

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. Unfortunately, as grids grow in number of sites and users participating, this method of establishing access to resources will not scale. For example, at the University of Michigan, over 120,000 users are registered and a significant amount of time and energy is spent managing this registry. As the grid grows beyond this scale, continued reliance on the existence of a local user account would engender the need to create a centralized bureaucracy to manage this registry, which is antithetical to stated grid goals.

It should be noted that if a site requires users to have local accounts for remote execution, then the site might not be able to use the full capabilities of the grid. The grid needs to be a fluid environment where sites can exchange cycles and provide access to users of other trusted participating grid sites. The overhead and time delay in requiring local user accounts could easily become the critical bottleneck in this process.

Distributed accounting on the grid assumes the existence of authentication and authorization mechanisms which securely and accurately establish the identity and credentials of user requesting access to grid resources. Once identity and credentials have been established,

distributed accounting methods must be able to map grid resource usage to the requesting user. Since it has already been established that local user accounts are not feasible in a grid environment, various methods of “accountless accounting” are being investigated.

1.1 Virtual Users

The Polish National Cluster is a collection of high performance computing resources distributed throughout Poland. User management is handled by a Virtual User Account System[2], which serves as an interface between a human user and the Polish HPC resources. Access is accomplished by assigning jobs to a Virtual User Account Manager user – a single account that exists on each HPC resource. The heart of the system is a Virtual User Account Server daemon. The daemon keeps track of the mapping between the real and the virtual user.

This takes care of the problem of setting up N accounts on each of the grid member systems. But it does open a potential security risk, since the Virtual User Account Manager needs to have access to *any* resource that *any* user might need at *any* time. Also, any specific licensing issues that might exist on a specific platform or for a specific piece of software will need to be implemented in the VUA Server daemon, so that they can be checked before the job is submitted to the HPC resources.

This is similar to the method used in Condor, middleware for grid management [3]. Jobs run via Condor are run on the target machines as user “nobody”, with accounting information returned to the middleware management platform. This makes user management easy, but is problematic when multiple grid users want to share a grid resource, since the underlying HPC system sees only one user – “nobody”.

1.2 Template Accounts

Another method being investigated is mapping a user to a template account, pulled from a pool of appropriately configured local accounts on each host[4]. Traditionally, a user’s account exists or doesn’t, based on an appropriate record in a password file. The mapping between a user and a template could be controlled by managing the state of the binding – active, inactive, pending, scratch, etc.

This would allow the user to exist indefinitely, without having to have an individual account on each resource. Retention of the binding would depend not only on the actual work being done by the user, but on other policies as well. These policies could include usage reporting and auditing, as well as classical authentication and authorization standards.

This method is being implemented at the University of Manchester, in England [4], and is under consideration as

the method of choice for the DataGrid project [5]. The DataGrid proposal acknowledges that implementing an appropriate accounting system is a complex undertaking. The prototypes and early implementations of these concepts should provide valuable insight into critical issues.

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 [6]. 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.

2.1 Central Control

Central control is the current standard for computational economy among HPC sites. Funding agencies provide financial resources for sites to acquire and provide supercomputing or other specialized technology resources. Researchers then request access to these resources through peer-reviewed access requests. Access is granted to requests deemed worthy, and the researchers’ accounts are established on the appropriate platforms. Accounting systems tend to be “home-grown” – developed in-house at the HPC site to perform the specific tasks required by the funding agency.

These home-grown systems get complicated when resources are funded by more than one agency, which tends to mean different allocation and reporting requirements, even on different fiscal calendars. As grid systems become more complicated, involving more sites and resources, with more sponsors, centralized control of the resources will become that much more complicated as well. All partners in a grid will need to agree from the start on what the accounting and usage requirements will be.

2.2 The Free Market Economy Model

In a free market economy, the allocation of resources is determined solely by supply and demand. Ideally, supply and demand are not subject to regulation other than normal competition, but property rights are allocated and upheld so that trade can occur. In the context of a grid computing system, usage accounting based on a free market economy could provide the following benefits:

Suppliers	Consumers
Resource Control: Each supplier site has control over the set of resources and the quantity of each resource that it chooses to make available on the grid.	Resource Selection: Consumers can choose from a variety of resources that might not otherwise be available to them.
Price: Each supplier site can modify the set of resources and the rates for each resource as needed.	Value: The costs incurred in utilizing resources from various suppliers can be compared prior to submitting a request for resources.
Implementation: Each site can implement as complex or as simple an accounting system as needed. New accounting systems can be easily prototyped. Supplier sites can change the way they charge as they desire with minimum impact or requirements on other sites.	Implementation: Resource requests can be submitted independent of implementation details. The consumer only needs to know a standard method for requesting resources and compensating the resource supplier.
Autonomy: Supplier sites need not agree nor negotiate the relative value of their resources as a prerequisite to making those resources available on the grid.	Independence: Consumers can compare resources available at various sites and select those that best meet their needs.
Exchange	
The free market model provides an automatic way to regulate the utilization of site resources by external resource consumers	

Table 1: Features of a free market for grid communities

2.3 Barter

Bartering is another possible economy for grid communities. This is currently the method being used in “consumer clusters” – e.g. SETI@home. Participants are “bartering” their available cycles for the opportunity to be involved in one of the largest computing projects currently extant. There’s still a strong sense of “doing it for the glory” in some areas of technology. Barter systems can leverage this effectively.

Bartering can also be viable in a grid system, if the participants establish guidelines when the grid is established. Bartering could be as straightforward as trading computing cycles, or be more complicated, involving exchanges of disparate resources – cycles for expertise, for example.

