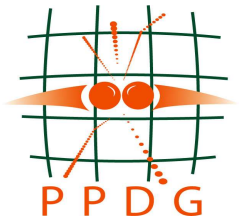

Information Services and Monitoring Projects

Jennifer M. Schopf
Argonne National Lab

James Magowan
IBM UK

LCG Workshop
March 13, 2002

Resource Discovery/Monitoring

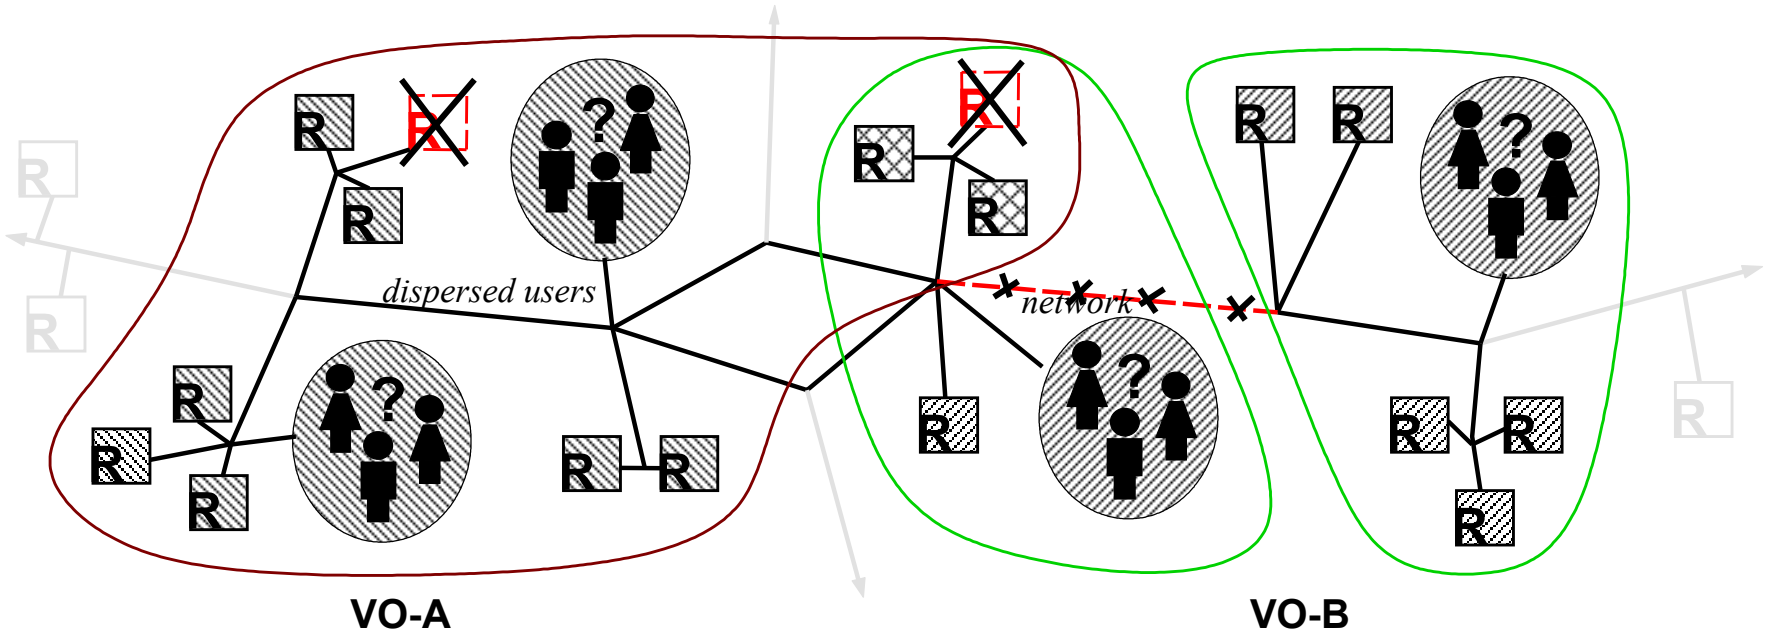

- Large numbers of distributed sensors with different properties, varying status
- Need different views of this information, depending on community membership, security constraints, intended purpose, sensor types, etc

Grid Monitoring and Information Services

- System information is critical to operation of the grid and construction of applications
- We need:
  - Systems that work through failures
  - Interoperablity between grids
  - Agreements/standards to allow continued use of local monitoring (if it already exists)

the globus project
www.globus.org

DataGRID

Overview

- Framework
  - MDS - Globus information service
  - R-GMA - EDG WP-3 information service
- Coordination and Schema Work
  - PPDG/GriPhyN joint monitoring effort
  - GLUE-schema effort
  - Schema work in GGF

