

Gold

User's Guide

Version 2.2.0

Gold User's Guide

version 2.2.0

Gold is an open source accounting system that tracks and manages resource usage on High Performance Computers. It acts much like a bank in which resource credits are deposited into accounts with access controls designating which users, projects, and machines may access the account. As jobs complete or as resources are utilized, accounts are charged and resource usage recorded. Gold supports familiar operations such as deposits, withdrawals, transfers, and refunds. It provides balance and usage feedback to users, managers, and system administrators.

Since accounting needs vary widely from organization to organization, Gold has been designed to be extremely flexible, featuring customizable accounting and supporting a variety of accounting models. Attention has been given to scalability, security, and fault tolerance. Gold facilitates the sharing of resources between organizations or within a Grid by providing distributed accounting while preserving local site autonomy.

[Legal Notices](#)

Table of Contents

- [Features](#)
- [Interfaces](#)
 - [Command Line Clients](#)
 - [Interactive Control Program](#)
 - [Web-based Graphical User Interface](#)
 - [Perl API](#)
 - [SSSRMAP Wire Protocol](#)
- [Installation](#)
 - [Preparation](#)
 - [Select a Database](#)
 - [Install Prerequisites](#)
 - [PostgreSQL database 7.2 or higher \(or other tested database\) \[REQUIRED\]](#)
 - [Perl 5.6.1 or higher \(with suidperl\) \[REQUIRED\]](#)
 - [libxml2 2.4.25 or higher \[REQUIRED\]](#)
 - [Gnu readline 2.0 or higher \[OPTIONAL\]](#)
 - [Apache Httpd Server 2.0 or higher \[OPTIONAL\]](#)
 - [OpenSSL 0.9.5a or higher \[OPTIONAL\]](#)
 - [mod_ssl 2.2.6 or higher \[OPTIONAL\]](#)
 - [Configuration](#)
 - [Compilation](#)
 - [Perl Module Dependencies](#)
 - [Installation](#)
 - [General Setup](#)
 - [Database Setup](#)
 - [Web Server Setup](#)
 - [Bootstrap](#)
 - [Startup](#)
 - [Initialization](#)

- **Getting Started**
 - [Define Users](#)
 - [Define Machines](#)
 - [Define Projects](#)
 - [Add Users to the Projects](#)
 - [Make Deposits](#)
 - [Check The Balance](#)
 - [Integrate Gold with your Resource Management System](#)
 - [Obtain A Job Quote](#)
 - [Make A Job Reservation](#)
 - [Charge for a Job](#)
 - [Refund a Job](#)
 - [List Transactions](#)
 - [Examine Account Statement](#)
 - [Examine Project Usage](#)
- **Getting More Advanced**
 - [Define Projects](#)
 - [Define Accounts](#)
 - [Make Deposits](#)
 - [Check The Balance](#)
 - [Define Charge Rates](#)
 - [Obtain A Guaranteed Job Quote](#)
 - [Make A Quoted Job Reservation](#)
 - [Charge for a Quoted Job](#)
 - [Partially Refund a Job](#)
 - [Examine Account Statement](#)
- **Managing Users**
 - [Creating Users](#)
 - [Querying Users](#)
 - [Modifying Users](#)
 - [Deleting Users](#)
- **Managing Machines**
 - [Creating Machines](#)
 - [Querying Machines](#)
 - [Modifying Machines](#)
 - [Deleting Machines](#)
- **Managing Projects**
 - [Creating Projects](#)
 - [Querying Projects](#)
 - [Modifying Projects](#)
 - [Deleting Projects](#)
 - [Project Usage Summary](#)
- **Managing Accounts**
 - [Creating Accounts](#)
 - [Querying Accounts](#)
 - [Modifying Accounts](#)
 - [Making Deposits](#)
 - [Querying The Balance](#)
 - [Personal Balance](#)
 - [Making Withdrawals](#)
 - [Making Transfers](#)
 - [Obtaining an Account Statement](#)
 - [Deleting Accounts](#)

- **Managing Allocations**
 - Creating Allocations
 - Querying Allocations
 - Modifying Allocations
 - Deleting Allocations
- **Managing Reservations**
 - Creating Reservations
 - Querying Reservations
 - Modifying Reservations
 - Deleting Reservations
- **Managing Quotations**
 - Creating Quotations
 - Querying Quotations
 - Modifying Quotations
 - Deleting Quotations
- **Managing Jobs**
 - Creating Jobs
 - Querying Jobs
 - Modifying Jobs
 - Deleting Jobs
 - Obtaining Job Quotes
 - Making Job Reservations
 - Charging Jobs
 - Issuing Job Refunds
- **Managing Charge Rates**
 - Creating ChargeRates
 - Querying ChargeRates
 - Modifying Charge Rates
 - Deleting Charge Rates
- **Managing Transactions**
 - Querying Transactions
- **Managing Roles**
 - Querying Roles
 - Querying Role Users
 - Querying Role Actions
 - Creating Roles
 - Associating an Action with a Role
 - Adding a Role to a User
 - Removing an Action from a Role
 - Removing a Role from a User
 - Deleting Roles
- **Managing Passwords**
 - Creating Passwords
 - Querying Passwords
 - Modifying Passwords
 - Deleting Passwords
- **Using the Gold Shell (goldsh)**
 - Usage
 - Command Syntax
 - Valid Objects
 - Valid Actions for an Object
 - Valid Predicates for an Object and Action
 - Common Options

