

Grid Computing

José M. Vidal

Thu Feb 26 15:56:37 EST 2004

We give a very high-level overview of grid computing and the Globus toolkit. This talk is based on:

- Ian Foster, Carl Kesselman, Jeffrey M. Nick, and Steven Tuecke. Grid Services for Distributed System Integration¹, *IEEE Computer*, 35(6), 2002.
- Ian Foster, Carl Kesselman, and Steven Tuecke. The Anatomy of the Grid: Enabling Scalable Virtual Organizations², 2001.
- Ian Foster, Carl Kesselman, Jeffrey M. Nick, and Steven Tuecke. The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration³, January, 2002.
- Introduction to Grid Computing and the Globus Toolkit⁴
- M. Mitchell Waldrop. Grid Computing.⁵ *Technology Review*, May; 105(4):30–37, 2002.

1 Introduction

- Computing power as a utility.
- A grid ties many computers into a virtual machine.
- A true grid must work
 - Across institutional barriers.
 - Indifferent to platform differences.
 - Robust in the face of network failures.
 - Provide security and authentication.

Note:

As you read in the papers, interest on the Grid stems largely from scientists with large computational problems which cannot be solved on any of the machines they can afford to buy. The Grid allows one person or institution to easily use the computational power of many others. This ability is alluring because the fact is that many of these high-end machines sit idle for large periods of time (wasted CPU power). Also, in a similar way, the Grid allows the sharing of other resources, namely research instruments. For example, it can allow a research group to access a particle accelerator, telescope, seismic sensor, etc. that is located at the other side of the world.

2 Existing Grids


- Many grids exist already. Mostly for scientific computing.
 - Access Grid⁶. ANL.
 - Information Power Grid⁷- Nasa.

- Unicore⁸
- TeraGrid⁹- NSF.
- *”TeraGrid is a multi-year effort to build and deploy the world’s largest, most comprehensive, distributed infrastructure for open scientific research. By 2004, the TeraGrid will include 20 teraflops of computing power distributed at five sites, facilities capable of managing and storing nearly 1 petabyte of data, high-resolution visualization environments, and toolkits for grid computing. Four new TeraGrid sites, announced in September 2003, will add more scientific instruments, large datasets, and additional computing power and storage capacity to the system. All the components will be tightly integrated and connected through a network that operates at 40 gigabits per second.”*
- Dr. Kameron heads a laboratory¹⁰ which has recently joined the Grid¹¹.

3 Definitions

- A **resource** is an entity to be shared: computer, storage, data, software. Defined by its interface.
- A **network enabled service** is a network protocol that provides us with some services: FTP server, Web server.
- A **virtual organization (VO)** is a dynamic grouping of resources and people in order to solve a problem. The Grid is meant to support VOs.
- **Virtualization** enables consistent resource access across heterogeneous platforms.
- The Globus Toolkit¹² is the de-facto standard implementation of a Grid.

4 Grid Architecture



- The resource protocol layer facilitates sharing of resources.
- The collective layer has services that involve the coordinate use of multiple resources.
- The connectivity protocols enable connections.
- The fabric layer refers to a set of resource types. This layer interfaces to local control.

4.1 Fabric Layer

- Provides the resources that are shared by the Grid: CPU time, storage, sensors.
- Fabric allows user to *ask* about available resources and provides *resource management* mechanism. For example
 - *Computational resources*: How fast is the hardware? Start, monitor, stop.
 - *Storage*: Can I read/write? how much? how fast? Read, write.
 - *Network*: What is the load? Reserve bandwidth.
 - *Code repositories*: CVS.
 - *Catalogs*: relational databases.
- **Globus** uses whatever is already on the machine, but provides its own software if there is none.

4.2 Connectivity

- Defines the core communication and authentication protocols required for Grid-specific network transactions.
- Communication uses TCP/IP, DNS.
- Authentication demands:
 - Single sign on.
 - Delegation: let this program access those files.
 - Integration: with Kerberos and Unix.
 - User-based trust relationships: If I can use A and B then my program should be able to use both together.
- **Globus** implements the public-key Grid Security Infrastructure (GSI) protocols which extend TLS.

4.3 Resource Layer

- There are **information protocols** that tells us about the state of the resource and **management protocols** that negotiate access to a resource.
- **Globus** provides:
 1. The Grid Resource Information Protocol (GRIP), based on LDAP¹³, provides information about resources. GRRP is for resource registration.
 2. The Grid Resource Access and Management (GRAM) protocol is HTTP-based and used for allocation of computational resource and monitoring and controlling them.
 3. The GridFTP is an extension of FTP¹⁴ allows partial file access and management of parallelism for high-speed transfers.
 4. LDAP.

4.4 Collective Layer

- These might include:
 1. Directory services. eg., resource-level GRRP and GRIP protocols.
 2. Co-allocation, scheduling, and brokering services. eg., DRM broker.
 3. Monitoring and diagnostics services.
 4. Data replication services.
 5. Grid-enabled programming systems. eg., Grid-enabled implementations of the Message Passing Interface and CORBA.
 6. Workload management systems and collaboration frameworks.
 7. Software discovery services.
 8. Community authorization servers.
 9. Community accounting and payment services.
 10. Collaboratory services.

5 Grid Security Infrastructure

- GSI provides:
- Uniform authentication, authorization, and message protection mechanisms.
- Single sign-on, delegation, identity mapping.
- Uses public key technology (SSL).
- Uses a supporting infrastructure of: certificate authorities, certificate and key management software.

6 How to Program?

- There is no language or simple library to link to! Instead you must obey all the protocols or, at least, the ones you need (no security?).
- Program must already be distributed, for example:
 - Grid-enabled CORBA.
 - MPICH-G2: grid enabled implementation of Message Passing Interface.
 - Condor-G: workflow management.
 - Legion: object models for grid computing.
 - Cactus: grid-aware numerical solver framework.
- Then, use it along with the globus toolkit.

7 Grid Resource Allocation Management

- The GRAM protocol allows programs to be started on remote resources, even if different.
- The Resource Specification Language (RSL) communicates requirements. RSL is similar to LDAP filters:

```
& (executable=myprog)
  ( - (&(count=5)(memory >= 74))
    (&(count=10)(memory >= 32)))
```

Create 5 instances of myprog on a machine that has at least 64MB memory or 10 instances on a machine with at least 32MB.

- We can also specify multiple resource needs:

```
+ (&(count=5)(memory>=64)
  (executable=p1))
  (&(network=atm)(executable=p2))
```

Execute 5 instances of p1 on a machine with 64MB or more, and execute p2 on a machine with an ATM connection.

- To submit the programs you use the commands:
 - globus-job-run: interactive jobs.

- globus-job-submit: batch/offline jobs.
- globusrun: flexible scripting infrastructure.
- GRAM-1 used HTTP-based RPC: request that a job be started/stopped. Returns a job contract.
- GRAM-2 uses SOAP.

8 Information Services

- The Grid Resource Information Service (GRIS) implements GRIP and GRRP.
- Use LDAP.
- GRIS tells us: machine load, process information, storage space, etc. It also lets us find specific machines: find machine with shortest queue.
- You can use the command

```
grid-info-host-search [options] filter [attributes]
```

to find suitable hosts.

9 GridFTP

- A superset of FTP.
- Implements often unused features of FTP: GSS binding, extended directory listing, simple restart.
- Extends FTp by adding: stripe/parallel data channels, partial file, automatic and manual TCP buffer setting, progress monitoring, extended restart.
- Globus also implements custom libraries, clients, and servers. All optimized for high-performance (big files fast).
 - gsi-ncftp
 - gsi-wuftp

10 Conclusion

- Globus toolkit is a long way from power grid analogy. Not surprising since automatic parallelization is unsolved problem.
- Mostly a collection of services: aggregation of existing technologies with some extra glue.
- New Open Grid Services Architecture (OGSA) hopes to snarf in more Web Services (SOAP, UDDI, WSDL) technologies.
- Joining a Grid requires personal invitation.
- Simple API is possible, but Grid rejects it because of need to support legacy software. *The Grid could be very simple if we could design it from scratch!*
- Hard to set up, according to Kirk.

Notes

- ¹<http://www.globus.org/research/papers/ieee-cs-2.pdf>
- ²<http://www.globus.org/research/papers/anatomy.pdf>
- ³<http://www.globus.org/research/papers/ogsa.pdf>
- ⁴<http://www.globus.org/training/grids-and-globus-toolkit/index.html>
- ⁵<http://jmvidal.cse.sc.edu/library/waldrop02a.pdf>
- ⁶<http://www-fp.mcs.anl.gov/fl/accessgrid>
- ⁷<http://www.ipg.nasa.gov>
- ⁸<http://www.unicore.de>
- ⁹<http://www.teragrid.org>
- ¹⁰<http://scape.cse.sc.edu/>
- ¹¹<http://uscnews.sc.edu/engr211.html>
- ¹²<http://www.globus.org>
- ¹³<http://www.wikipedia.org/wiki/LDAP>
- ¹⁴<http://www.wikipedia.org/wiki/Ftp>

This talk is available at <http://jmvidal.cse.sc.edu/talks/gridcomputing>

Copyright © 2004 Jose M Vidal. All rights reserved.