Metacomputing Directory Service (MDS)

- Globus Information Service
- Used by iVDGL, GriPhyN, PPDG, EDG, NMI, Grads, etc.
- Requirements and characteristics
  - Uniform, flexible access to information
  - Scalable, efficient access to dynamic data
  - Access to multiple information sources
  - Decentralized maintenance
  - Secure information provision
- Main contact John McGee (mcgee@isi.edu)

MDS Architecture

- Resources run a standard information service (GRIS) that speaks LDAP and provides information about the resource
- GIIS provides a "caching" service
  - Resources register with GIIS
  - GIIS pulls information when requested by a client (when out of date)
- GIIS provides the collective-level indexing/searching function

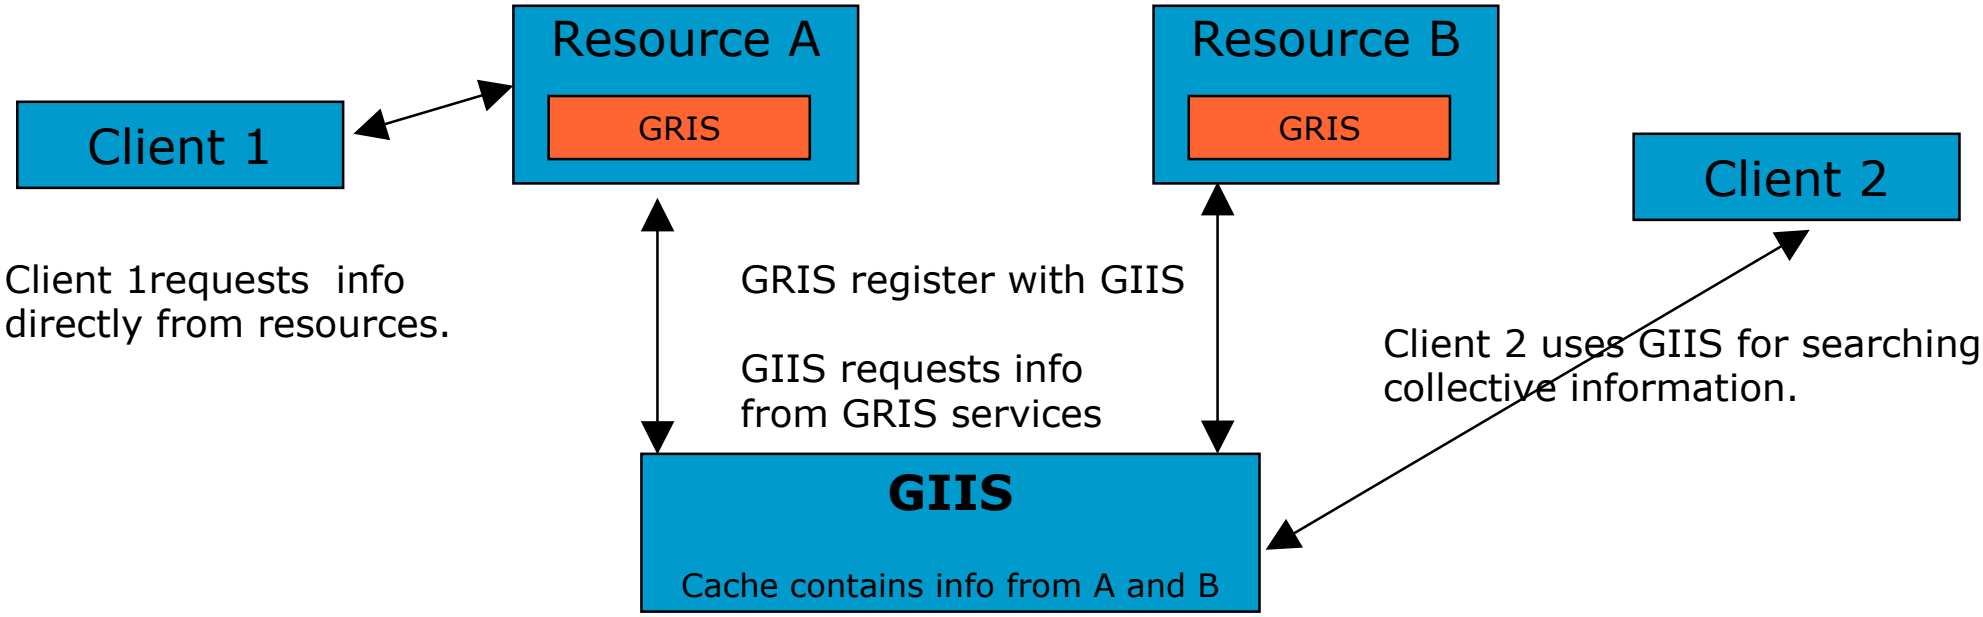

Protocols

- MDS protocols based on LDAP
- Dynamic Registration via Reg. Protocol (GRRP)
  - soft-state protocol
- Resource Inquiry via Info. Protocol (GRIP)
  - Co-located with resource on network
- Resource Discovery (via GRIP or other)
  - Using GRIP allows resource/directory hierarchy
- Also well defined interfaces to add new sensor data