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

The grid resource supplier must be able to provide its resource rates, quotes for resource requests, and resource usage. The following mechanisms should be implemented at a grid resource supplier site:

- **Grid Resource Provider Rates:** The grid resource provider should have a way to set and maintain the rates for the use of its resources. There is no need for an agreement between members on how this information is stored or provided, but they should agree from the start to the standards that will be used to calculate and maintain the rates.
- **Provide Quotes for resource allocation requests:** Grid members need to agree on the format of this message. The response to the resource allocation quote request will provide a cost for the requested resources. The final charge for the resource usage by the job should not exceed the quote if the job resource requirements did not exceed the estimates provided for the quote. The response to the resource quote request will contain the requester's authorization identifier, an expiration date/time that describe when the quote expires, and a server unique identifier. The resource utilization quote provided in response to a resource allocation request will be a total cost and will not be broken down by resource categories. If the request stipulated a range of

charges, all ranges will be provided with a separate unique ID.

- **Track Resource Utilization:** Each grid resource provider can choose to gather information on resources consumed by local and remote users. Grid resource providers must (it's in their own interest) collect information on grid credits collected from resource consumers. Each site must have the access and ability to track the information it will charge for against the particular job request.
- **Job Account Information:** The functionality required to package and transfer the data pertaining to resources utilized by a resource consumer must be defined and agreed upon by all grid participants. For maximum flexibility, sites should be able to provide an accounting record (either pull or push). When the job completes (normally or abnormally), the accounting information is gathered and sent back to the resource-consuming site. This accounting should be broken down by resource category and must include the requesting site's unique ID. Error checking (ack/nack) should be implemented to be sure that this information is delivered. If for some reason the delivery fails, the information must go to the supplier's accounting authority to handle manually.

3.2 Consumer Sites

The grid resource consumer 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.

- **Resource Usage Quote Query:** This should be a request in a common agreed upon format that specifies the resources requested. This resource quote request does not obligate the requester to use the requested services; it is simply a mechanism that the potential resource consumer can use to ascertain potential costs for utilizing the resources that the resource provider can provide. The resource quote request should have a requesting site unique identifier and a description of the resources required. The resource quote request can request a range of charges based upon additional qualifiers such as quality of service if provided by the resource provider site.
- **Accountable Resource Use Request:** If a resource consuming entity decides to use a resource provider site the resource-consuming request should include a unique requester ID and will include the server ID associated with the resource quote provided to the resource consuming entity.

- **Resource Request Quote Cancellation:** Although not required, it is suggested if the requesting site decides to use the successful bidder for a job, a cancellation should be sent to all the resource providers that provided a quote whose resources are not going to be used. This would include canceling unused resource requests from the resource provider site that won the bid. If this cancellation were not sent, the reservation should be removed automatically when the quote expires.

3.3 Valuing Resources

The local resource provider determines the base value of resources within their administrative purview. This resource valuation can be used as a mechanism to attract or deter external users by utilizing the laws of supply and demand. Submitted jobs must therefore contain sufficient resource requirement information to allow local resource allocation software to determine the cost of the local resources that will be consumed by the job.

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. If local resources are provided to remote users without local accounts/accounting, the local resource provider must provide a full accounting of each resource used and the costs charged for each resource for the job. This accounting can be performed immediately (e.g. at the completion of the job), later (i.e. when the accounting software is run), or upon request from the requesting site.

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. Resource quotes should contain a time frame for which the resource quote is valid. The quote process will facilitate an open bidding process for resources that will allow the user to comparison shop. This raises the additional question of how to release a quote that has not been accepted.

3.4 Chargeable Items

Current research has shown great variety in the specific data that is collected for usage accounting. These are some of the major types of metrics used for managing resources on the grid:

Resource Metrics
CPU billing unit
Wall clock or usage billing unit
Memory
If usage is tied to CPUs, amount / CPU
Megabyte of on-line storage
Premium rate(s) for special handling

Higher job queue priority within a job class
Network bandwidth usage (if bandwidth is pre-allocated and reserved)
Special applications
Local consultant, programmer or administrator time utilized beyond normal operation duties
Transportable media

Table 2: Usage metrics currently in use

This list is not definitive nor obligatory. For instance, a site may decide that it will only charge for CPU utilization. It is also not exhaustive. Supplier sites who calculate "usage" using resource metrics not included here are welcome to define their "charges" to meet the unique requirements of their site or particular resource. 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

Each site must implement and publish its conflict resolution procedures 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

Regardless of the economic model agreed upon among the grid participants, each participating site will try to maintain a "zero balance" in the aggregate. 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.

Using standard accounting practices, the following scenario is offered as an example of account balancing in a free market scenario. When a site submits a job that will consume grid resources:

- The resulting resource utilization charges are viewed as a debit to the submitting (consumer) site. The consumer's home site can then decide how to charge the user's authorized project and individual account.
- The resulting resource utilization charge is handled as a credit to the resource provider (supplier) site. This entitles jobs at the supplier's site to use an

equivalent amount of grid resources at the consumer's site.

Ultimately, the credits provide at a supplier site should be balanced by debits incurred as usage at other grid sites.

A resource supplier could potentially increase demand for its resources (and gather more grid credits) by lowering its rates and reduce grid demand by raising its resource rates.

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] M.Y. Koo, "Market-based Grid Computing – Survey Results", draft in discussion at Global Grid Forum 3, Frascati, Italy, October, 2001
- [4] T. J. Hacker and B. D. Athey, "Account Allocations on the Grid", <http://www.nas.nasa.gov/~thigpen/accounts-wg/Documents/accounttemplates.pdf>
- [5] 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
- [6] 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>