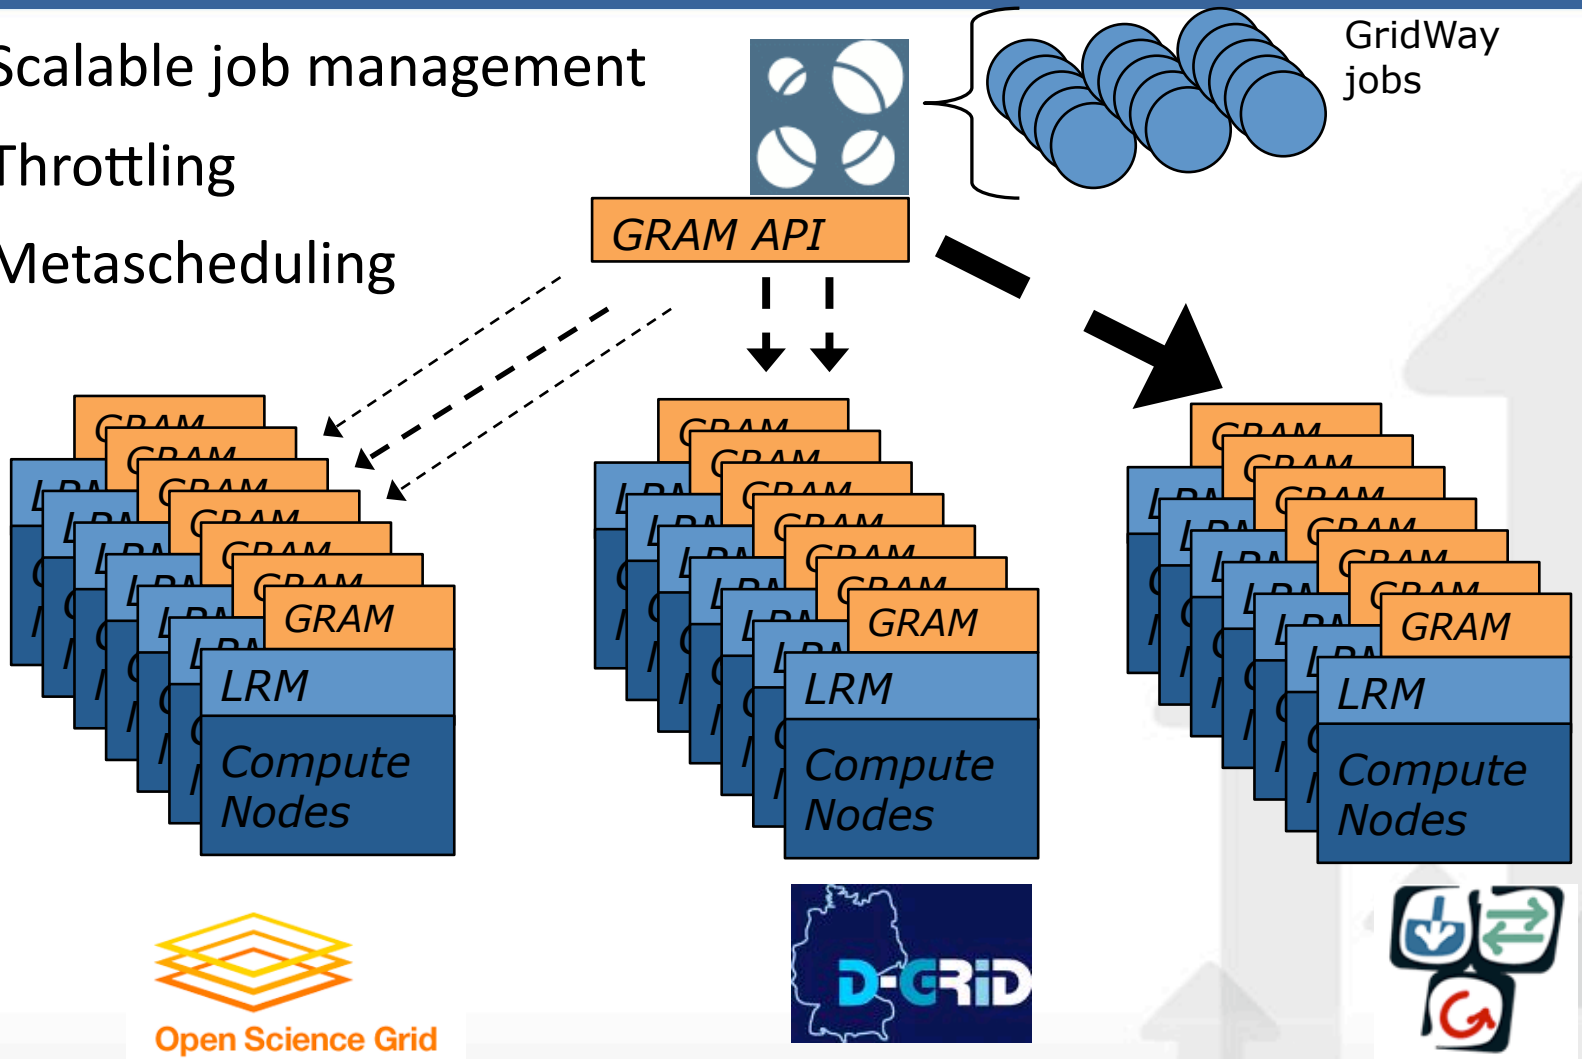

Ninf-G (grid-rcp)

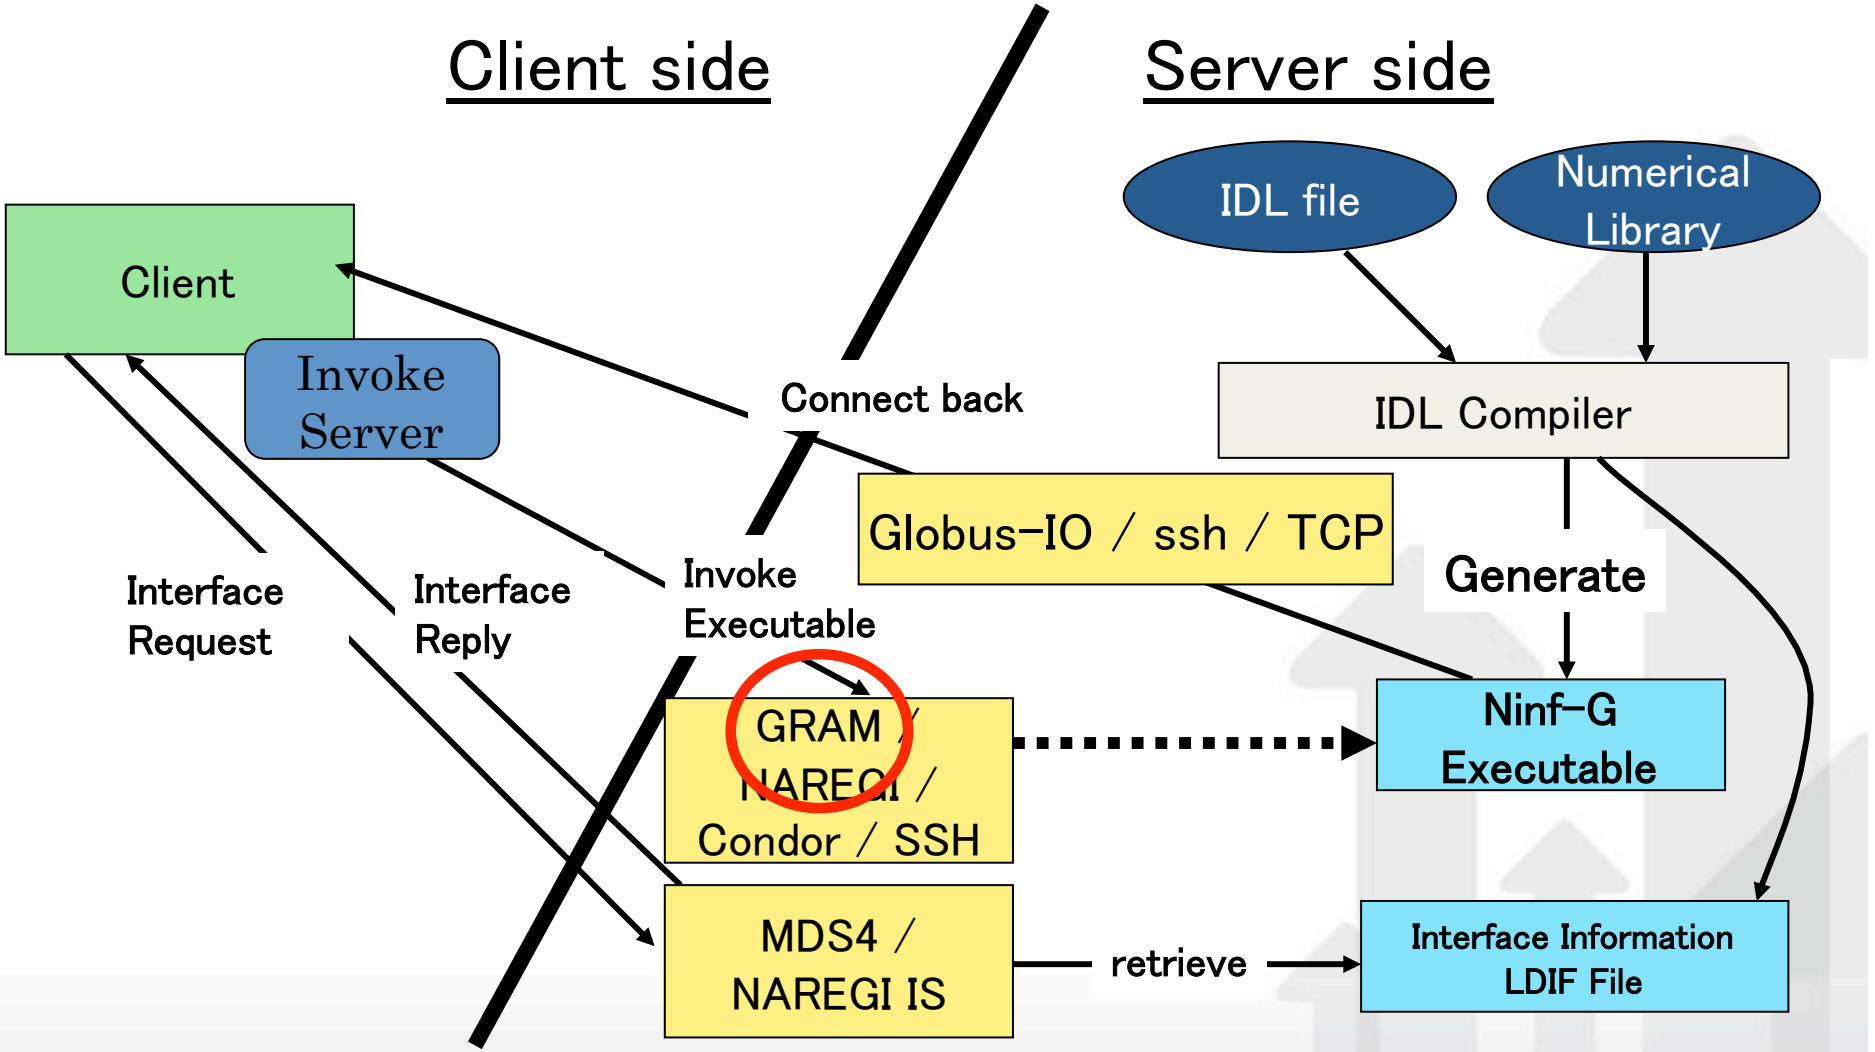

CaBIG and Globus

- caGrid is built on top of Globus 4 WSRF Java Core and Security

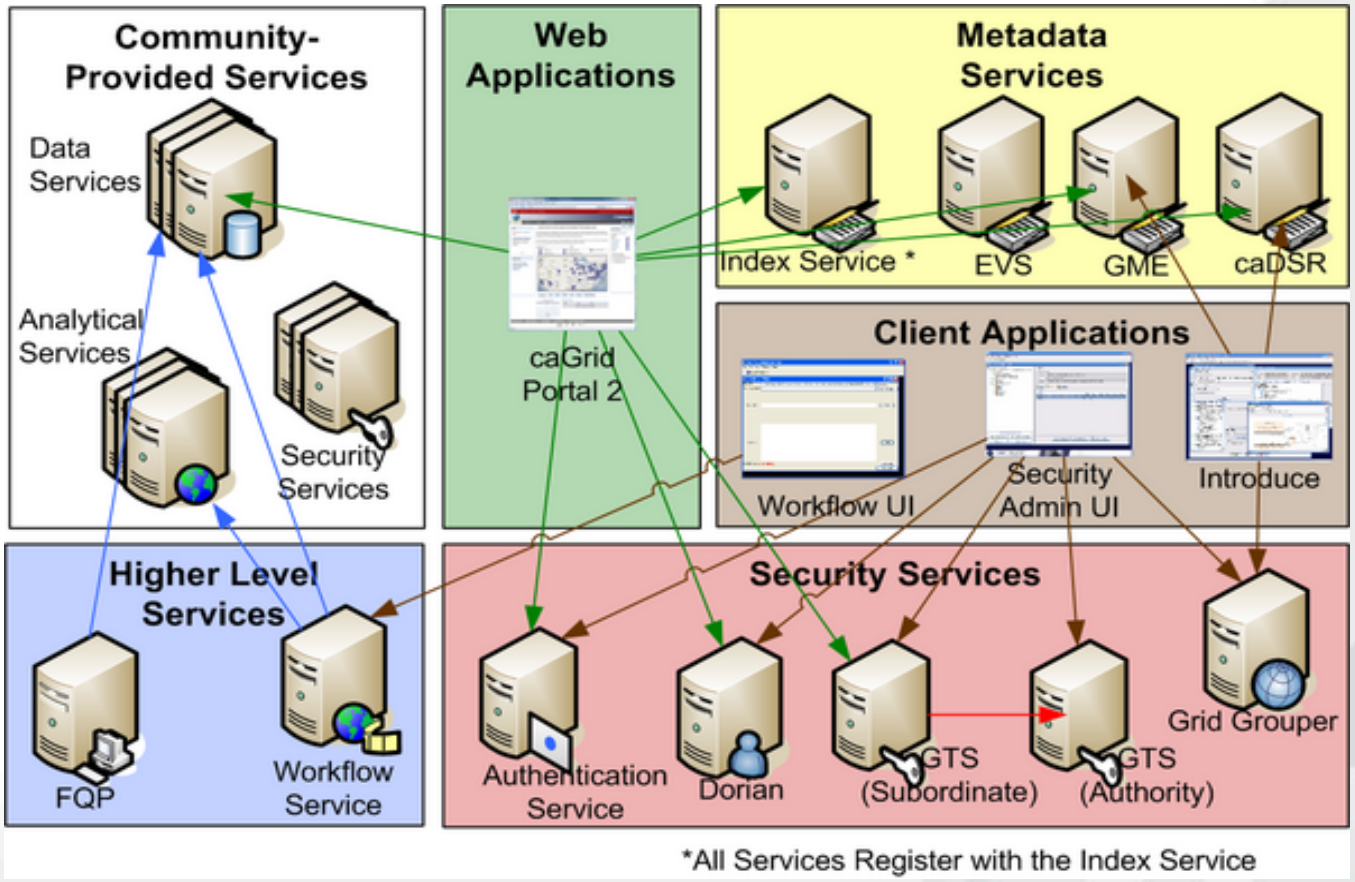

CaBIG – TeraGrid Integration

- Leave caGrid service infrastructure as is with the exception of the analytical services.

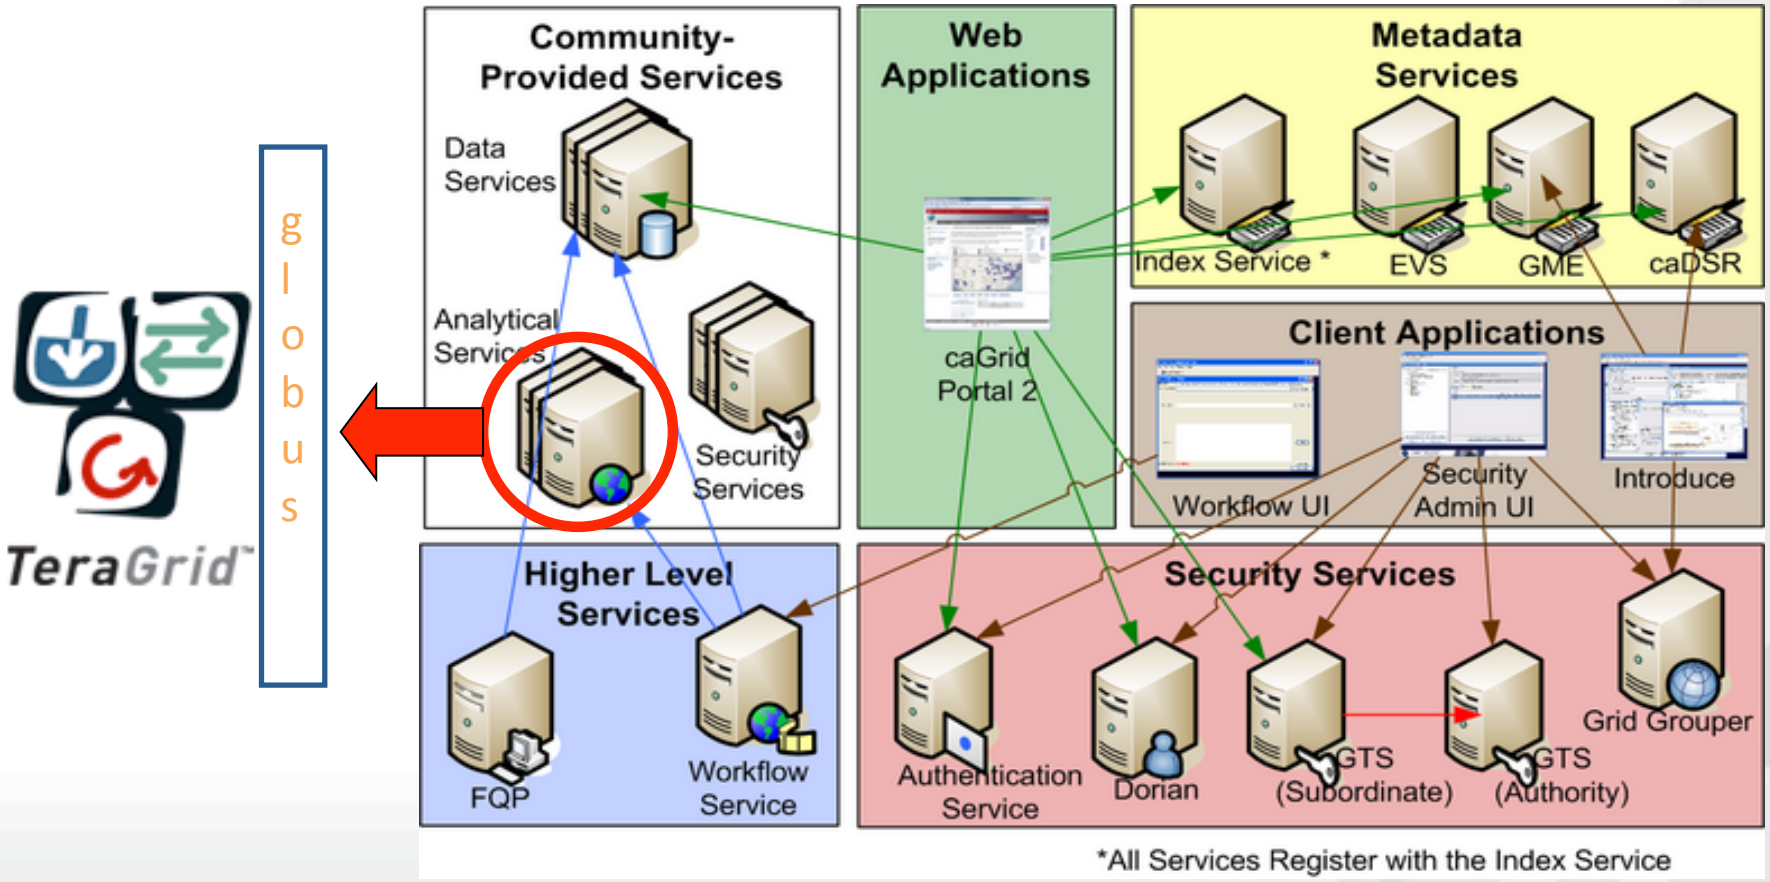

Hierarchical Clustering Results

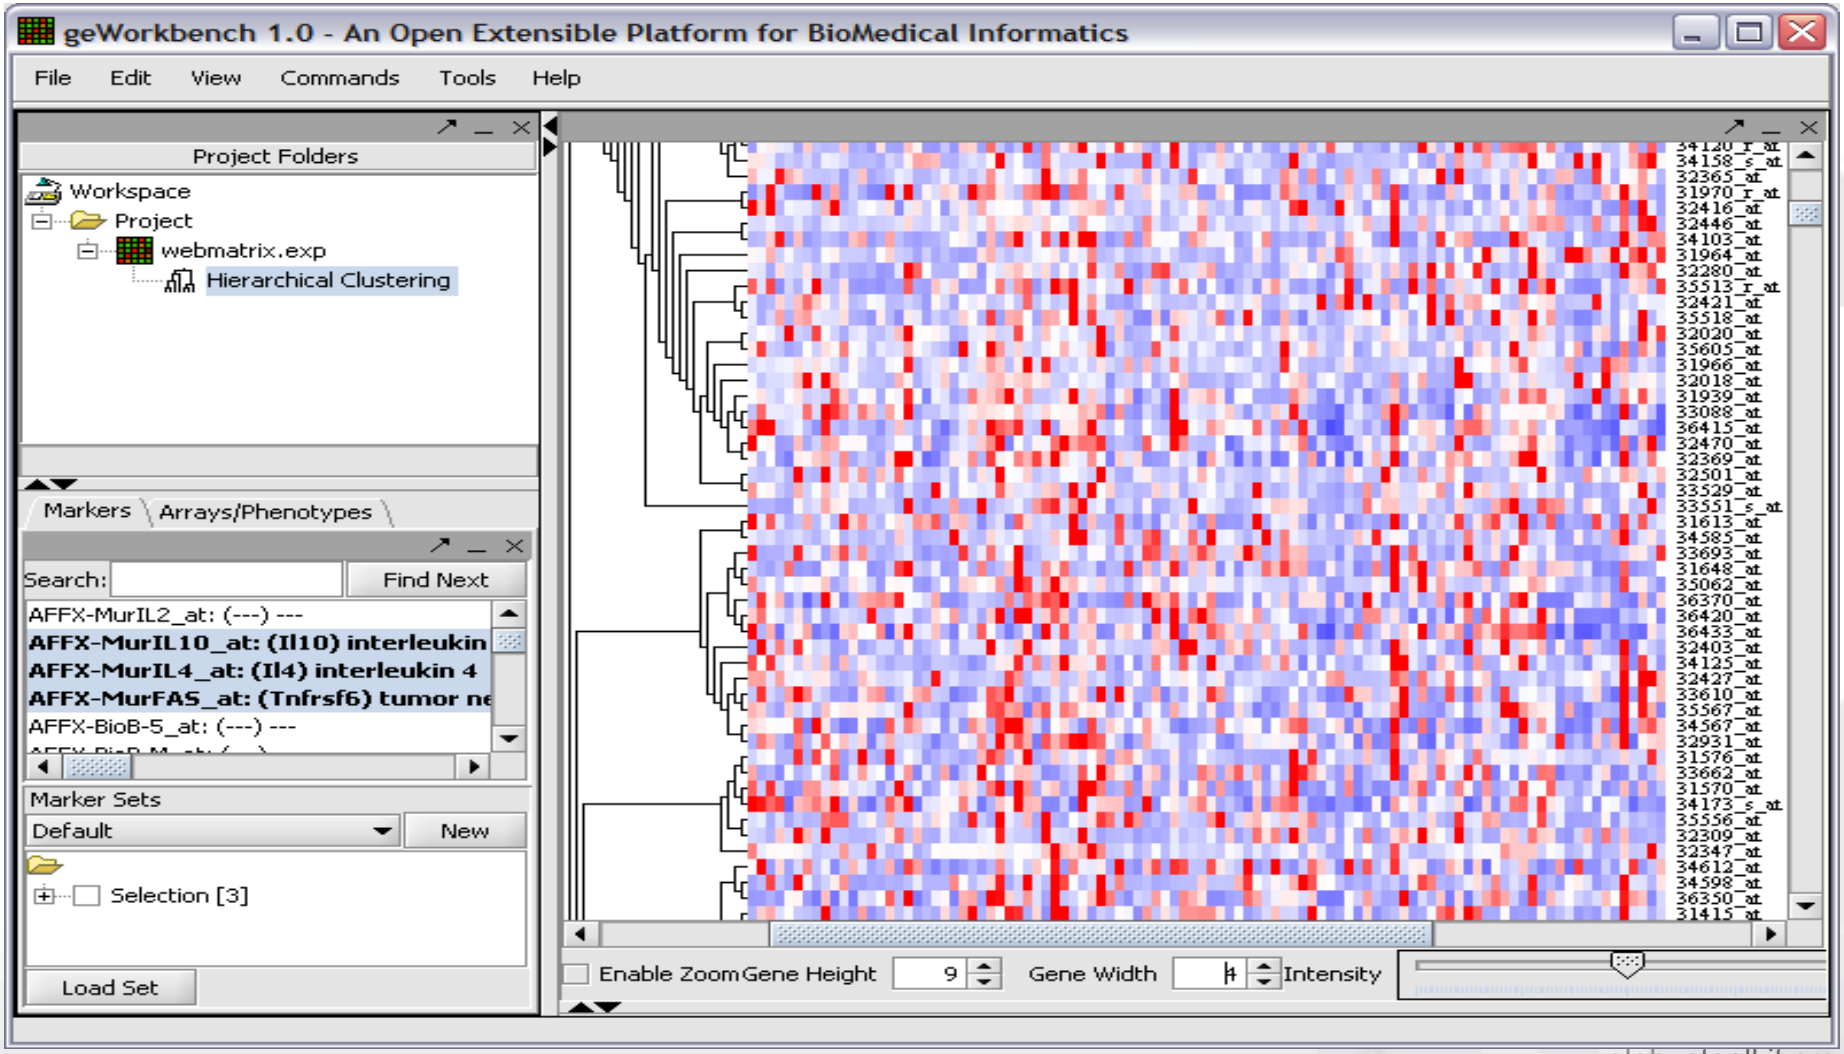

From GRAM2 to GRAM5

Changes Made to Improve Scalability

- **Reduced memory and processor usage**
  - One JM process per user for all their jobs
  - Throttled and efficient perl script executions
  - Job state paged to disk while jobs are executing
- **Removed extra listening port per job for MPIg jobs**
  - Functionality can be re-implemented around GRAM
- **Removed active monitoring of stdout/err files for streaming during LRM job execution**
  - LRMs: transfer stdout/err at the end of job execution
  - Retain for Fork for Condor-G grid monitor

GRAM2 Architecture

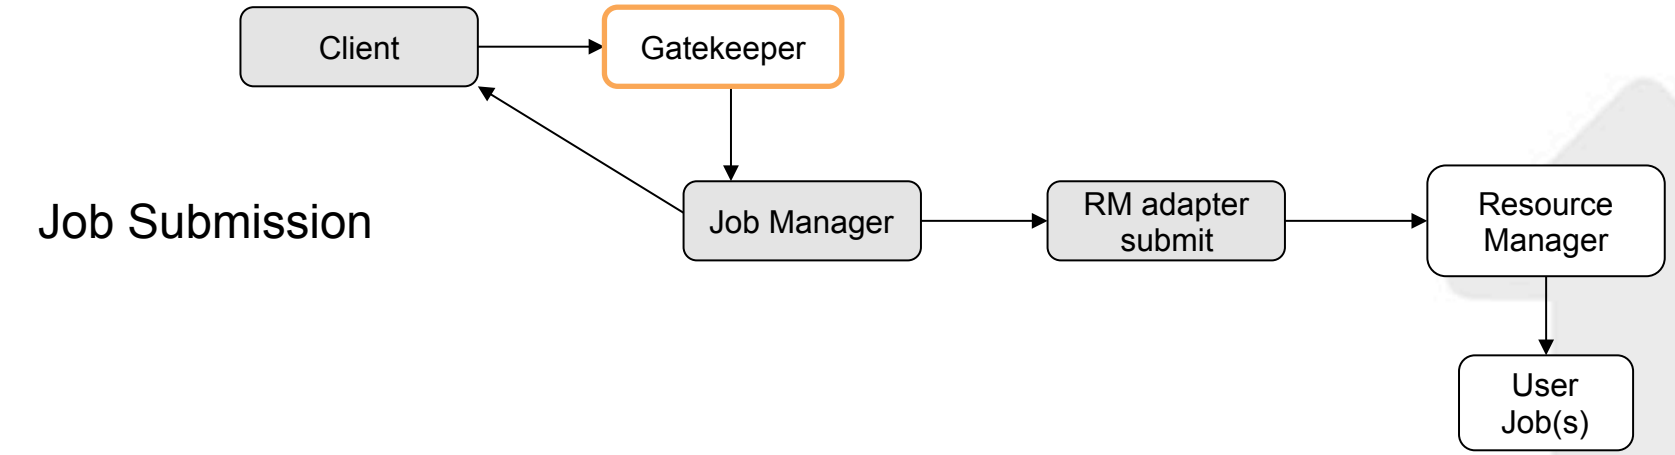

Job Submission

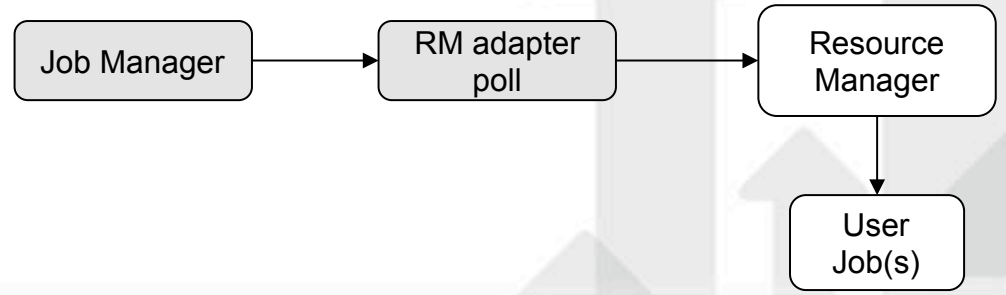

Job Monitoring

GRAM2 Scaling Issues

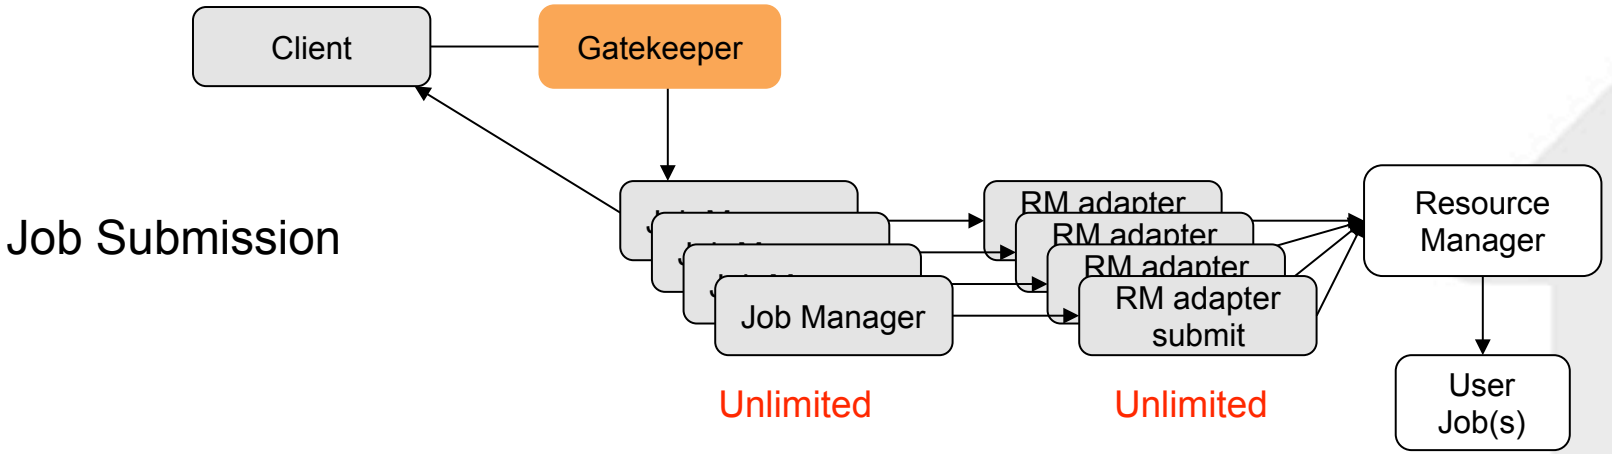

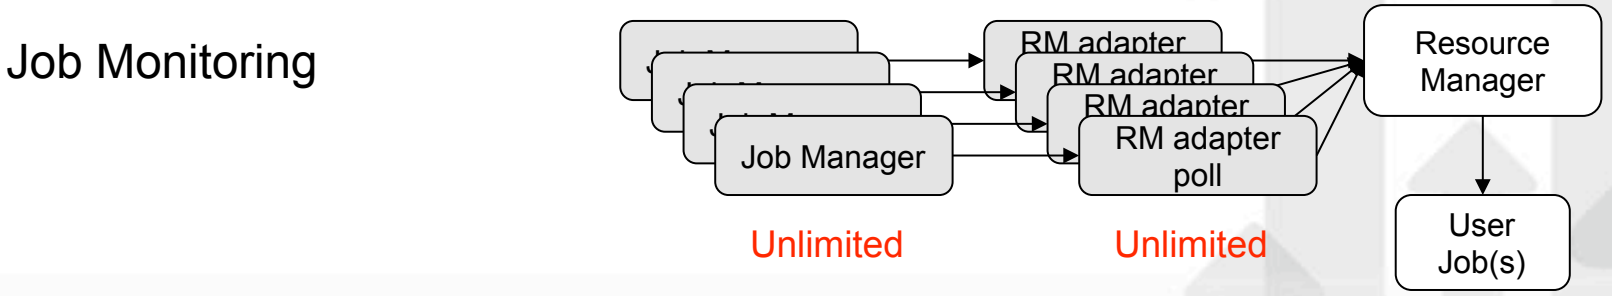

GRAM5 Architecture

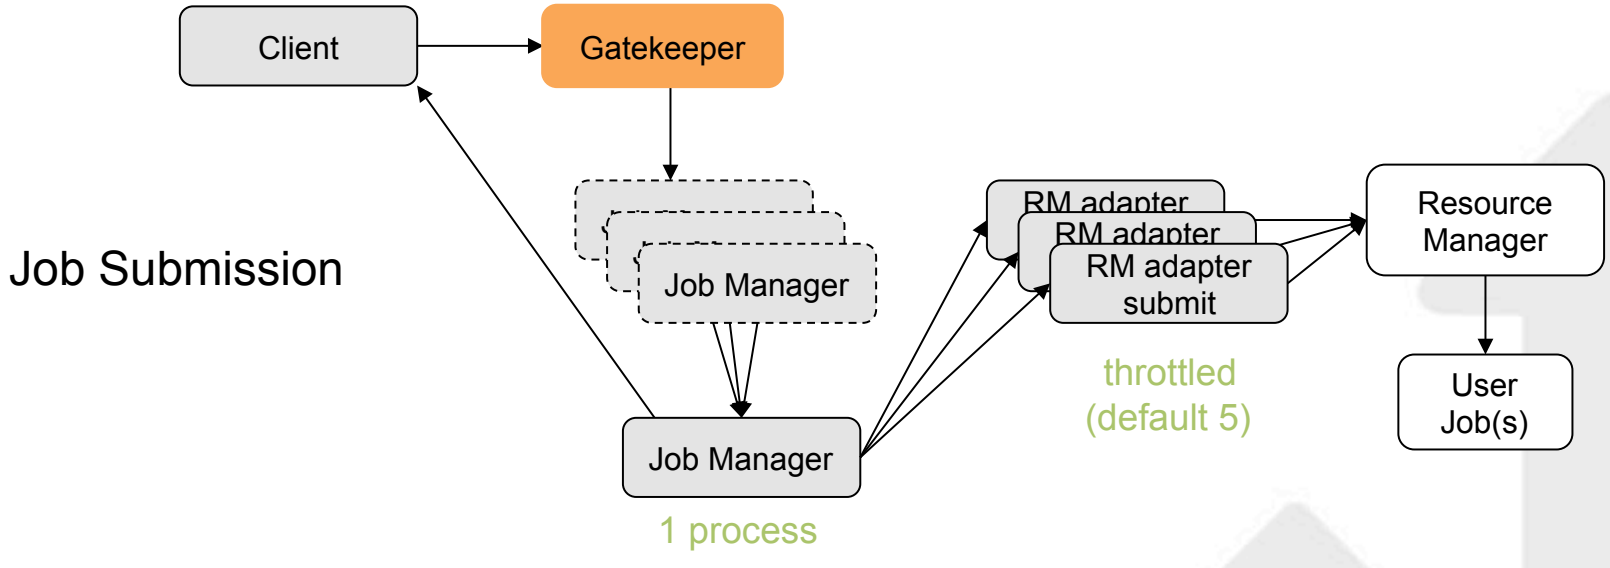

Job Submission

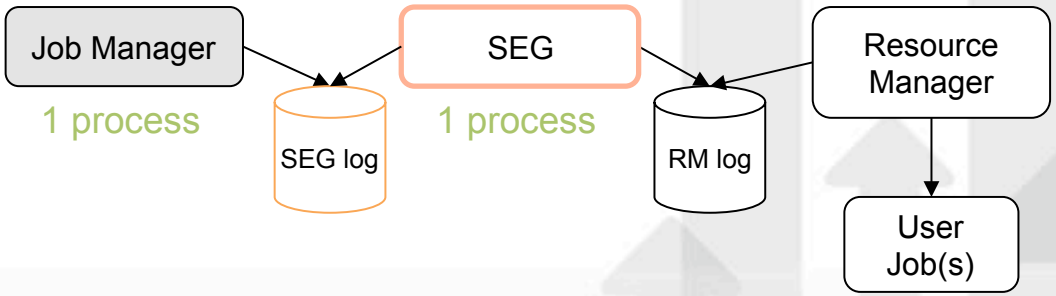

Job Monitoring

Other Improvements

- **Re-implemented Job Manager Logging**
  - Now based on CEDPS logging standard
  - Log levels: fatal, error, warn, info, debug, trace
- **Added job exit codes**
- **Added service version detection**
- **Added usage statistics**
- **Added auditing of TG gateway user attribute**
- **Added SGE adapter**
- **Updated admin, user, developer guides**

Testing Results

- Client: 50 simultaneous reqs; 2000 total jobs; sleep 5 min
- Service: 10 LRM batch slots
- GRAM2 would fail at ~300 jobs, GRAM5 is fine with 7000+

Experiment	Total Jobs	Termination Tasks After 1 Hour	Termination Duration (hh:mm:ss)	Master Node Max 1 min. Load Average	Master Node Average 1 min. Load Average	Errors
1-client-poll	2110	2000	00:14:18	11.46	7.96	None
1-client-seg	2110	2000	00:10:36	2.82	0.93	None
1-client-seg-uncapped	6664	6584	00:42:46	3.26	2.75	None
5-client-seg	6800	6434	00:57:20	3.19	2.49	Connection refused during termination
5-client-seg-diffusers	7226	6720	00:45:41	3.79	3.13	None

Testing Results

- Client: 50 simultaneous reqs; 2000 total jobs; sleep 5 min
- Service: 10 LRM batch slots
- GRAM2 would fail at ~300 jobs, GRAM5 is fine with 7000+

Experiment	Total Jobs	Termination Tasks After 1 Hour	Termination Duration (hh:mm:ss)	Master Node Max 1 min. Load Average	Master Node Average 1 min. Load Average	Errors
1-client-poll	2110	2000	00:14:18	11.46	7.96	None
1-client-seg	2110	2000	00:10:36	2.82	0.93	None
1-client-seg-uncapped	6664	6584	00:42:46	3.26	2.75	None
5-client-seg	6800	6434	00:57:20	3.19	2.49	Connection refused during termination
5-client-seg-diffusers	7226			3.79	3.13	None

SEG is more efficient

Testing Results

- Client: 50 simultaneous reqs; 2000 total jobs; sleep 5 min
- Service: 10 LRM batch slots
- GRAM2 would fail at ~300 jobs, GRAM5 is fine with 7000+

Experiment	Total Jobs	Termination Tasks After 1 Hour	Termination Duration (hh:mm:ss)	Master Node Max 1 min. Load Average	Master Node Average 1 min. Load Average	Errors
1-client-poll	2110	2000	00:14:18	11.46	7.96	None
1-client-seg	2110	2000	00:10:36	2.82	0.93	None
1-client-seg-uncapped	6664					
5-client-seg	6800	6434	00:57:20	3.19	2.49	Connection refused during termination
5-client-seg-diffusers	7226	6720	00:45:41	3.79	3.13	None

250 concurrent job submissions to the same JM proc

Load average remained low

Testing Results

- Client: 50 simultaneous reqs; 2000 total jobs; sleep 5 min
- Service: 10 LRM batch slots
- GRAM2 would fail at ~300 jobs, GRAM5 is fine with 7000+