Soft-state Registration

- Periodic notification
  - "Service/resource is available"
  - Expected-frequency metadata
- Automatic directory construction
  - Add new resources to directory
  - Invite resources to join new directory
- Self-cleaning
  - Reduce occurrence of "dead" references

MDS-2 Implementation

- Grid Information Service (GRIS)
  - Provides resource description
  - Modular content gateway
- Grid Index Information Service (GIIS)
  - Provides aggregate directory
  - Hierarchical groups of resources
- NOTE: EDG Ftree
  - flexible backend for MDS by EDG WP3
  - currently in EDG testbed installation of MDS

What's next for MDS (core)

- MDS 2.2 expected in April 2002
  - Will harden the 2.1 code
  - Additional performance improvements
  - Address some of the packaging difficulties that occurred with 2.1
- Notification is currently not included in MDS 2.2 because it was not required as part of the use cases gathered by the joint monitoring group
- The following release will likely be the OGSA-ified version, currently under discussion, and will include notification

What next for MDS Monitoring

- Leading the effort to have common schemas between EDG and US projects
- Defining additional information providers in cooperation with sensor developers
- One example - GridFTP logging data
  - source address, file name, file size, number of parallel streams, TCP buffer size, start and end timestamps, total time consumed, aggregate bandwidth achieved, read/write
  - www-unix.mcs.anl.gov/~vazhkuda/GridFTP-Information-Provider/

the globus project
www.globus.org

DataGRID

Overview

- Framework
  - MDS - Globus information service
  - R-GMA - EDG WP-3 information service
- Coordination and Schema Work
  - PPDG/GriPhyN joint monitoring effort
  - GLUE-schema effort
  - Schema work in GGF

Introduction

- Europe
  - European DataGrid Project
  - DataTAG
  - CrossGrid (new)
- GMA
  - from GGF Perf Working Group
  - several implementations
- R-GMA
- Implications of OGSA
- EDG TestBeds

Grid Monitoring Architecture (GMA)

- We use it not only for monitoring but also as the basis of an information system

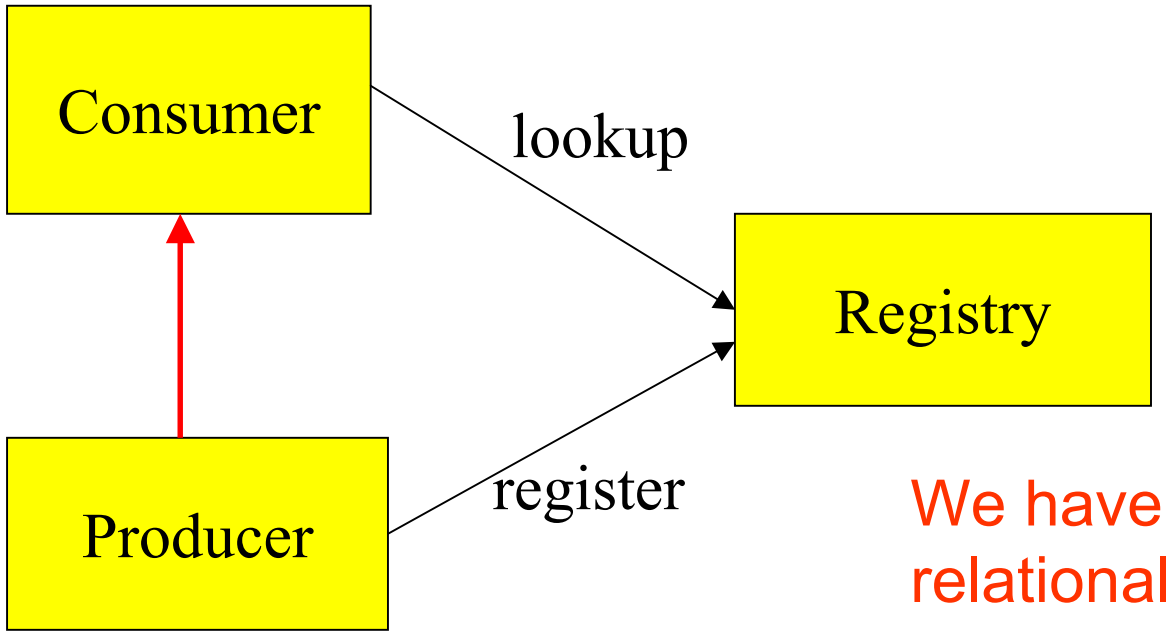

We have chosen a relational implementation

R-GMA

- Not a general distributed RDBMS system, but a way to use the relational model in a distributed environment where ACID properties are not generally important.
- Producers announce: SQL "CREATE TABLE"
  publish: SQL "INSERT"
- Consumers collect: SQL "SELECT"

Table Example

The round trip time from RAL to Lyon as measured pinger was 500 ms on the 9th November at 11:00. :

(RAL, Lyon, Pinger, 1500, 001-11-09T11:00Z)

A set of such tuples could be stored in a table:

NetworkMonitor(NM1, NM2, Tool, Measurement, Time/Date)

Any structured data can be represented in tables in this manner.

Complex queries can be formulated with SQL.

the globus project
www.globus.org

Power behind R-GMA

- The set of tables producers publish constitute a global relational schema
- Individual producers contribute “views” on this global virtual database
- Views can themselves be expressed in SQL
- Queries against the global schema have to be matched with suitable producers
- This soon becomes complicated due to issues of soundness and completeness of producers so we need a mediator.

Schema and views I

CPULoad (Global View)				
Country	Site	Facility	Load	Timestamp
UK	RAL	CDF	0.3	19055711022002
UK	RAL	ATLAS	1.6	19055611022002
UK	GLA	CDF	0.4	19055811022002
UK	GLA	ALICE	0.5	19055611022002
CH	CERN	ALICE	0.9	19055611022002
CH	CERN	CDF	0.6	19055511022002

CPULoad (Producer 1)				
UK	RAL	CDF	0.3	19055711022002
UK	RAL	ATLAS	1.6	19055611022002

CPULoad (Producer 2)				
UK	GLA	CDF	0.4	19055811022002
UK	GLA	ALICE	0.5	19055611022002

CPULoad (Producer3)				
CH	CERN	ATLAS	1.6	19055611022002
CH	CERN	CDF	0.6	19055511022002

Schema and Views II

CPULoad (Producer 1)				
UK	RAL	CDF	0.3	19055711022002
UK	RAL	ATLAS	1.6	19055611022002

WHERE Country = ‘UK’ AND Site = ‘RAL’

CPULoad (Producer 2)				
UK	GLA	CDF	0.4	19055811022002
UK	GLA	ALICE	0.5	19055611022002

WHERE Country = ‘UK’ AND Site = ‘GLA’

SELECT load FROM CPULoad WHERE Site=‘GLA’

SELECT load FROM CPULoad WHERE Facility=‘CDF’

Example

NetworkMonitor			
NM1	**NM2**	**Tool**	**Measurement**
RAL	Lyon	Pinger	1500
RAL	CERN	Pinger	800

StorageElement		
Name	**NM**	**…**
RALS1	RAL	…
CERNS1	CERN	…
LyonS1	Lyon	…

ComputingElement		
Name	**NM**	**…**
RALQ1	RAL	…
CERNQ1	CERN	…
LyonQ1	Lyon	…

```
SELECT S.Name, C.Name, N.Measurement
FROM NetworkMonitor N, ComputingElement C, StorageElement S
WHERE S.NM = N.NM1 AND C.NM = N.NM2 AND MEASUREMENT < 1000
```

S.Name	C.Name	Measurement
RALS1	CERNQ1	800

OGSA - Future

- OGSA and GMA originate from GGF
- R-GMA is conceptually close to OGSA
- OGSA - replaces resources with services
- OGSA has an index for locating services
- R-GMA might (eventually) use OGSA index to provide the functionality needed for the Schema/Registry
- In the short term the R-GMA Registry and Schema will instead be OGSA services
- Notification (Source and Sink) - may provide some of the functionality we require for R-GMA
- We plan to be quick off the mark with this work