- [Common Actions Available for most Objects](#)
 - [Query Action](#)
 - [Create Action](#)
 - [Modify Action](#)
 - [Delete Action](#)
 - [Undelete Action](#)
 - [Multi-Object Queries](#)
- [**Customizing Gold Objects**](#)
 - [Removing an Attribute from an Object](#)
 - [Adding an Attribute to an Object](#)
 - [Modifying an Attribute](#)
 - [Creating a Custom Object](#)
 - [Adding an Action to an Object](#)
 - [Examples Creating Custom Objects](#)
- [**Integration with the Resource Management System**](#)
 - [**Dynamic versus Delayed Accounting**](#)
 - [Delayed Accounting](#)
 - [Dynamic Accounting](#)
 - [**Interaction Points**](#)
 - [Job Quotation @ Job Submission Time \[Optional — Recommended\]](#)
 - [Job Reservation @ Job Start Time \[Optional — Highly Recommended\]](#)
 - [Job Charge @ Job End Time \[Required\]](#)
 - [**Methods of interacting with Gold**](#)
 - [Configuring an application that already has hooks for Gold](#)
 - [Using the appropriate command-line client](#)
 - [Using the Gold control program](#)
 - [Use the Perl API](#)
 - [Communicating via the SSSRMAP Protocol](#)
- [**Configuration Files**](#)
 - [Server Configuration](#)
 - [Client Configuration](#)

Legal Notices

Copyright

© 2010 Adaptive Computing Enterprises, Inc. All rights reserved. Distribution of this document for commercial purposes in either hard or soft copy form is strictly prohibited without prior written consent from Adaptive Computing Enterprises, Inc.

Trademarks

Adaptive Computing, Cluster Resources, Moab, Moab Workload Manager, Moab Cluster Manager, Moab Cluster Suite, Moab Grid Scheduler, Moab Grid Suite, Moab Access Portal, and other Adaptive Computing products are either registered trademarks or trademarks of Adaptive Computing Enterprises, Inc. The Adaptive Computing logo and the Cluster Resources logo are trademarks of Adaptive Computing Enterprises, Inc. All other company and product names may be trademarks of their respective companies.

Acknowledgments

Gold includes software developed by [Pacific Northwest National Laboratory](#) and [Battelle Memorial Institute](#).

Features

- **Dynamic Charging** — Rather than post-processing resource usage records on a periodic basis to rectify project balances, accounts are updated immediately at job completion.
- **Reservations** — A hold is placed against the account for the estimated number of resource credits before the job runs, followed by an appropriate charge at the moment the job completes, thereby preventing projects from using more resources than were allocated to them.
- **Flexible Accounts** — A uniquely flexible account design allows resource credits to be allocated to specific projects, users, and machines.
- **Expiring Allocations** — Resource credits may be restricted for use within a designated time period allowing sites to implement a use-it-or-lose-it policy to prevent year-end resource exhaustion and establishing a project cycle.
- **Flexible Charging** — The system can track and charge for composite resource usage (memory, disk, CPU, etc) and custom charge multipliers can be applied (Quality of Service, Node Type, Time of Day, etc).
- **Guaranteed Quotes** — Users and resource brokers can determine ahead of time the cost of using resources.
- **Credit and Debit Accounts** — Accounts feature an optional credit limit allowing support for both debit and credit models. This feature can also be used to enable overdraft protection for specific accounts.
- **Nested Accounts** — A hierarchical relationship may be created between accounts. This allows for the delegation of management responsibilities, the establishment of automatic rules for the distribution of downstream resource credits, and the option of making higher level credits available to lower level accounts.
- **Powerful Querying** — Gold supports a powerful querying and update mechanism that facilitates flexible reporting and streamlines administrative tasks.
- **Transparency** — Gold allows the establishment of default projects, machines, and users. Additionally Gold can allow user, machines, and projects to be automatically created the first time they are seen by the resource management system. These features allow job submitters to use the system without even knowing it.
- **Security** — Gold supports multiple security mechanisms for strong authentication and encryption.
- **Role Based Authorization** — Gold provides fine-grained (instance-level) