



the globus alliance
www.globus.org

Globus and TeraGrid: Current Experience and Near Term Plans

Jennifer M. Schopf

Argonne National Lab

<http://www.mcs.anl.gov/~jms/Talks/>



What is a Grid?

- Resource sharing
 - Computers, storage, sensors, networks, ...
 - Sharing always conditional: issues of trust, policy, negotiation, payment, ...
- Coordinated problem solving
 - Beyond client-server: distributed data analysis, computation, collaboration, ...
- Dynamic, multi-institutional virtual orgs
 - Community overlays on classic org structures
 - Large or small, static or dynamic



Applications and Grids

- Why is this hard?
 - Lack of central control
 - Shared resources
 - Communication and coordination
- So why do it?
 - Computations that need to be done with a time limit
 - Data that can't fit on one site
 - Data owned by multiple sites
- Applications that need to be run bigger, faster, more



Globus is an Hour Glass

- Local sites have their own policies, installs – heterogeneity!

- Queuing systems, monitors, network protocols, etc

- Globus unifies – standards!

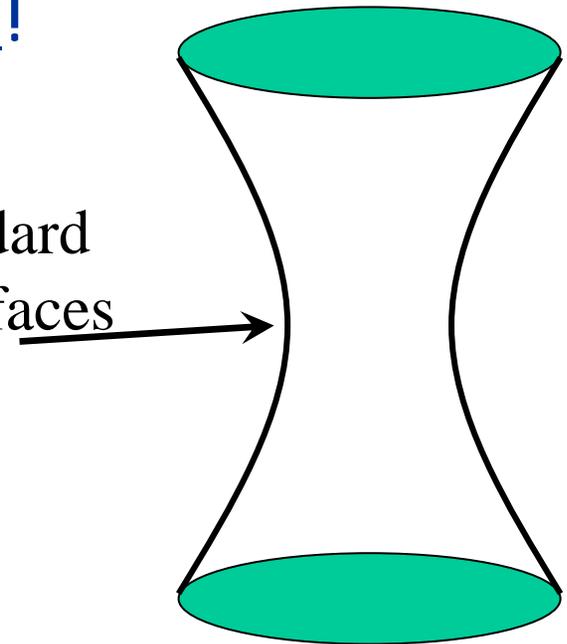
- Build on Web services

- Use WS-RF, WS-Notification to represent/access state

- Common management abstractions & interfaces

Higher-Level Services
and Users

Standard
Interfaces



Local heterogeneity



Globus is a Building Block

- Basic components for Grid functionality
 - Not turnkey solutions, but building blocks & tools for application developers & system integrators
- Highest-level services are often application specific, we let aps concentrate there
- Easier to reuse than to reinvent
 - Compatibility with other Grid systems comes for free
- We provide basic infrastructure to get you one step closer



Globus Philosophy

- Globus was first established as an open source project in 1996
- The Globus Toolkit is open source to:
 - Allow for inspection
 - > for consideration in standardization processes
 - Encourage adoption
 - > in pursuit of ubiquity and interoperability
 - Encourage contributions
 - > harness the expertise of the community
- The Globus Toolkit is distributed under the (BSD-style) Apache License version 2



dev.globus

- Governance model based on Apache Jakarta
 - Consensus based decision making
- Globus software is organized as several dozen “Globus Projects”
 - Each project has its own “Committers” responsible for their products
 - Cross-project coordination through shared interactions and committers meetings
- A “Globus Management Committee”
 - Overall guidance and conflict resolution

- [Welcome](#)
- [List of projects](#)
- [Guidelines](#)
- [Infrastructure](#)
- [How to contribute](#)
- [GlobDev events](#)
- [Recent changes](#)
- [GlobDev FAQ](#)

common runtime projects

- [C Core Utilities](#)
- [C WS Core](#)
- [CoG jglobus](#)
- [Core WS Schema](#)
- [Java WS Core](#)
- [Python Core](#)
- [XIO](#)

data projects

- [GridFTP](#)
- [OGSA-DAI](#)
- [Reliable File Transfer](#)
- [Replica Location](#)

execution projects

- [GRAM](#)

information projects

- [MDS4](#)

security projects

- [C Security](#)
- [CAS/SAML Utilities](#)
- [Delegation Service](#)

Guidelines
(Apache
Jakarta)

Infrastructure
(CVS, email,
bugzilla, Wiki)

Projects
Include
...

Welcome

This is the new home Globus software development; it is still under construction. The current status of our efforts to build this environment can be found [on this page](#). Comments regarding this site can be sent to info@globus.org. Thank you for your interest in Globus development!

Globus was first established as an open source software project in 1996. Since that time, the Globus development team has expanded from a few individuals to a distributed, international community. In response to this growth, the Globus community (the "Globus Alliance") established in October 2005 a new source code development *infrastructure* and meritocratic *governance model*, which together make the process by which a developer joins the Globus community both easier and more transparent.

The Globus governance model and infrastructure are based on those of [Apache Jakarta](#). In brief, the governance model places control over each individual software component (*project*) in the hands of its most active and respected *contributors* (*committers*), with a *Globus Management Committee* (GMC) providing overall guidance and conflict resolution. The infrastructure comprises *repositories*, *email lists*, Wikis, and *bug trackers* configured to support per-project community access and management.

For more information, see:

- [The Globus Alliance Guidelines](#), which address various aspects of the Globus governance model and the Globus community.
- A description of the Globus Alliance [Infrastructure](#).
- A list of current Globus projects.
- Information about Globus community events.
- The [conventions and guidelines](#) that apply to contributions



Globus Technology Areas

- Core runtime
 - Infrastructure for building new services
- Security
 - Apply uniform policy across distinct systems
- Execution management
 - Provision, deploy, & manage services
- Data management
 - Discover, transfer, & access large data
- Monitoring
 - Discover & monitor dynamic services



Non-Technology Projects

- Distribution Projects
 - Globus Toolkit Distribution
 - Process in use since April '07
- Documentation Projects
 - GT Release Manuals
- Incubation Projects
 - Incubation management project
 - And any new projects wanting to join



Globus Projects

MPICH-G2

GridWay

Incubator Mgmt

Java Runtime

C Runtime

Python Runtime

Delegation

CAS

C Sec

MyProxy

GSI-OpenSSH

GRAM

OGSA-DAI

Data Rep

GridFTP

Reliable File Transfer

GT4

Replica Location

MDS4

GT4 Docs

Incubator Projects

Common Runtime

Security

Execution Mgmt

Data Mgmt

Info Services

Other



Globus Projects

MPICH-G2

GridWay

Incubator Mgmt

Java Runtime

Delegation

MyProxy

OGSA-DAI

GT4

Data Rep

Replica Location

C Runtime

CAS

GSI-OpenSSH

GridFTP

MDS4

Python Runtime

C Sec

GRAM

Reliable File Transfer

GT4 Docs

Incubator Projects

GAARDS

MEDICUS

Cog WF

Virt WkSp

GDTE

GridShib

OGRO

UGP

Dyn Acct

Gavia JSC

DDM

Metrics

Introduce

PURSE

HOC-SA

LRMA

WEEP

Gavia MS

SGGC

ServMark

Common Runtime

Security

Execution Mgmt

Data Mgmt

Info Services

Other



Globus Software: dev.globus.org

Globus Projects

MPICH-G2

GridWay

Incubator Mgmt

Java Runtime

C Runtime

Python Runtime

Delegation

CAS

C Sec

MyProxy

GSI-OpenSSH

GRAM

OGSA-DAI

Data Rep

GridFTP

Reliable File Transfer

GT4

Replica Location

MDS4

GT4 Docs

Incubator Projects

GAARDS

MEDICUS

Cog WF

Virt WkSp

GDTE

GridShib

OGRO

UGP

Dyn Acct

Gavia JSC

DDM

Metrics

Introduce

PURSE

HOC-SA

LRMA

WEEP

Gavia MS

SGGC

ServMark

Common Runtime

Security

Execution Mgmt

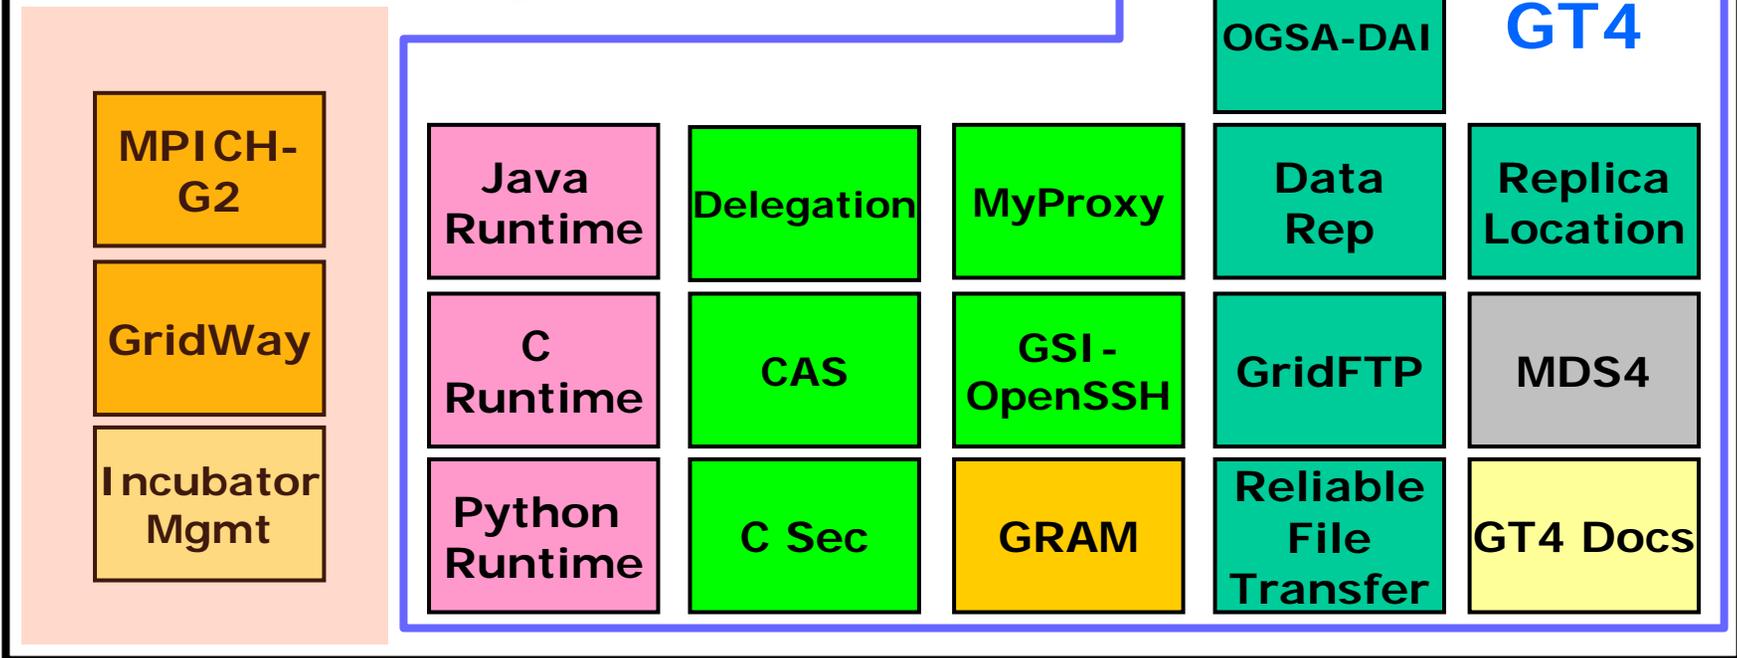
Data Mgmt

Info Services

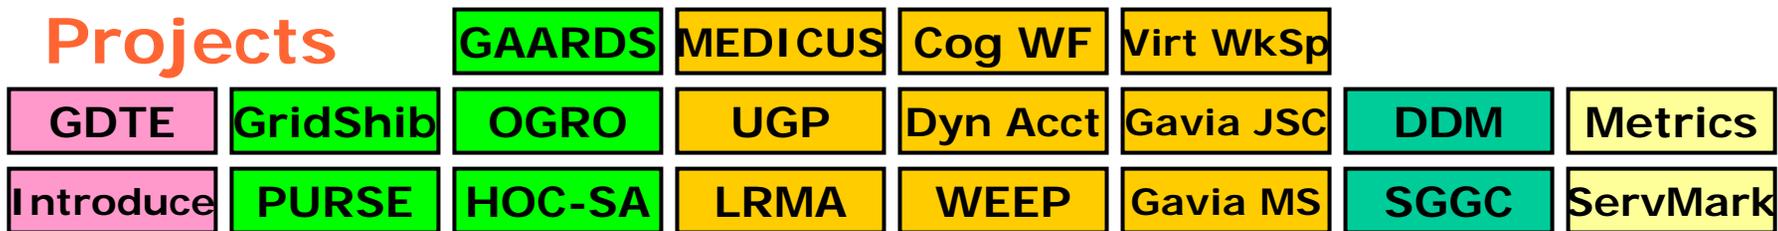
Other



Globus Projects



Incubator Projects





GT4 Core Functionality

Reference implementation of WSRF and WS-N functions

- Naming and bindings (basis for virtualization)
 - Every resource can be uniquely referenced and has one or more associated services for interacting
- Lifecycle (basis for resilient state management)
 - Resources created by svcs following a factory pattern
 - Resource destroyed immediately or scheduled
- Information model (basis for monitoring & discovery)
 - Resource properties associated with resources
 - Operations for querying and setting this info
 - Asynchronous notification of changes to properties
- Service groups (basis for registries & collective svcs)
 - Group membership rules and membership management
- Base fault type



WSRF vs XML/SOAP

- The definition of WSRF means that the Grid and Web services communities can move forward on a common base
- Why Not Just Use XML/SOAP?
 - WSRF and WS-N *are* just XML and SOAP
 - WSRF and WS-N are just Web services
- Benefits of following the specs:
 - These patterns represent best practices that have been learned in many Grid applications
 - There is a community behind them
 - Why reinvent the wheel?
 - Standards facilitate interoperability



Upcoming Changes to Core

- Currently working on specification upgrade
 - Final versions of WS-RF, WS-N, WS-A
 - -Required for compliance with some OGF profiles (E.g. OGSA WSRF Basic Profile, OGSA Basic Security profile)
 - To be included in GT 4.2 release upcoming
- Also evaluating implications of an update to Axis2
 - Determine if truly needed
 - Will have an open call to solicit timeline
 - Would likely delay 4.2 release if included



Globus Security

- Control access to shared services
 - Address autonomous management, e.g., different policy in different work-groups
- Support multi-user collaborations
 - Federate through mutually trusted services
 - Local policy authorities rule
- Allow users and application communities to set up dynamic trust domains
 - Personal/VO collection of resources working together based on trust of user/VO



Security Tools

- Basic Grid Security Mechanisms
- Certificate Generation Tools
- Certificate Management Tools
 - Getting users “registered” to use a Grid
 - Getting Grid credentials to wherever they’re needed in the system
- Authorization/Access Control Tools
 - Storing and providing access to system-wide authorization information
- Credential management service
 - MyProxy (One time password support)

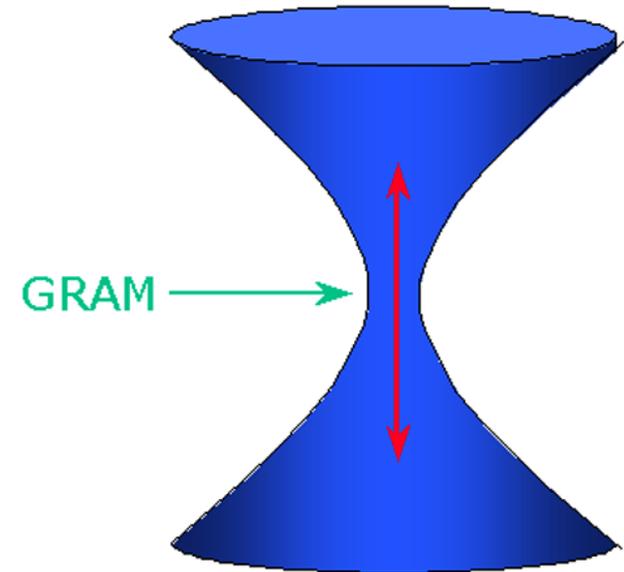


GRAM - Basic Job Submission and Control Service

- A uniform service interface for remote job submission and control
 - Includes file staging and I/O management
 - Includes reliability features
 - Supports basic Grid security mechanisms
 - Available in Pre-WS and WS
- GRAM is *not* a scheduler.
 - No scheduling
 - No metascheduling/brokering
 - Often used as a front-end to schedulers, and often used to simplify metaschedulers/brokers

Applications

Metaschedulers, Brokers



Local Management Mechanisms



GT4 WS GRAM

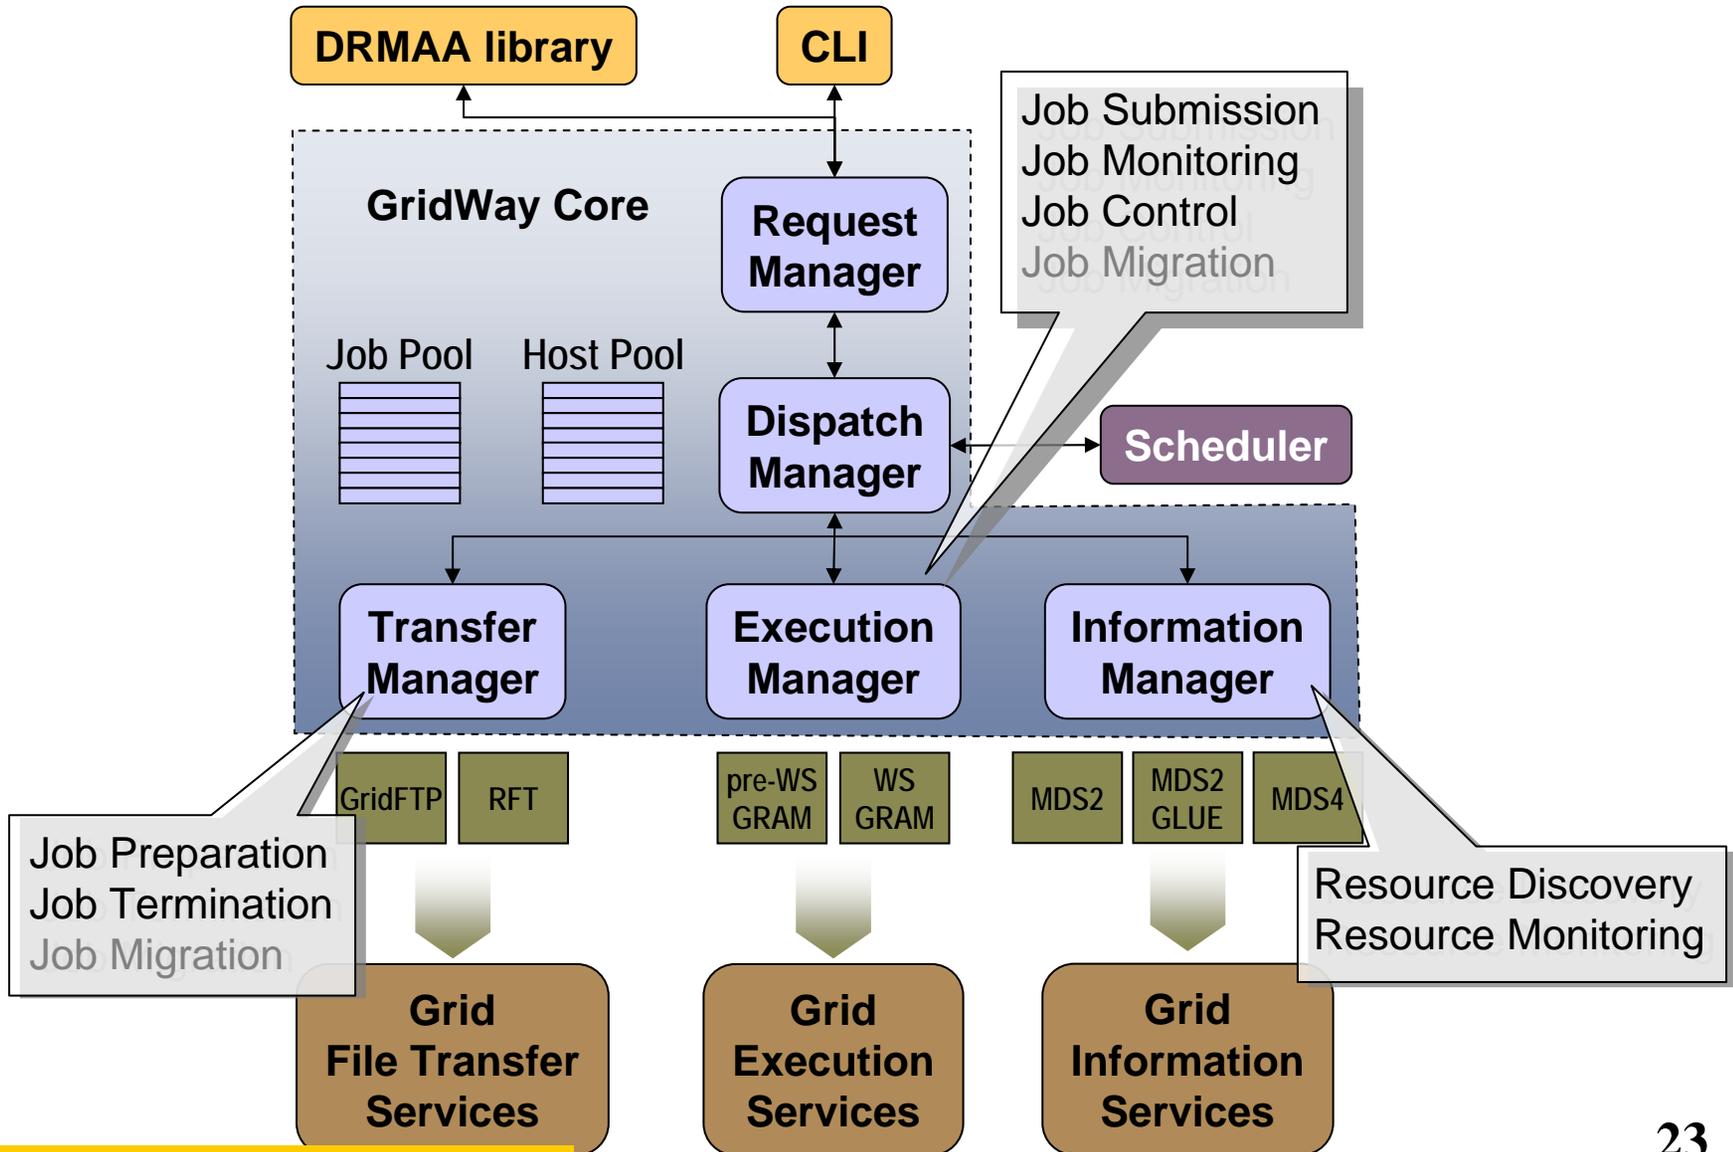
- 2nd-generation WS implementation optimized for performance, flexibility, stability, scalability
- Streamlined critical path
 - Use only what you need
- Flexible credential management
 - Credential cache & delegation service
- GridFTP & RFT used for data operations
 - Data staging & streaming output
 - Eliminates redundant GASS code



GridWay Meta-Scheduler

- Scheduler virtualization layer on top of Globus services
 - A LRM-like environment for submitting, monitoring, and controlling jobs
 - A way to submit jobs to the Grid, without having to worry about the details of exactly which local resource will run the job
 - A policy-driven job scheduler, implementing a variety of access and Grid-aware load balancing policies
 - Accounting

GridWay Meta-Scheduler

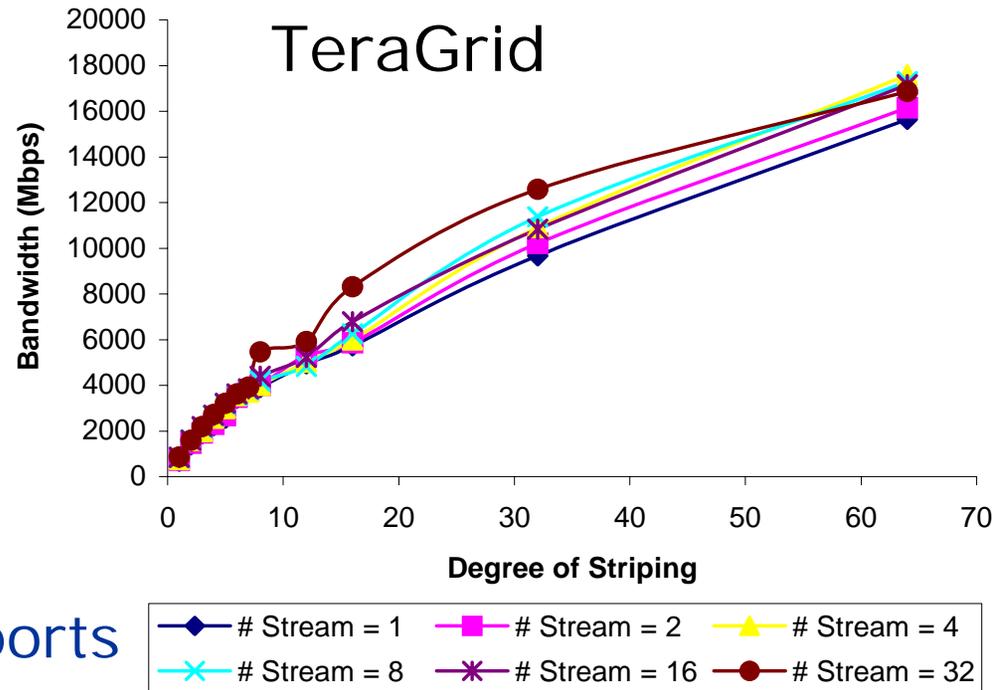




GridFTP in GT4

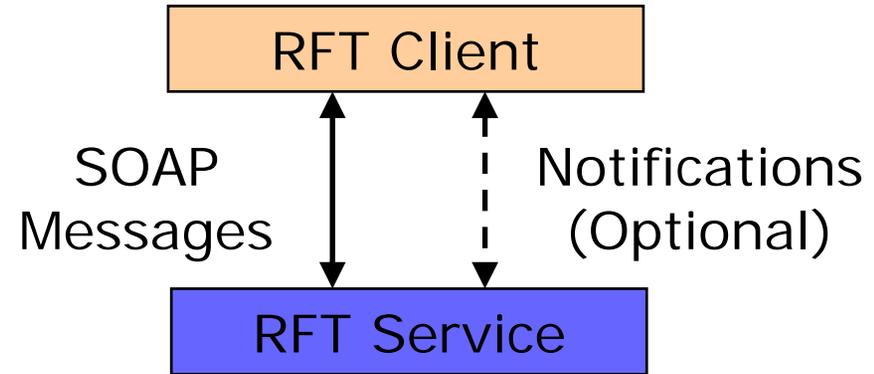
- 100% Globus code
 - No licensing issues
 - Stable, extensible
- IPv6 Support
- XIO for different transports
- Striping → multi-Gb/sec wide area transport
- Pluggable
 - Front-end: e.g., future WS control channel
 - Back-end: e.g., HPSS, cluster file systems
 - Transfer: e.g., UDP, NetBLT transport

Disk-to-disk on TeraGrid



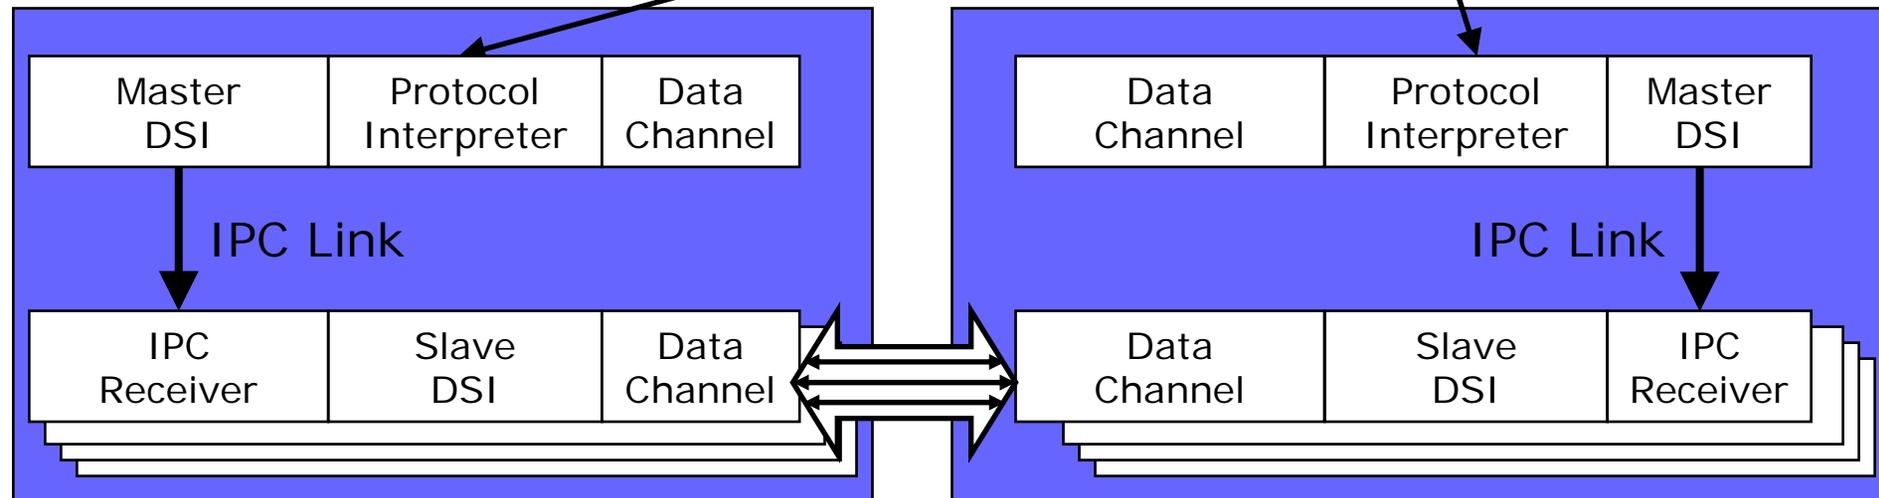
Reliable File Transfer: Third Party Transfer

- Fire-and-forget transfer
- Web services interface
- Many files & directories
- Integrated failure recovery
- Has transferred 900K files



GridFTP Server

GridFTP Server





TGCP - TeraGrid Copy

- Applies a set of transformation rules to source and destination
 - Local admin supplies the rules
 - Adds host/port and appropriate path information, puts into GridFTP URL format
- When source/dest sites are identified, add network tuning parameters based on a table maintained by admins
- Invoke either g-u-c or rft to perform the transfer



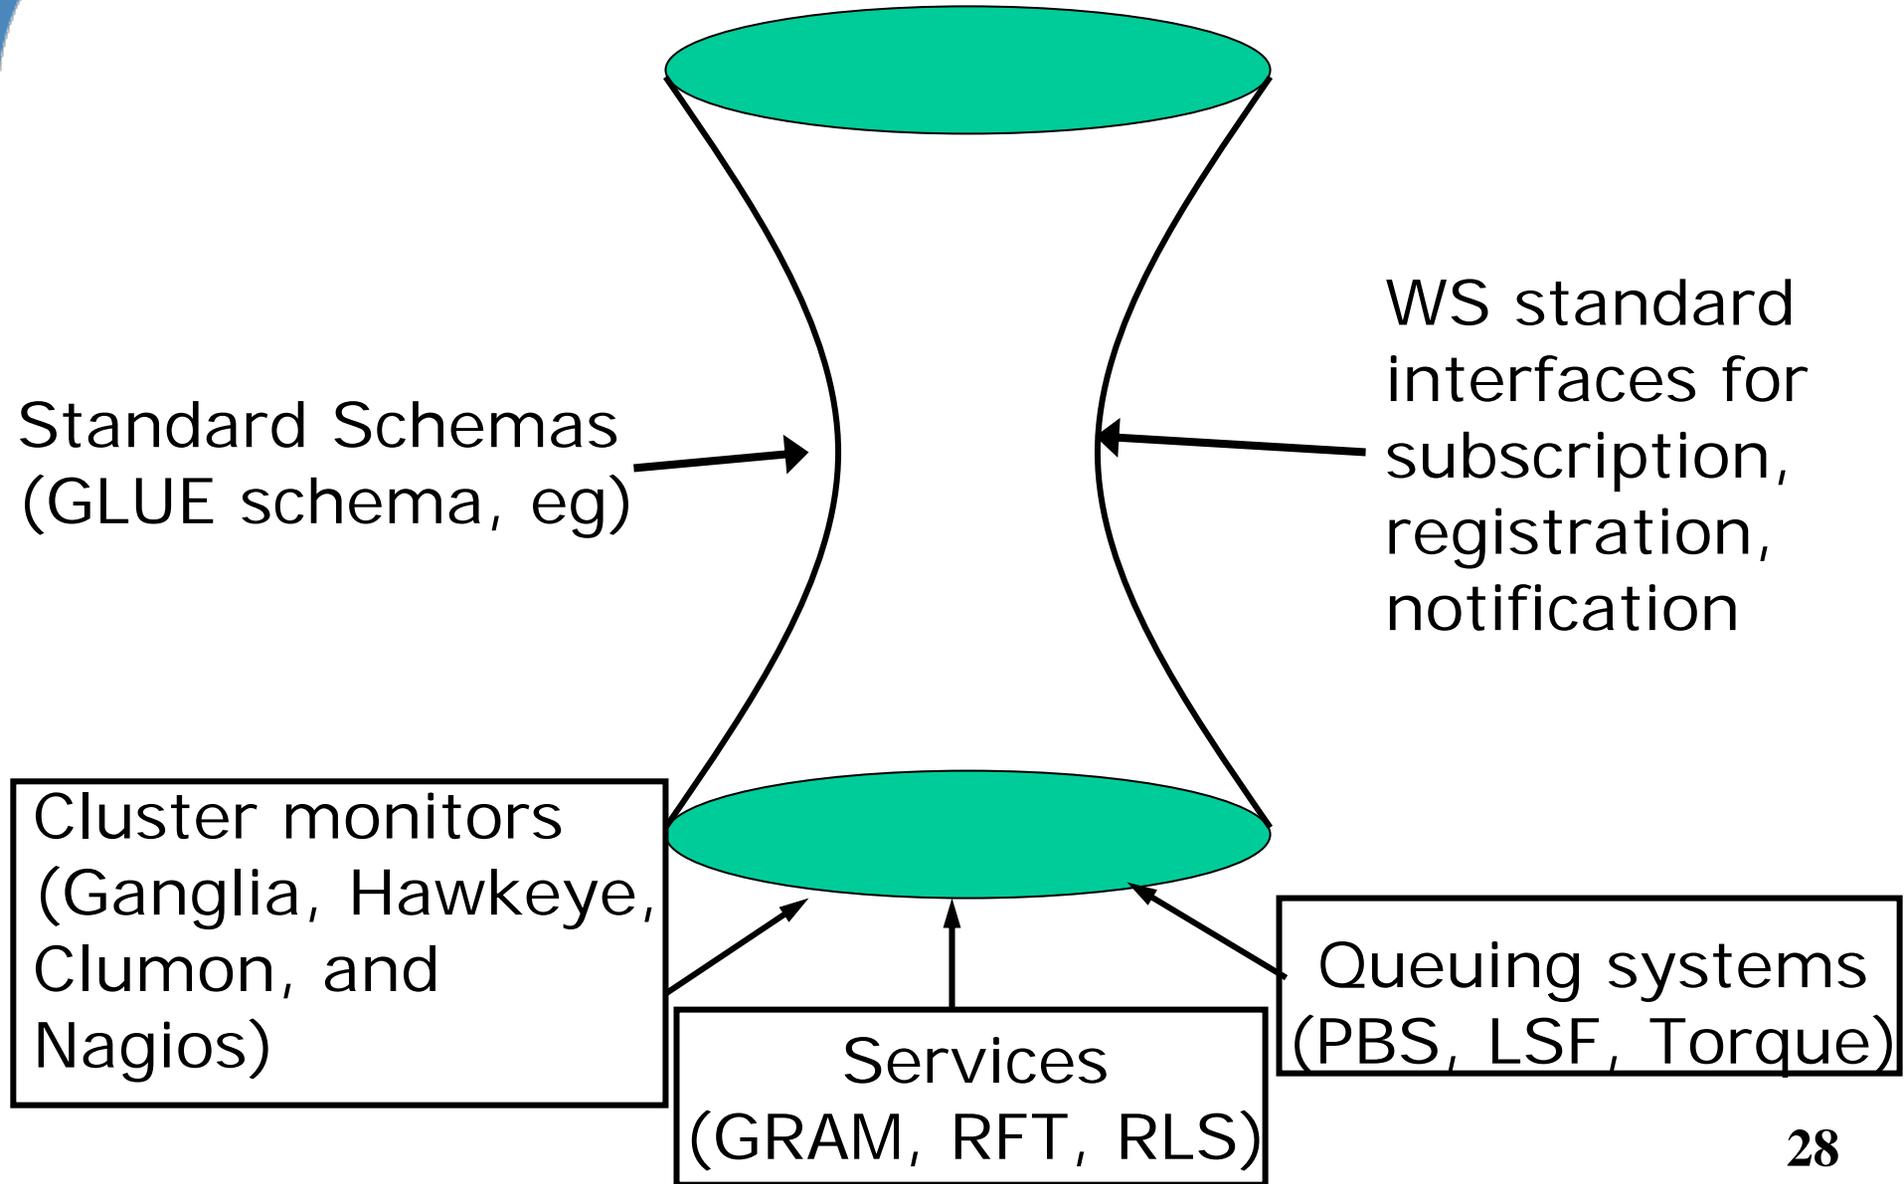
Monitoring and Discovery System (MDS4)

- Grid-level monitoring system used most often for resource selection
 - Aid user/agent to identify host(s) on which to run an application
- Uses standard interfaces to provide publishing of data, discovery, and data access, including subscription/notification
 - WS-ResourceProperties, WS-BaseNotification, WS-ServiceGroup
- Functions as an hourglass to provide a common interface to lower-level monitoring tools



Information Users :

Schedulers, Portals, Warning Systems, etc.



File Edit View Favorites Tools Help

Back Forward Stop Refresh Home Search Favorites

Address <http://128.9.64.250:8080/webmds/webmds?info=openEndedQuery&xmlSource.openEndedQuery.param.endpoint=http%3A%2F%2F141.142.48.5%3A20202%2Fwsrf%2Fservices%2FDefaultInd> Go Links

Queue Overview

Name	UniqueId	Gram Information			LRMS		CPUs		Status	Jobs			Policy Limits			
		Version	Host	Port/URL	Type	Version	Total	Free		Total	Running	Waiting	Wall Clock Time	CPU Time	Total Jobs	Running Jobs
big	big	4.0.1	tg-login1.ncsa.teragrid.org	2019	PBS-Torque	2.0.0p7	891	538	enabled	0	0	0	2880	-1	-1	-1
dque	dque	4.0.1	tg-login1.ncsa.teragrid.org	2019	PBS-Torque	2.0.0p7	891	538	enabled	171	50	121	1440	-1	-1	-1
long	long	4.0.1	tg-login1.ncsa.teragrid.org	2019	PBS-Torque	2.0.0p7	891	538	enabled	0	0	0	5760	-1	-1	-1
priority	priority	4.0.1	tg-login1.ncsa.teragrid.org	2019	PBS-Torque	2.0.0p7	891	538	enabled	0	0	0	1440	-1	-1	-1
debug	debug	4.0.1	tg-login1.ncsa.teragrid.org	2019	PBS-Torque	2.0.0p7	891	538	enabled	0	0	0	30	-1	-1	-1
quake	quake	4.0.1	tg-login1.ncsa.teragrid.org	2019	PBS-Torque	2.0.0p7	891	538	enabled	4	0	4	2880	-1	-1	-1
gpfs-wan	gpfs-wan	4.0.1	tg-login1.ncsa.teragrid.org	2019	PBS-Torque	2.0.0p7	891	538	enabled	0	0	0	1440	-1	-1	-1

Cluster / Subcluster Overview

Type	Name	UniqueId	Processor		Total Memory	Operating System	SMP Size	Storage Device			TeraGrid Extensions
			Type	Clock Speed				Name	Size	Available Space	Total Nodes
Cluster	NCSA-TeraGrid	NCSA-TG									891
SubCluster	NCSA-TG-IA64CPU13-FASTIO-HIMEM	IA64CPU13-FASTIO-HIMEM	IA-64	1296	4061	Linux 2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	entire-system	353385	91439	128
SubCluster	NCSA-TG-IA64CPU13-FASTIO-LOMEM.ncsa.teragrid.org	IA64CPU13-FASTIO-LOMEM	IA-64	1296	4101	Linux 2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	entire-system	353384	91435	128
SubCluster	NCSA-TG-IA64CPU15-FASTCPU-GPFSWAN.ncsa.teragrid.org	IA64CPU15-FASTCPU-GPFSWAN	IA-64	1496	4106	Linux 2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	entire-system	260036	10620	16
SubCluster	NCSA-TG-IA64CPU15-FASTCPU.ncsa.teragrid.org	IA64CPU15-FASTCPU	IA-64	1496	4106	Linux 2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	entire-system	260036	10619	615
SubCluster	NCSA-TG-IA64CPU13-FASTIO-HIMEM-SPARE	IA64CPU13-FASTIO-HIMEM-SPARE	IA-64	1296	4056	Linux 2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	entire-system	353372	91423	1
SubCluster	NCSA-TG-IA64CPU13-FASTIO-LOMEM-SPARE	IA64CPU13-FASTIO-LOMEM-SPARE	IA-64	1296	4061	Linux 2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	entire-system	353385	91439	1
SubCluster	NCSA-TG-IA64CPU15-PHASE2-FASTCPU-SPARE2	IA64CPU15-PHASE2-FASTCPU-SPARE2	IA-64	1496	4106	Linux 2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	entire-system	260036	10620	2

Hosts in Subcluster NCSA-TG-IA64CPU13-FASTIO-HIMEM

Name	UniqueId	TeraGrid Extensions
		Node Properties
tg-c001.ncsa.teragrid.org	tg-c001	all,ia64-compute,compute,ia64-cpu13,fastio,himem,rack40,clos12,stage
tg-c002.ncsa.teragrid.org	tg-c002	all,ia64-compute,compute,ia64-cpu13,fastio,himem,rack40,clos12
tg-c003.ncsa.teragrid.org	tg-c003	all,ia64-compute,compute,ia64-cpu13,fastio,himem,rack40,clos12
tg-c004.ncsa.teragrid.org	tg-c004	all,ia64-compute,compute,ia64-cpu13,fastio,himem,rack40,clos12
tg-c005.ncsa.teragrid.org	tg-c005	all,ia64-compute,compute,ia64-cpu13,fastio,himem,rack40,clos12



Globus Toolkit Distribution

Globus Projects

MPICH-G2

GridWay

Incubator Mgmt

Java Runtime

Delegation

MyProxy

OGSA-DAI

GT4

C Runtime

CAS

GSI-OpenSSH

Data Rep

Replica Location

Python Runtime

C Sec

GRAM

GridFTP

MDS4

Reliable File Transfer

GT4 Docs

Incubator Projects

GAARDS

MEDICUS

Cog WF

Virt WkSp

GDTE

GridShib

OGRO

UGP

Dyn Acct

Gavia JSC

DDM

Metrics

Introduce

PURSE

HOC-SA

LRMA

WEEP

Gavia MS

SGGC

ServMark

Common Runtime

Security

Execution Mgmt

Data Mgmt

Info Services

Other



Versioning and Support

- Versioning
 - Evens are production (4.0.x, 4.2.x),
 - Odds are development (4.1.x)
- We support this version and the one previous
 - Currently stable version 4.0.4
 - We support 3.2.x and 4.0.x
 - We've also got the 4.1.2 dev release available (1 June '07)



Several "Next" Versions

- 4.0.5 – stable release
 - 100% same interfaces
 - Expected mid June
- 4.1.3 – development release(s)
 - New functionality
 - Expected every 6-8 weeks (mid July)
- 4.2.0 - stable release
 - Tested, documented 4.1.x branch
 - Likely late summer or early fall
 - Discussed on gt-dev@globus.org
- 5.0 – substantial code base change
 - With any luck, not for years :)



Incubator Process in dev.globus

- Entry point for new Globus projects
- Incubator Management Project (IMP)
 - Oversees incubator process from first contact to becoming a Globus project
 - Quarterly reviews of current projects
 - Process being debugged by “Incubator Pioneers”

http://dev.globus.org/wiki/Incubator/Incubator_Process



Current Incubator Projects

dev.globus.org/wiki/Welcome#

Incubator Projects

- Distributed Data Management (DDM)
- Dynamic Accounts
- Gavia-Meta Scheduler
- Gavia- Job Submission Client
- Grid Authentication and Authorization with Reliably Distributed Services (GAARDS)
- Grid Development Tools for Eclipse (GDTE)
- GridShib
- Grid Toolkit Handle System (gt-hs)
- Higher Order Component Service Architecture (HOC-SA)
- Introduce
- Local Resource Manager Adaptors (LRMA)
- Metrics
- MEDICUS
- Open GRid OCSP (Online Certificate Status Protocol)
- Portal-based User Registration Service (PURSe)
- ServMark
- SJTU GridFTP GUI Client (SGGC)
- UCLA Grid Portal Software (UGP)
- WEEP
- Cog Workflow
- Virtual Workspaces



Contribute to an Existing Project

- Contribute code, documentation, design ideas, and feature requests
- Joining the mailing lists
 - *-dev, *-user, *-announce for each project
 - See the project wiki page at dev.globus.org
- Chime in at any time
- Regular contributors can become committers, with a role in defining project directions.



For More Information

- Jennifer Schopf
 - jms@mcs.anl.gov
 - <http://www.mcs.anl.gov/~jms>
- Globus Main Website
 - <http://www.globus.org>
- Dev.globus
 - <http://dev.globus.org>

I should mention some open positions....

- Open positions as part of the Computation Institute, a joint institute between University of Chicago and Argonne National Laboratory
- Scientific research programmer, Grid computing research and applications
 - (requisition # 075544)
- Computer systems programmer, Grid computing research and workflow
 - (requisition # 075338)
- Software and senior software developers
 - (requisition #075287 and 074800)

<http://jobs.uchicago.edu/>, click "Job Opportunities" and search for requisition number listed



Globus at TG '07

June 6, Wednesday

- 9:30-10, Portal-based User Registration Service (PURSe)
- 10-10:30, Resource Discovery on TeraGrid with MDS4
- 10:30-11, Globus 4 GRAM & GRAM Audit
- 11:30-12, GT4 GRAM: A Functionality and Performance Study
- 2:45-3:30, GridFTP/RFT/tgcp

June 7, Thursday

- 9:30-10 GridWay: A Metascheduler for Globus-based Grids