R-GMA Architecture

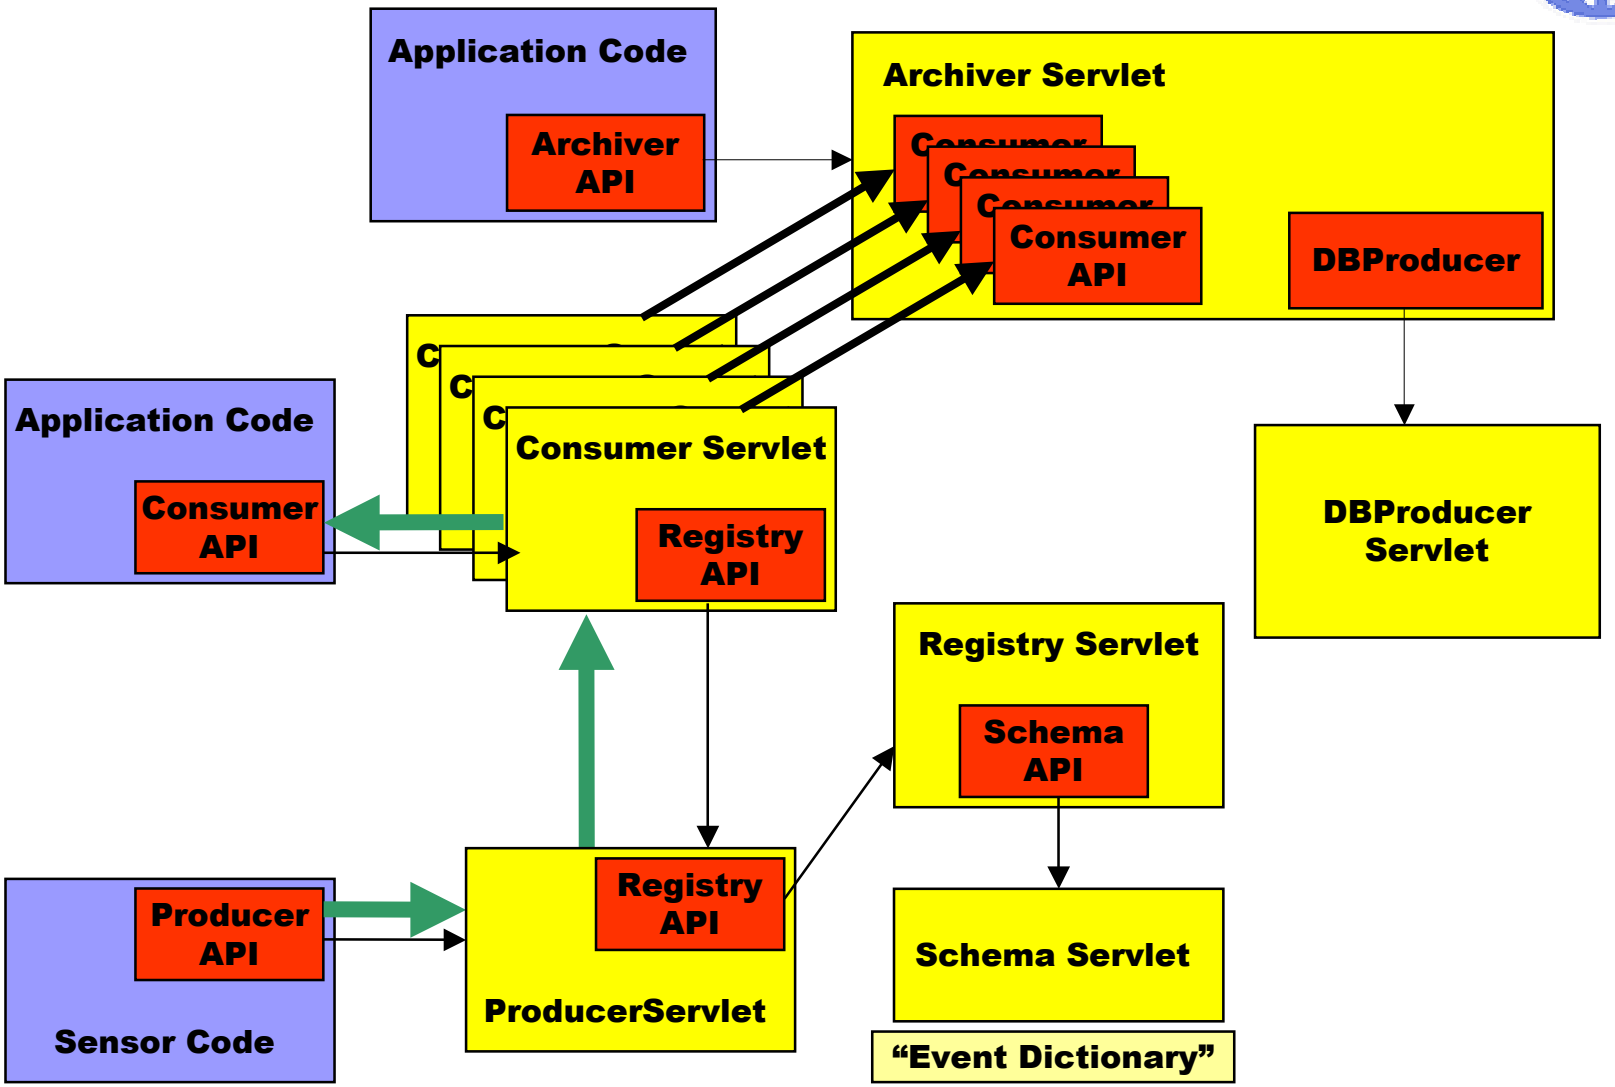

EDG Testbeds

- Existing testbed contains MDS, Ftree and R-GMA
- R-GMA will be used in upcoming testbeds to feed information into MDS
  - Minimum disruption, Resource Broker
  - Use existing backend scripts
  - Test R-GMA with proper workload
  - Gain experience in configuring R-GMA

TestBed 2

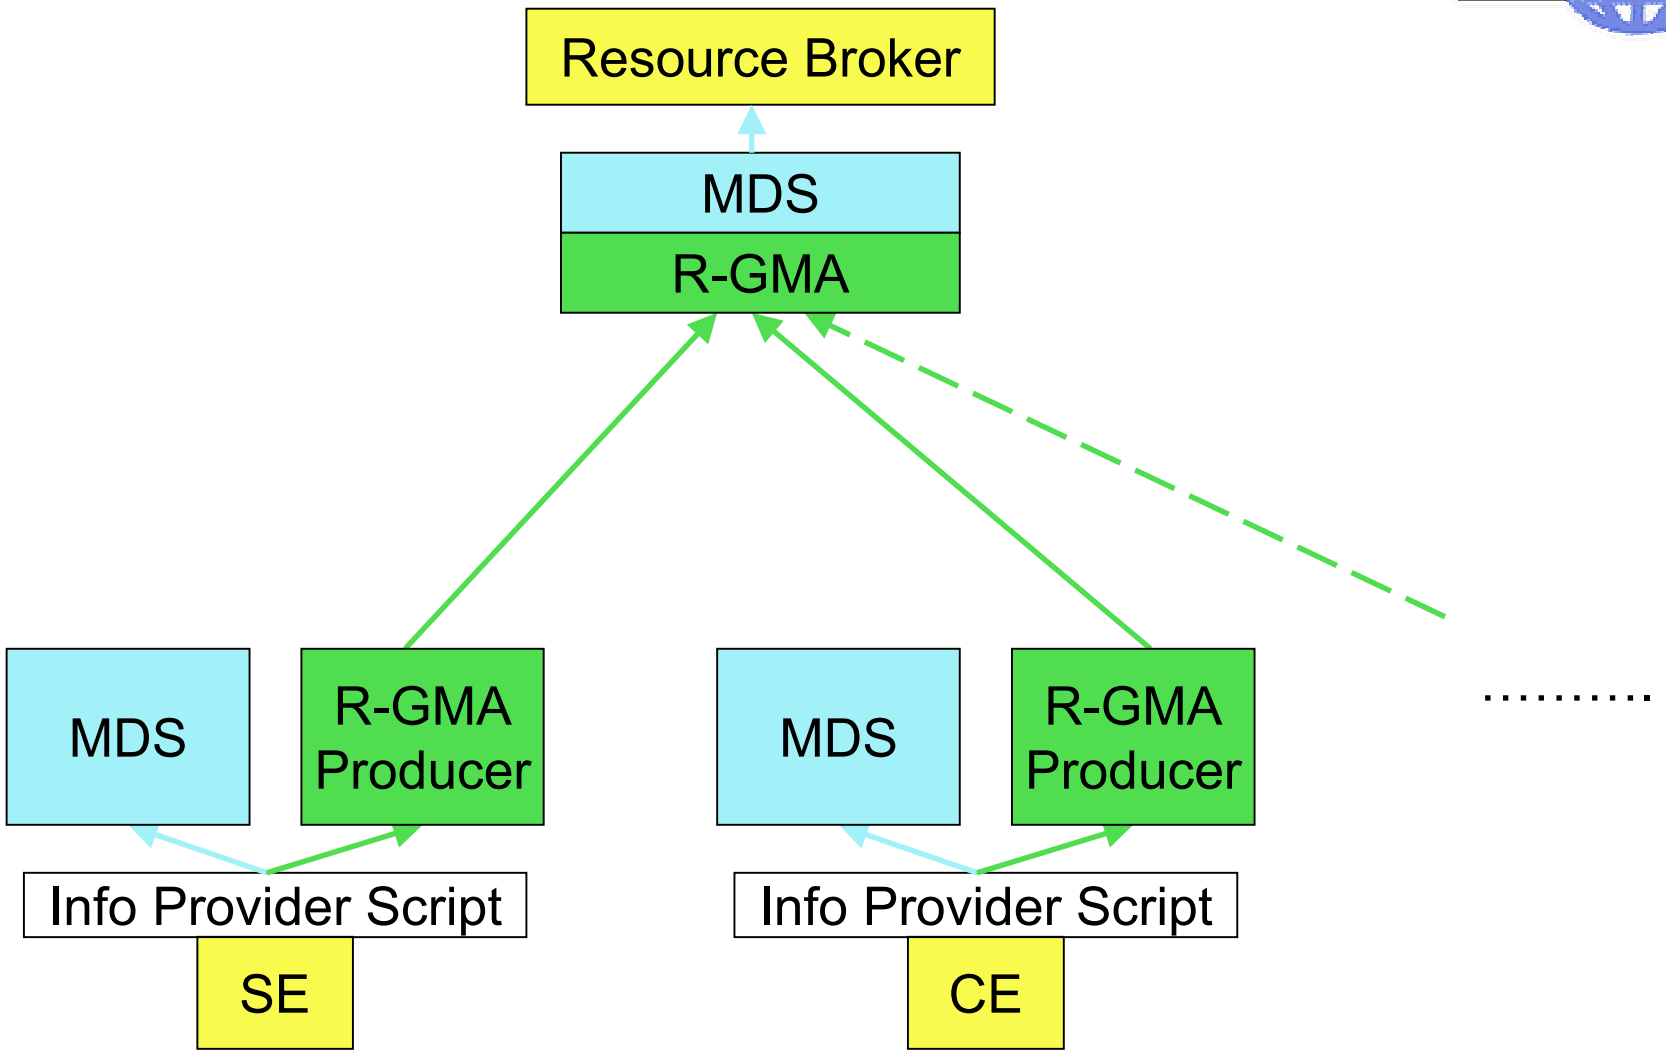
Resource Broker
MDS
R-GMA
MDS
R-GMA Producer
MDS
R-GMA Producer
.....
Info Provider Script
SE
Info Provider Script
CE

the globus project
www.globus.org

DataGrid Testbed 2

- The main info services will be R-GMA
- Security - via Spitfire Mechanism (EDG WP2)
- Shift towards Web Services
- More powerful Mediator
- GRM/Prove using R-GMA as transport
- Improved documentation

the globus project
www.globus.org

Summary

- Moving towards R-GMA from Ftree and MDS
  - R-GMA to initially feed into MDS
- Using GMA from GGF
- Relational view of the world
- As powerful as the mediator
- Very similar to upcoming OGSA
  - servlets not web services at the moment
  - move towards Web Services and OGSA

the globus project
www.globus.org

DataGRID

Overview

- Framework
  - MDS - Globus information service
  - R-GMA - EDG's information service
- Coordination and Schema Work
  - PPDG/GriPhyN joint monitoring effort
  - GLUE-schema effort
  - Schema work in GGF

Joint Monitoring Project PPDG/GriPhyN/iVDGL

- PPDG
  - Delivery of end-to-end applications and integrated production systems
- GriPhyN
  - Virtual Data for Data Grids - resource management, security, fault tolerance research in support of virtual data
- iVDGL
  - Create, operate and evaluate international research lab for data-intensive science
- All using common Grid middleware
  - Globus, Condor, etc.

the globus project
www.globus.org

How can we coordinate our efforts?

- October 2001 formed the joint PPDG/GriPhyN/iVDGL monitoring working group
- Cross-cutting effort between experiments to have a common framework

- Co-chairs: Jennifer Schopf, Brian Tierney
- www.mcs.anl.gov/~jms/pg-monitoring
- mail to pg-monitoring@mcs.anl.gov

Charter

- The goal of this group is to coordinate the monitoring effort between GriPhyN/iVDGL and PPDG. There is a large body of work involving monitoring systems for distributed or Grid resources, and we plan to leverage heavily from them whenever possible.

Gather Use Cases and Requirements

- Defined a template
- 19 use cases from ~9 groups
- Roughly 4 categories
  - Health of system
  - System upgrade evaluation
  - Resource selection
  - Application specific progress
- Working on requirements document now

Deploy a Common Framework

- MDS GIIS for joint project use
- Possibility of using WP-4 network monitoring framework (currently under discussion)?
- Need for joint schemas at the Grid level- see GLUE-schema work in JTB

the globus project
www.globus.org

GLUE-Schema Effort

- Part of HICB/JTB GLUE framework
- To address need to common schemas between projects
  - framework independent
  - something to translate into, not a requirement within fabric layer
- Co-chairs: Cristina Vistoli and Brian Tierney
- maillist: glue-schema@hicb.org
- www.hicb.org/glue/glue-schema/schema.html

Problems with Schema's

- Naming MetaData implies a structure, e.g. relational, hierarchical, etc
- Attribute/Measurement units
- Attribute/Measurement naming
- Object/Entity naming
- Representation of a schema (LDAP, SQL, XML, CIM, etc)
- A lot of schemas out there
  - information, monitoring, scheduling, accounting, security, and more!

Goal

- Define a minimum common schema requirement for interoperability
  - Compute Elements, Network Elements, Storage Elements
  - EDG, DataTag, PPDG, GriPhyN, iVDGL, etc.!
- Proposed first step: creation of a table comparing some of the initial approaches
  - First draft sent to mailing list Monday

Example

EDG Object class	EDG attribute	desc.	required	single/mutli	type	globus object class	globus attribute	desc.	required	single/mutli
Computing Element	Architecture	The architecture of the hosts composing the CE - assumed all are identical	y	single	cis	MdsComputer	Mds-Computer-isa	informally names the Instruction Set Architecture (ISA) of the computing element	y	multi
						MdsComputer	Mds-Computer-platform	informally describes the platform type of the computing element	y	multi
						MdsCpu	Mds-Cpu-vendor	informally names the CPU vendor	y	multi
						MdsCpu	Mds-Cpu-model	informally names the CPU model	y	multi
						MdsCpu	Mds-Cpu-version	informally names the CPU version or stepping	y	multi
						MdsCpu	Mds-Cpu-features	informally names optional CPU features	n	multi
						MdsCpu	Mds-Cpu-speedMHz	inficates the clock speed of a CPU	n	multi

the globus project
www.globus.org

DataGRID

Overview

- Framework
  - MDS - Globus information service
  - R-GMA - EDG WP-3 information service
- Coordination and Schema Work
  - PPDG/GriPhyN joint monitoring effort
  - GLUE-schema effort
  - Schema work in GGF

DAMED WG

- Discovery And Monitoring Event Description Working Group
- Chairs
  - > Jennifer Schopf, ANL
  - > James Magowan, IBM
- Dammed if we do
  - Not everyone will be happy
- Dammed if we don't
  - Never reach our goal of seamless interoperability of grids (one big grid e.g. internet)

Charter

- Define a basic set of monitoring event descriptions
  - information (attributes) associated with a particular data element
  - conventions for the representation of the value associated with it.
- Develop standard representations of the most widely used measurement values (the "top N".)
- Emergence of a set of conventions and recommendations that will ease the task of defining richer, domain-specific schemas

the globus project
www.globus.org

Milestones

- Research existing schemas
  - draft of comparison of some schemas currently in process
  - Completion by April 2002
- Develop "English" descriptions for the "top N" monitoring elements.
  - Completion by May 2002
- Map English definitions to basic technologies including LDAP, SQL, XML.
  - Draft for discussion by GGF-5

Simple Steps

- Links to current Grid schema effort
- Units and names for basic top n attributes/Measurements
- More complex entities involving more than one measurement/attribute.
- Solve Entity Identity Problems

the globus project
www.globus.org

How to get involved

- Subscribe to mailing list
  - damed - damed-wg@gridforum.org
- WebSite
  - http://www-didc.lbl.gov/damed
- EDG input through WP3 to DAMED
  - wp3 (EDG) - edgschemas@jiscmail.ac.uk

Other efforts

- DataTAG WP4 (Ghiselli, Vistoli)
  - http://datatag.web.cern.ch/datatag/
- Distributed Monitoring Framework (Tierney), Web 100
  - http://www-didc.lbl.gov/DMF/
- Network Weather Service (Wolski)
  - http://nws.cs.ucsb.edu/
- PingER (Cotrell)
  - http://www-iepm.slac.stanford.edu/pinger/
- CrossGrid WP3
  - http://www.crossgrid.org/
- HENP working group in Internet2
  - http://www.internet2.edu/henp/
- PyGMA (LBL), NASA-GMA (Smith, NASA Ames)
- and more!

the globus project
www.globus.org

Additional Information

- MDS
  - www.globus.org/gt2/mds2.1/
- PPDG/GriPhyN/iVDGL joint effort
  - www.mcs.anl.gov/~jms/pg-monitoring/
- Glue-schema
  - www.hicb.org/glue/glue-schema/schema.html
- R-GMA
  - hepunx.rl.ac.uk/grid/wp3/
- damed-wg
  - http://www-didc.lbl.gov/damed