Grid Data Management

Seminar on Grid Computing

Amira Zaki

International School of Informatics
Hagenberg, Austria

29/01/2009

Outline

The Grid

Is a service for sharing computer power and data storage capacity over the Internet. The Grid goes well beyond simple communication between computers, and aims ultimately to turn the global network of computers into one vast computational resource.

Why Data Management?

- Grid systems are formed to solve problems too complex or too expensive to solve with local resources. Problem solving is concerned with the consumption and production of information. Thus information on the Grid is important
- The Grid is a complex and information-rich environment. Grid middleware uses information about the availability of services; their purpose; ways in which they can be combined and configured; and how they are discovered, invoked and evolve. Thus information about the Grid is also important

Why Data Management?

- Grid systems are formed to solve problems too complex or too expensive to solve with local resources. Problem solving is concerned with the consumption and production of information. Thus information on the Grid is important
- The Grid is a complex and information-rich environment. Grid middleware uses information about the availability of services; their purpose; ways in which they can be combined and configured; and how they are discovered, invoked and evolve. Thus information about the Grid is also important

Data Management

Definition

Mechanisms to transfer, public, replicate, discover, share and analyse data

Example

Business applications (ex.. online business services) need to maintain database consistency worldwide, manage data replication, facilitate data discovery, and respond dynamically to changed in the load applied to databases by users

Data Management

Definition

Mechanisms to transfer, public, replicate, discover, share and analyse data

Example

Business applications (ex.. online business services) need to maintain database consistency worldwide, manage data replication, facilitate data discovery, and respond dynamically to changed in the load applied to databases by users

Challenges of Data Management

- Diverse usage scenarios
- Heterogeneity at all system levels
- Performance demands associated with access, manipulation and analysis of large quantities of data

Key Concepts

- **Data Movement**
- **Data replication**

Key Concepts

- **Data Movement**: between storage systems or between programs and data storage.
  Challenge: large size of datasets and wide area data transfer delays
- **Data replication**

Key Concepts

- **Data Movement**: between storage systems or between programs and data storage.
  Challenge: large size of datasets and wide area data transfer delays
- **Data replication**: creating replicas to reduce access latency, maintain local control over necessary data, improve reliability and local balancing.
  Challenge: locating existing replicas, selecting among available replicas, proactively replicating data items to satisfy demand, satisfying consistency requirements

Movement of Grid Data

Involves moving data between storage systems or between programs and data storage

In Globus Toolkit:

There are two components related to data movement:

- GridFTP tools
- Reliable File Transfer (RFT) service

Grid FTP explained

- Protocol for secure, robust, fast and efficient transfer of bulky data
- It has the following functionalities:
  - Grid Security Infrastructure support
  - Third-party control of data transfer
  - Parallel data transfer
  - Striped data transfer
  - Partial file transfer
  - Automatic negotiation of TCP buffer/window sizes
  - Support for reliable and restartable data transfer

Grid FTP explained

- Protocol for secure, robust, fast and efficient transfer of bulky data
- It has the following functionalities:
  - Grid Security Infrastructure support
  - Third-party control of data transfer
  - Parallel data transfer
  - Striped data transfer
  - Partial file transfer
  - Automatic negotiation of TCP buffer/window sizes
  - Support for reliable and restartable data transfer

