

Accounting and Accountability for Distributed and Grid Systems

Laura F. McGinnis
Pittsburgh Supercomputing
Center
lfm@psc.edu

William Thigpen
NASA Ames
bthigpen@arc.nasa.gov

Thomas J. Hacker
University of Michigan
hacker@umich.edu

Abstract

While the advent of distributed and grid computing systems will open new opportunities for scientific exploration, the reality of such implementations could prove to be a system administrator's nightmare. A lot of effort is being spent on identifying and resolving the obvious problems of security, scheduling, authentication and authorization. Lurking in the background, though, are the largely unaddressed issues of accountability and usage accounting:

- *Mapping resource usage to resource users;*
- *Defining usage economies or methods for resource exchange;*
- *Describing implementation standards that minimize and compartmentalize the tasks required for a site to participate in a grid.*

For an accounting system to be functional in a grid environment, it needs to be decentralized, scalable and flexible. It must have a minimum impact on local accounting and should not make any limiting assumptions about whether accounting is done by user, group, project, or site. The requirements on the remote site will be to track the resources used by the requesting job and then pass this information back to the requesting site in some standardized format. At the requesting site, the information can then be accrued as needed for local requirements. A distributed allocation and accounting approach, using a consumer/supplier or client/server structure will work across multiple sites and satisfy the needs of the participating administrative and policy domains.

A survey of current practices [1] shows that the only thing many sites have in common is their diversity. The Distributed Accounting Working Group, a research group in the Global Grid Forum has discovered that

- Most HPC sites are already supporting a variety of resources. This makes them “mini-grids”, at least as far as current practices go.

- Resource allocation requests are reviewed before they are granted. No one just shows up and starts computing without first being vetted by a peer group or other responsible authority. Review criteria and timing vary from site to site.
- Usage must be reported to the site's funding or sponsoring organization. The format and timing of this accountability, though, is as diverse as the sites, agencies, and platforms.

Commonality does not necessarily smooth the implementation of accounting systems, nor is diversity necessarily a barrier. What is critical is that the current practices as the participating sites be examined when the grid is being formed, not just added as an afterthought when a problem among members and/or users arises.

1. Mapping Usage to Users

The current situation at most potential grid sites is that to run jobs on a machine, the user needs to have a local user account on that machine. This will not scale for a grid installation. For this reason, various methods of “accountless accounting” are being investigated. Examples include “Virtual Users” at the Polish National Cluster [2], and Template Accounts [3], which are being implemented at the University of Manchester, in England and are being considered for the European DataGrid project [4].

2. Usage Economies and Methods of Exchange

In the context of a grid, certain fundamental concepts must be defined for resources to be equitably and efficiently allocated and utilized:

- **Supplier:** A provider of grid resources
- **Consumer:** A user of grid resources
- **Value:** A measurement of the usage of grid resources. In the consumer's perspective, this could be seen as *cost* or *price*.
- **Exchange:** The act of utilizing grid resources provided by a grid supplier and received by a grid consumer

A number of economic models are being investigated as potential frameworks for managing grid resource

economies [5]. Since most grids are being developed in response to specific scientific needs and are still falling under the purview of closely held management teams, there are many opportunities to decide what model would best serve a particular grid community. Some examples are “*central control*”, the current standard for computational economy among HPC sites, “*free market economies*” where the allocation of resources is determined solely by supply and demand, and “*bartering*”, where participants make cycles available in exchange for the opportunity to be involved in one of the largest computing projects currently extant

3. Functionality and Methodology

Regardless of the economic model implemented for a given grid, there are certain minimal functions that need to be met for the grid to meet the needs of the member sites. Implementation details are dependent on the platforms and member requirements. There are a number of systems in development, and some in production, that meet these functional requirements.

- 3.1 Supplier Sites must be able to provide resource rates, quotes for resource requests,
- 3.2 Consumer Sites and consumers must be able to obtain quotes for future resource consumption and either request that the resource-consuming job be executed or inform the resource provider that the resource quote was rejected.
- 3.3 The Value of the Resources should be determined by the local resource provider for the resources within their administrative purview. The local authority will also need to decide, for their administrative purview, if a remote user is required to have a local account to utilize local resources. The rates determined by local resource providers for resources, while flexible, must be made available to a potential grid user upon request for a quote. This raises the additional question of how to release a quote that has not been accepted.
- 3.4 Chargeable Items vary greatly from site to site. Ultimately, the only requirement is that the resource usage be presented to the “consumer” in an understandable and decomposable fashion – the user needs to know what the measures are for using a site’s resources so that an informed decision can be made before submitting a job.
- 3.5 Conflict Resolution procedures need to be developed, implemented and published by each site, for disputes over charges incurred. An overall procedure establishing minimum resolution standards must be agreed to and implemented. This will be strongly based on the methods of exchange that have been agreed upon by the participating

sites, but it should not be overlooked when the grid community is developing its charter or service agreements.

- 3.6 Account Balancing is strongly recommended for grid participants. In a centrally controlled system, this will imply maximizing usage on the funded resources; in a barter economy, sites will not participate if they do not believe they are receiving at least as much as they are providing.

4. Conclusion

Accounting and accountability are often overlooked in the excitement of implementing a distributed high performance computational system, but they are critical to the success of the endeavor. In the world of demoware, hypotheses can be tested on the basis of handshakes and email conversations. As grids move into production and begin addressing significant questions, agreements on standards for allocation, access, and accounting will become more important.

Middleware developers are addressing accounting in their packages in a number of ways, and this is a good approach. It is the middleware that will bring the diverse resources together into a cohesive, functioning system. For the sake of ease of use and centralization, it makes sense to collect, maintain and distribute usage data from the same administrative point that is managing other aspects of the interoperating system.

What is most critical is that grid member sites agree upfront on the model and method for resource exchange. This is standard procedure in a closed site, and a key to effectively managing critical resources. As sites become more open, accounting and accountability should not be overlooked.

References

- [1] L. F. McGinnis, “Resource Accounting – Current Practices”, http://www.psc.edu/~lfm/dawg_current_practices.pdf
- [2] M. Kupczyk, N. Meyer, and P. Wolniewicz; “Simplifying Administration and Management Processes in the Polish National Cluster”, <http://www.nas.nasa.gov/~thigpen/accounts-wg/Documents/VUS.pdf>
- [3] T. J. Hacker and B. D. Athey, “Account Allocations on the Grid”, <http://www.nas.nasa.gov/~thigpen/accounts-wg/Documents/accounttemplates.pdf>
- [4] C. Anglano, S. Barale, L. Gaido, A. Guarise, S. Lusso, and A. Werbrouck, “An Accounting System for the DataGrid Project – Preliminary Proposal”, draft in discussion at Global Grid Forum 3, Frascati, Italy, October, 2001
- [5] R. Buyya, D. Abramson, and J. Giddy, “A Case for Economy Grid Architecture for Service Oriented Grid Computing”, <http://www.nas.nasa.gov/~thigpen/accounts-wg/Documents/ecogrid.pdf>