

Distributed Accounting on the GRID

Bill Thigpen, Tom Hacker

thigpen@nas.nasa.gov hacker@umich.edu

This paper is very preliminary. We would greatly appreciate any and all input.

This paper proposes a fully distributed view of accounting and a methodology for allocating GRID resources to computation for use on a grid system. It is thought that the grid will contain a large number of unconnected sites, and these sites will need to exchange accounting and bid/quote information. The current situation at most potential GRID sites is that to run jobs on a machine, the user needs to have an account on that machine. Unfortunately, as the GRID grows in the number of sites and users, this method of establishing access to resources will not scale. At the University of Michigan, over 120,000 users are registered and a significant amount of time and energy is spent managing this registry. If the GRID grows beyond this scale, continued reliance on the existence of an account will engender the need to create a centralized bureaucracy to manage this registry, which is antithetical to stated GRID goals. We do not propose that this currently used method go away, but would like to propose another, more flexible way of running a job. This new, proposed method would allow the user to run a job at a remote site without requiring an account on that machine. This assumes that the existence of other authorization mechanisms that can securely and accurately establish the identity and "credentials" of the owner of the job request exists. This paper does not address that mechanism, but assumes its existence for purposes of definition and discussion. Refer to the security GRID working group for work being done on authorization and authentication.

This paper is divided into a discussion of who sets the "value" of the resources and how the assignment of value to resources should occur, a description of the basic functionality required to support client and server operations and a discussion the needed currency balance system.

For an accounting system to be functional in a grid environment, it should be decentralized, scalable and flexible. It should have a minimum impact on local accounting and should not make any limiting assumptions about whether accounting is done by user, group or site. The requirements on the remote site will be to track the resources used by the requesting job. This information will be passed back to the requesting site and that site the information can be accrued in any way that they desire. To do this, it is suggested that we look at a distributed allocation and accounting client/server approach that will work across multiple sites and administrative and policy domains.

Resources

This section discusses the interaction between job resource usage, resources, jobs resource allocation requests and a methodology for controlling the relationship between these components.

The local resource provider authority determines the 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 contain enough 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 decide for their administrative purview if they will require a remote user 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 or later when the accounting software is run (or upon request from the requesting site).

I personally feel that if a site requires users to have a local account for remote execution, it will 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.

To participate in this free-market environment, several interrelated activities must occur. The owners of the grid resources need to come to an agreement on a mechanism to authorize a user for access to a resource. The security-working group is looking into this issue. Once an agreement is reached, the GRID resource allocation system will authorize the resource use.

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. The resource quotes should contain a time frame that the resource quote is valid. The quote process will facilitate an open bidding process for resources that will allow the user to comparison shop.

Required Functionality to Provide Client/Server Services

The basic functionality required to provide client/server operations (and the services that depend upon those operations) is the ability to move the accounting information, resource allocation, resource quotes, and resource account management information from the server or resource provider to the representative of the resource consumer or user. As we mentioned above, if we assume that the requesting site has an allocation for the user the job is being submitted from, then the requesting site must have a mechanism to obtain the maximum cost of a job prior to submission to the GRID to verify adequate resources. The resource requirement information is generated at the client (or resource consumer) site,

then there must be a mechanism for gathering and accepting resource quotes prior to job submission. There must be a mechanism to provide the accounting information to the requesting site.

Server

The GRID resource provider must be able to provide its resource rates, provide quotes and provide resource usage. The following mechanisms should be implemented at a GRID resource provider site.

1. **GRID Resource Provider Rates.** The GRID resource provider site should have a way to set and maintain the rates for the use of its resources. There is no need for an agreement on how this information is stored or provided.
2. **Provide Quotes for resource allocation requests.** An agreement on the format of this message must be a deliverable for the Grid Form. 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.
3. **Track Resource Utilization.** Each GRID resource provider can choose to gather information on resources consumers by local and remote users. GRID resource providers must (it's in their own interest) collect information on GRID credits collected from resource consumers. This has to happen, but does not need to be specified in the GRID community. Each site must have the access and ability to track the information it will charge for against the particular job request.
4. **Job Account Information.** The functionality required to package and transfer the resources utilized by a resource consumer must be defined and agreed upon by all GRID participants. For maximum flexibility, sites must be able to provide an accounting record (either pull or a push). When the job completes (normal or abnormal), 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. This information must be delivered. If for some reason, it can't be, the information must go to the local resource provider authority to handle manually.

Client

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.

1. **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.
2. **Accountable Resource Use Request.** If a resource consuming entity decides to use a resource provider site the resource consuming request will include a unique requester ID and will include the server ID associated with the resource quote provided to the resource consuming entity.
3. **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 would be removed automatically when the quote expires.

Chargeable Items. This list will eventually be an exhaustive list of the major types of resources accounted for on the Grid. As a start, the following resource items are submitted for possible inclusion:

- Premium rate(s) for special handling
- Higher job queue priority within a job class
- Charge per CPU billing unit
- Charge by wall clock or usage billing unit
- Charge for memory
- If ties to CPUs, amount / CPU
- Charge per megabyte of on-line storage
- Charge for network bandwidth usage (if bandwidth pre-allocated and reserved)
- Special Application Charges
Local consultant, programmer or administrator charges for time utilized beyond operation duties to provide assistance
- Transportable media charges (if data moved to tape and shipped to requestor)

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 list will grow, but the fact that an item is on the list does not mean that any resource provider must charge for it. For instance, a site may decide that it will only charge for CPU utilization.

This completes the specification of how an individual job is handled, but does not address GRID resource provider and consumer resource GRID credit account balancing. It is up to a site to try to maintain a zero balance in the aggregate. When a GRID site submits jobs to the grid that will consume GRID resources, the resulting resource utilization charges would be viewed as a negative charge to the GRID site, authorized project and individual users' accounts. When a job from the grid is run on local resources, the resulting resource utilization charge is handled as a positive credit to the resource provider site. A resource providing site 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.

In essence what we are proposing is a free market economy for the GRID. This free market system provides the following benefits:

1. Each site has total control over the set of resources and the quantity of each resource that it chooses to make available on the GRID.
2. No agreements or negotiations between resource-providing sites on the relative value of their resources are required as a prerequisite to making resources available on the GRID.
3. Each site can modify the set of resources and the rates for each resource, as they desire.
4. Provides an automatic way to regulate the utilization of site resources by external resource consumers.
5. Allows each site to implement as complex or as simple an accounting system as they need.
6. Provides an environment that facilitates the prototyping of new accounting systems. Each site can change the way they charge, as they desire with minimum impact or requirements on other sites. The only need would be if they added a new type of charge, the format of the exchange packet would change.

GRID Accounting Policies - Making the Market Work

To create a single metric that represents resource usage, a formula is required that will map from a set of resource utilization or allocation figures to a fixed resource unit. This resource unit will then comprise the basic GRID allocation unit. Project Administrators will be provided with a number of resource allocation units that they can distribute between resources that are made available to them from a GRID producer. The weighting factors in the formula for each GRID site should reflect the supply/demand curve for that resource at that computing site. For example, if a GRID producer provides a high-speed solid-state storage-device resource that is in great demand at their site, they should have a large weighting factor for the use of that resource in the cost formula. Correspondingly, if a GRID producer has a resource that is underused, the weighting factor should be very small for that resource. The utilization of supply and demand curves for all GRID

resource will help to create a resource market force that will equalize and load-balance the use of resource across the entire GRID.

The establishment of policies and procedures to handle GRID sites that will not or can not maintain an aggregate zero balance must be established prior to the production implementation of this accounting system. If sites don't have resources to offer the GRID and just want to consume GRID resources, they should establish a partnership with an existing or new GRID member that can "sell" enough resources to balance their demands. In the case where a contributing site can not maintain an aggregate zero balance, a method must be established to transfer balances from one site to another, possibly involving the transfer of real funds. It is also possible for a group of sites to establish a GRID partnership block to declare a single site for balancing purposes.

When examining the question of accounting balance between GRID computing sites, there are several factors to consider. The first factor is what the GRID "fiscal year" or fiscal period ought to be. If it is too small, unnecessary time and effort is spent attempting to "close the books" to achieve some parity between sites heavy with GRID consumers and sites that heavily supply GRID cycles. If the period is too long, large disparities will arise between GRID consumers and producers, and may jeopardize the willingness of GRID producers to participate in a computational grid. Thus, a suitable "GRID fiscal year" must be identified that will allow the GRID participants to close the books and address any disparities between GRID consumers and producers. How the disparities will be taken care of is not yet defined - it may require the transfer of real funds between net consumers and producers, or it may result in the reduction of resources made available to a group of consumers.