Third-party Transfer Operation

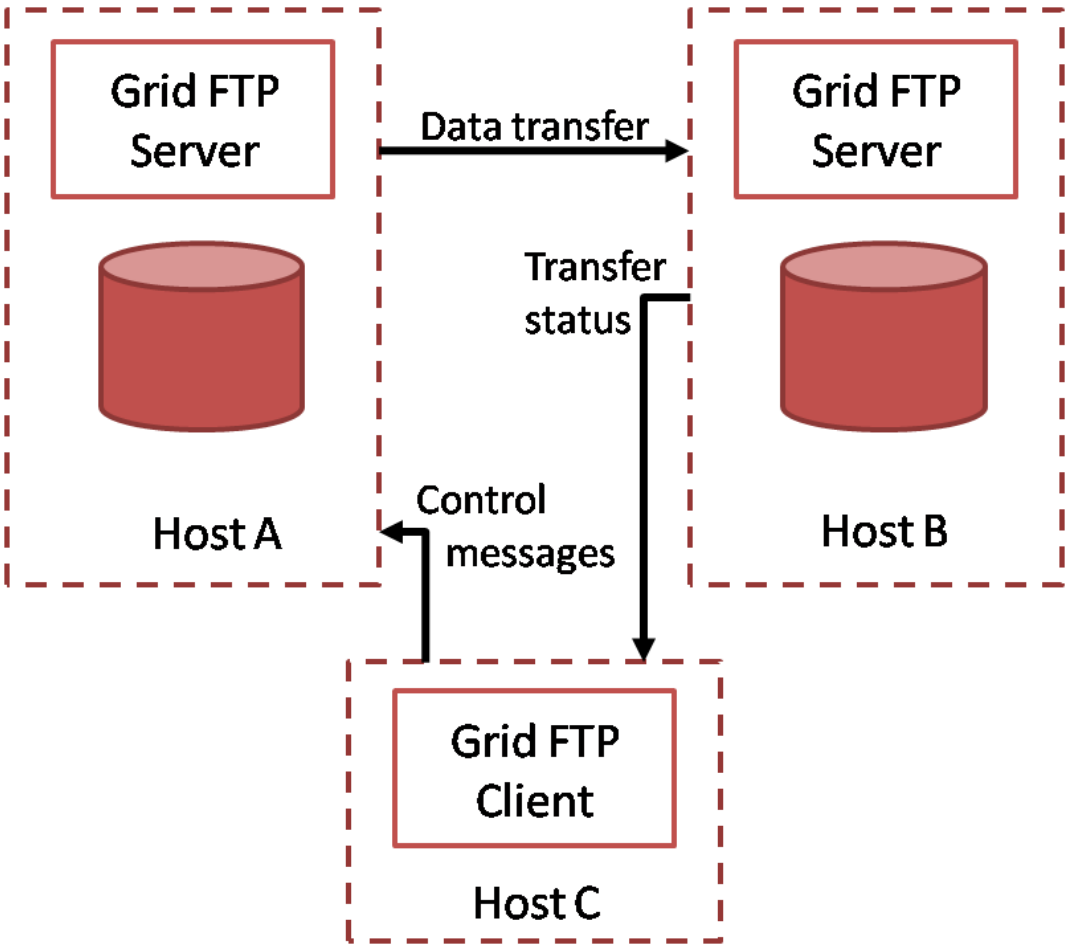

Providing data

To make data available to others, you need to install a server on a host that can access that data and make sure that there is an appropriate Data Storage Interface (DSI) available for the storage system holding the data
The basic syntax for `globus-gridftp-server` is:

```
globus-gridftp-server [optional command line switches]
```

File Transfer

GridFTP client is used to access data that others have made available via many protocols (http, https, ftp, gsiftp, and file). It is not an interactive client, but a command line interface, suitable for scripting.

Basic Notation

```
globus-url-copy -vb -tcp-bs 2097152 -p 4 src_url dest_url
```

where

`-vb`	verbose mode (specifying performance, transferred bytes)
`-tcp-bs`	size in bytes of TCP buffer used by FTP channels
`-p`	number of parallel data connections
`-url`	`file:///path/to/my/file`
	`gsiftp://hostname/path/to/remote/file`

Putting Files

Moving the file `/tmp/foo` from a file system accessible to your client to a file name `/tmp/bar` on a host named `remote.machine.my.edu` running a GridFTP server:

```
globus-url-copy -vb -tcp-bs 2097152 -p 4
   file:///tmp/foo gsiftp://remote.machine.my.edu/tmp/bar
```

Getting Files

Moving a file from a server to your file system, would just reverse the source and destination URLs

```
globus-url-copy -vb -tcp-bs 2097152 -p 4
   gsiftp://remote.machine.my.edu/tmp/bar file:///tmp/foo
```

Third-party Transfer

To move a file between two GridFTP servers (third-party transfer), both URLs would use `gsiftp:` as the protocol:

```
globus-url-copy -vb -tcp-bs 2097152 -p 4
    gsiftp://other.machine.my.edu/tmp/foo
    gsiftp://remote.machine.my.edu/tmp/bar
```

Pipelining

Pipelining allows the client to have many outstanding, unacknowledged transfer commands at once. Instead of being forced to wait for the finished response message, the client is free to send transfer commands at any time.
To enable Pipelining:

```
globus-url-copy -pp
```

Multicasting

To transfer a single file to many destinations in a multicast, use `-mc` option:
Pipelining is enabled by using the -pp option:

```
globus-url-copy -vb -tcp-bs 209715 -p 4 -mc filename src_url
```

The filename must contain a line-separated list of destination urls:

```
gsiftp://localhost:5000/home/user/tst1
gsiftp://localhost:5000/home/user/tst3
gsiftp://localhost:5000/home/user/tst4
```

Or:

```
globus-url-copy -MC multicast.file src_url
```

Multicasting

To transfer a single file to many destinations in a multicast, use `-mc` option:
Pipelining is enabled by using the -pp option:

```
globus-url-copy -vb -tcp-bs 209715 -p 4 -mc filename src_url
```

The filename must contain a line-separated list of destination urls:

```
gsiftp://localhost:5000/home/user/tst1
gsiftp://localhost:5000/home/user/tst3
gsiftp://localhost:5000/home/user/tst4
```

Or:

```
globus-url-copy -MC multicast.file src_url
```

GridFTP Limitations

Although powerful, the GridFTP has the following limitations:

- GridFTP protocol is not a web service protocol
- GridFTP requires that the client maintain an open socket connection to the server throughout the transfer

RFT explained

- Provids a "job scheduler"-like functionality
- Enables clients to transfer files or directories and to monitor transfer status or to subscribe to receive notifications if transfer status changes

Submitting a Transfer

To submit a transfer request the user must first create a transfer file. Each line of this ASCII text file is a source/destination URL pair. There can be any number of lines per file.
An example file follows:

```
gsiftp://localhost:2811/etc/group  gsiftp://localhost:2811/tmp/test_crft
gsiftp://ftp.globus.org:2811/pub/README gsiftp://myhost.here/home/user/file
```

Submitting a Transfer

Once the transfer file is created `globus-crft` can be used in a variety of ways to transfer a file. The most simple is the blocking transfer:

```
globus-crft -c -s -m -vb -f <trans_file> -e <contact>
```

where

`-c`	Create a new RFT server
`-s`	Submit the transfer request
`-m`	Monitor the transfers: client blocks until all transfers complete
`-vb`	Display verbose output
`-f <trans_file>`	Pointer to the transfer file described earlier
`-e <contact>`	Contact string form: `https://hostname.com:8443/wsrf/services/`

Non-blocking Transfer

The client can do non-blocking RFT submission. It can submit an RFT request and then terminate, returning later to monitor the status of the request. To accomplish this the client saves the EPR of the newly created RFT service to disk.

```
globus-crft -c -s -f <trans_file> -e <contact>
                -ef <epr_output)file>
```

At some point later the client uses this same file to monitor the state of the transfer:

```
globus-crft -ef <epr input file> --getOverallStatus
```

Cleaning Up

Once a transfer request completes, the user should destroy the resources associated with it. If the user stored the EPR of the service it created, this can be done with:

```
globus-crft -ef <epr input file> --destroy
```

Replication of Grid Data

Replicas: Copies of data available in various locations in order to reduce access latency, and improve reliability and load balancing

Replica Management Services include:

- creating new copies of files
- registering these new copies in a Replica Catalog
- finding all existing copies of a file when a cataolg is queried

In Globus Toolkit:

There is one component related to data replication:

- Replica Location Service (RLS)

Replication of Grid Data

Replicas: Copies of data available in various locations in order to reduce access latency, and improve reliability and load balancing

Replica Management Services include:

- creating new copies of files
- registering these new copies in a Replica Catalog
- finding all existing copies of a file when a cataolg is queried

In Globus Toolkit:

There is one component related to data replication:

- Replica Location Service (RLS)

Logical vs. Physical file names

- Logical file name: a unique identifier for the contents of a file
- Physical file name: the location of a copy f the file on a stroage system

The purpose of the Replica Management Service is to map a unique logical file name to a possibly different physical name for the file on a particular storage device.

Logical vs. Physical file names

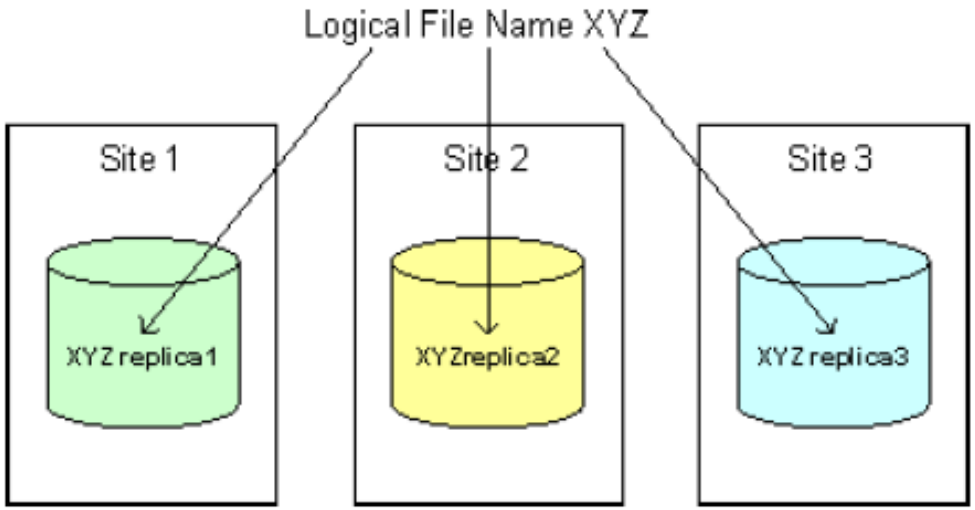

Features

Features of a Replica Management Service include:

- Separation of Replication and Matadata information
- Replica Management Service Consistency
- Rollback
- Distributed locking mechanism

Data Selection Scenario

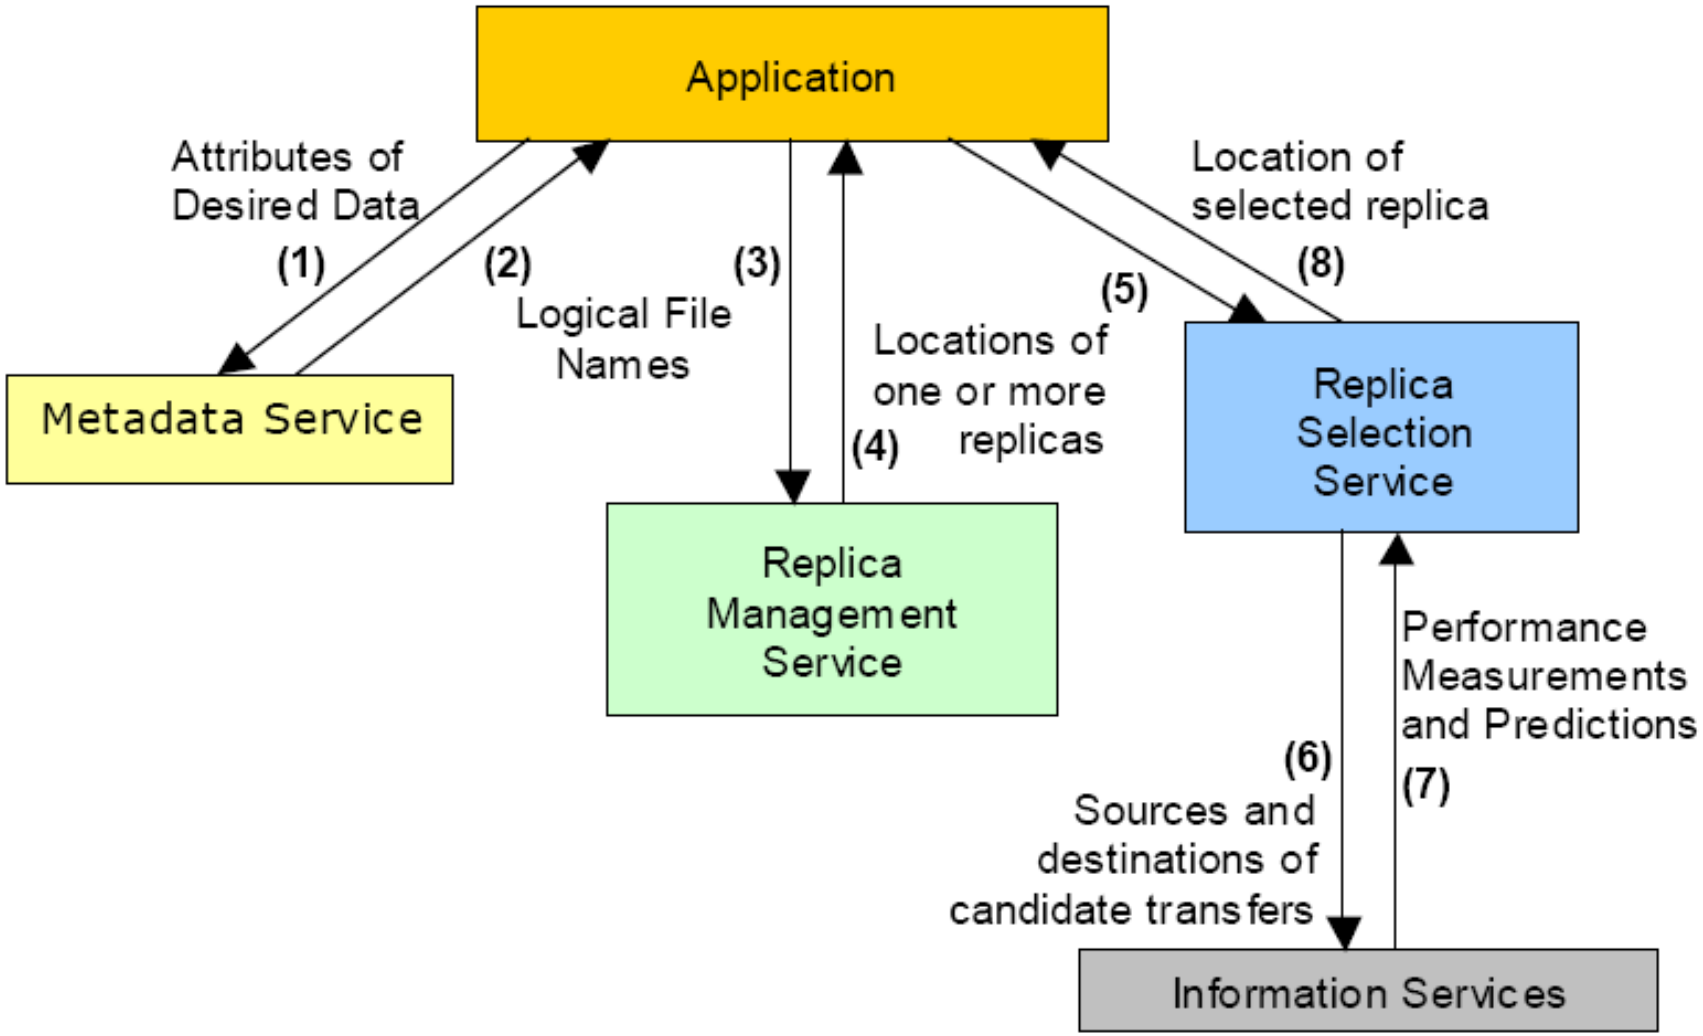

RLS explained

- RLS is a distributed registry, meaning that it may consist of multiple servers at different sites
- RLS provides ability keep track of one or more copies (replicas) of files in a Grid environment
- Helpful to find where existing files are located in the Grid

`globus-rls-server` can be configured to be a Location Replica Catalog (LRC) server, which manages Logical FileName (LFN) to Physical FileName (PFN) mappings in a database

Creating Replica Location Mappings

When the RLS server is first installed its database of replica location information will be empty, as expected. To create a replica location mapping:

```
globus-rls-cli create my-logical-name-1
                 url-for-target-name-1 rls://localhost
```

Adding Replica Location Mappings

To map additional target names to a logical name created by the previously described create command:

```
globus-rls-cli add my-logical-name-1
                  url-for-target-name-2 rls://localhost
```

Querying Replica Location Mappings

Once your RLS server is populated with replica location mappings, you can query the server for useful information using

```
query lrc lfn my-logical-name-1 rls://localhost
```

The result obtained for the commands so far results in:

```
my-logical-name-1: url-for-target-name-1
my-logical-name-1: url-for-target-name-2
```

Querying Replica Location Mappings

Once your RLS server is populated with replica location mappings, you can query the server for useful information using

```
query lrc lfn my-logical-name-1 rls://localhost
```

The result obtained for the commands so far results in:

```
my-logical-name-1: url-for-target-name-1
my-logical-name-1: url-for-target-name-2
```

Deleting Replica Location Mappings

To remove unwanted replica location mappings from the RLS server (i.e. associations between the specified logical name and the specified target name):

```
delete my-logical-name-1 url-for-target-name-1
       rls://localhost
```

Only when all mapping associations for a given logical name are eliminated will the logical name be deleted from the RLS server.

Conclusion

- Collecting, managing, sharing, publishing and exploiting large volumes of data is a challenge
- Globus toolkit offers GridFTP and RFT as a means to move grid data
- Globus toolkit offers RLS as a means to replicate grid data

References


Bill Allcock, Joe Bester, John Bresnahan, Ann L. Chervenak, Ian Foster, Carl Kesselman, Sam Meder, Veronika Nefedova, Darcy Quesnel, and Steven Tuecke.
Data Management and Transfer in High Performance Computational Grid Environments.
Parallel Computing Journal, 28:749–771, 2002.


Ian Foster and Carl Kesselman.
The Grid 2: Blueprint for a New Computing Infrastructure.
Elsevier, 2004.


Grid Data Management.
http://www.ci.uchicago.edu/osgedu/schools/gridlab/data.html.


GT 4.2 Developer's Guide.
http://www.globus.org/toolkit/docs/4.2/4.2.0/developer/.


GT 4.2 Grid FTP User's Guide.
http://www.globus.org/toolkit/docs/4.2/4.2.0/data/gridftp/user/index.html.


GT 4.2 RFT User's Guide.
http://www.globus.org/toolkit/docs/4.2/4.2.0/data/rft/user/index.html.


GT 4.2 RLS User's Guide.
http://www.globus.org/toolkit/docs/4.2/4.2.0/data/rls/user/.


GT 4.2 User's Guide.
http://www-unix.globus.org/toolkit/docs/4.2/4.2.1/user/.