Experiment	Total Jobs	Termination Tasks After 1 Hour	Termination Duration (hh:mm:ss)	Master Node Max 1 min. Load Average	Master Node Average 1 min. Load Average	Errors
1-client-poll	2110	2000	00:14:18	11.46	7.96	None
1-client-seg	2110	2000	00:10:36	2.82	0.93	None
						None
5-client-seg	6800	6434	00:57:20	3.19	2.49	Connection refused during termination
5-client-seg-diffusers	7226	6720	00:45:41	3.79	3.13	None

1000s of concurrent cancel requests to the same JM process; clients must handle this

Testing Results

- Client: 50 simultaneous reqs; 2000 total jobs; sleep 5 min
- Service: 10 LRM batch slots
- GRAM2 would fail at ~300 jobs, GRAM5 is fine with 7000+

Experiment	Total Jobs	Termination Tasks After 1 Hour	Termination Duration (hh:mm:ss)	Master Node Max 1 min. Load Average	Master Node Average 1 min. Load Average	Errors
1-client-poll	2110	2000	00:14:18	11.46	7.96	None
1-client-seg	2110	2000	00:10:36	2.82	0.93	None
1-client-seg-uncapped						
5-client-seg	6800	6434	00:57:20	3.19	2.49	Connection refused during termination
5-client-seg-diffusers	7226	6720	00:45:41	3.79	3.13	None

5 different users, 5 different JM process, no error

Releases

- **GT 5.0.0 (Jan 2010)**
- **GT 5.0.1 (March)**
- **GT 5.0.2 (July)**
  - Addressed known TG and OSG blocker issues
- **GT 5.0.3 (Feb, 2011)**
- **GT 5.0.4 (Soon early May)**
  - TG request - added per job logging

Next Improvements

- **GT 5.2 Native Packaging**
  - Globus components in Linux distributions
  - RPM and Debian packages
- **Reduce complexity for sites updating GRAM**
  - Modify LRM adapter interface to use callouts instead of patching source directly

Thanks to the Developers!!

- **Joe Bester - ANL**
- **Mike Link - ANL**

Integrated Information Services

JP Navarro, Aaron Diestelkamp

What are Information Services?

A system that enables distributed communities to share information about: computation, storage, and other resources, software, services, science gateways, data collections, system outages; etc.

Grids, in particular, need local information to flow across administrative domains to grid wide information services.

Globus Toolkit offers tools to build Grids:

- GridFTP is grid enabled ***data movement***
- GRAM is grid enabled ***execute***
- GSI OpenSSH is grid enabled ***login***

Integrated Information Services is grid enabled ***information exchange***.

Why a new Information Service?

Community Needs

- Existing communities that rely on GT MDS4 need a long-term replacement.
- New communities need a future looking information services implementation.

Technology Opportunity

- Web 2.0 and REST offer a compelling new paradigm for distributed information sharing that wasn't as obvious when MDS2 and MDS4 were developed.

As we evolved the TeraGrid's MDS4 based information service in the Web 2.0 and REST direction it became clear we had a path forward for a general purpose information services with attractive characteristics for our existing communities.

IIS 3 pronged approach

XML technology and standards

- Continued XML centric storage/exchange (like MDS4) and XSD schemas
  - Arguably the predominant internet information exchange format
- Simple information exchange and discovery using REST APIs
- Generalized query interfaces using Xpath/Xquery
- Flexible transformations using XSLT (can produce CSV, HTML, JSON, …)
- Many XML tools available for buildind a flexible modular system with modest development and maintenance effort

Information Schemas

- Develop and/or integrate important schemas along with the service
- Integrated with other Globus Toolkit components and Globus Online services

Community/Ecosystem

- Build a community around RESTfull APIs/standards, standard schemas, integrating schemas, publishing adaptors, and shared XSLT views on top of standard schemas

Differences with WS-MDS/WebMDS

As a successor to WS-MDS4 / WebMDS, how is it different?

Service Implementation

- Primary CXF JAX-RS interface; may implement JAX-WS eventually
- File-system and XML database (eXist) persistent store versus in-memory
- One service instead of two (WS-MDS and Tomcat/WebMDS)
- XMLSH -> XProc application logic versus custom Java

  XProc is a W3C Pipeline Language at the recommendation stage

Focus on Information Schemas

Focus on Community Building around common information requirements

Architectural Overview

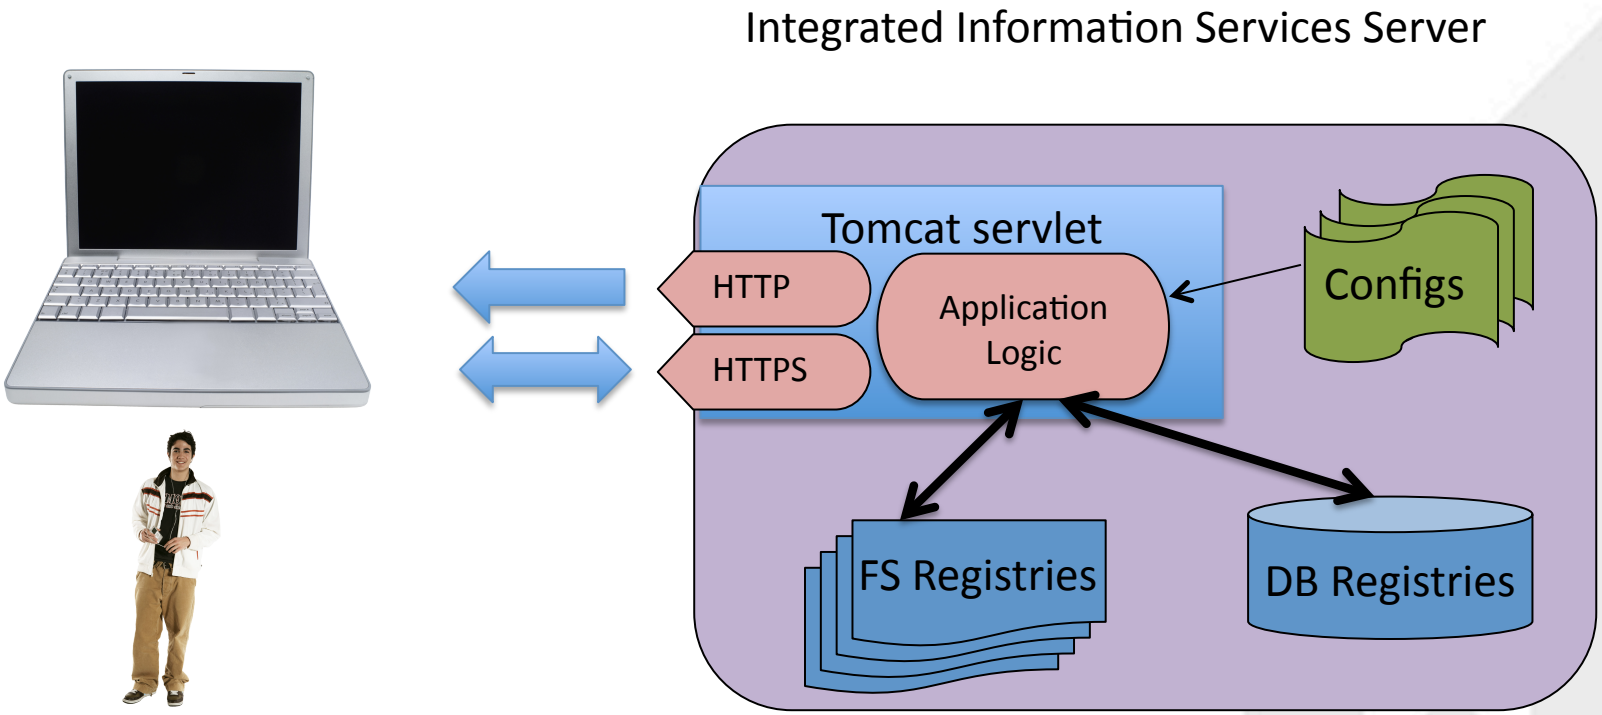

Each Registry is a forest of XML documents that share a common schema

Roadmap – Alpha 0.1.0

Integrated Information Services Lite (looking for a better name)

- Alpha version 0.1.0 – Released this week
  - http://www.ci.uchicago.edu/iis/)
- Simple deployable Tomcat based information service
- With HTTPS basic authentication
- RESTfull API
- Both File-system and eXist (XMLDB) XML document repositories

Information Schemas

- Sample and testing capability registries
- TeraGrid/BIRN Capability Definition and Deployment registries

REST API Overview

BASE_URL = http[s]://{server}:{port}/registry/rest/

BASE_URL
- List available registries

BASE_URL/{registry_name}/.index
- List registry's document index

BASE_URL/{registry_name}/
- GET: Retrieve all registry documents (a)
- POST: store a new registry document

BASE_URL/{registry_name}/{unique_id}
- GET: Retrieve a single document based on unique document ID (a)
- PUT: store a registry document
- DELETE: delete a registry document

BASE_URL/{registry_name}/{element}/{value}
- Retrieve all documents with first order element that has specified value (a)

BASE_URL/{registry_name}/{@attribute}/{value}
- Retrieve all documents with base attribute that has specified value (a)

BASE_URL/{registry_name}/?xpath={xpath_expression}
- Access entire registry based on xpath expression
- Any URL labeled with (a) can have an xpath

BASE_URL/{registry_name}/?xslt="{xslt_filename}"
- Access entire registry and transform using specified xslt file (on server)
- Any URL labeled with (a) can specify an xslt transform

Example: http://info.test.globus.org/registry/rest/

Roadmap – Future plans

Service Implementation

- More authentication methods and finer-grained authorization
- Improved performance and scalability (current weakness)
- Information synchronization across instances (push and pull)
- API based registry administration
- AMQP publish/subscribe integration (implemented by XD)
- WS-* support (possibly required by XD)
- Globus Toolkit packaging and distribution at some point

Information Schemas

- TeraGrid/XD and OGF GLUE2 implementations
- TeraGrid/XD capability information model implementation
- Community driven schema, adaptor, and view repository

If there is demand, Globus Online Hosted IIS?

More information

Web home

- http://www.ci.uchicago.edu/wiki/bin/view/IIS/WebHome

The software

- Download: http://www.ci.uchicago.edu/iis/
- Browse: https://trac.ci.uchicago.edu/iis/browser

E-mail lists

- iis-user@globus.org
- iis-dev@globus.org

Distributed Architecture

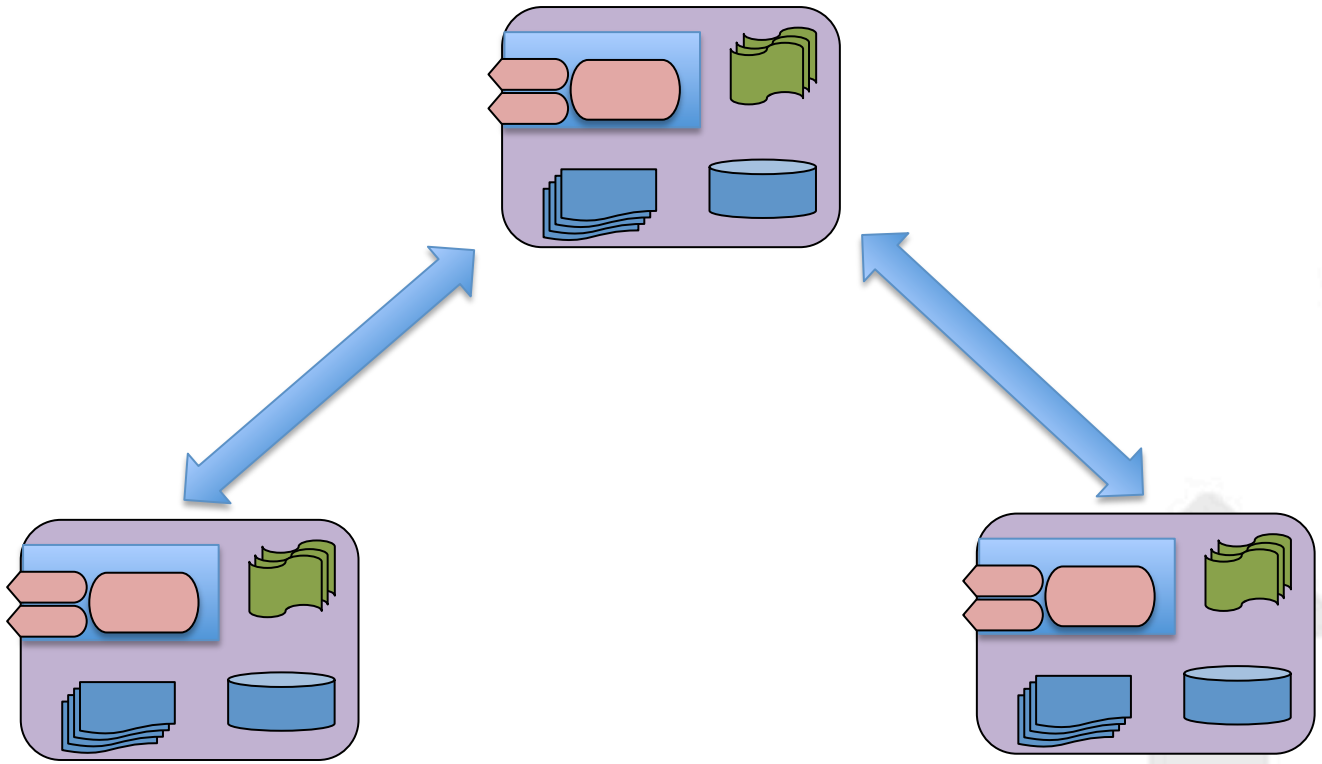

Technical Approach

- XML centric information registry
- Leverages mainstream tools:
  - Service: Java, Tomcat, Apache CXF, JAX-RS
  - Application logic: XMLSH ; *eventually replaced by Xproc Calabash*
  - Registry storage: file-system or XMLDB(eXist) storage; *eventually maybe Sedna and others*
- Supported standards:
  - REST
  - XML with full XPath and xslt support; *plus Xquery eventually*
  - HTML, JSON, RSS, and CSV output support
- Simple modular design
- Simple to deploy and operate
- Customizable with application logic, additional storage adaptors, etc.
- Open source, open to community contribution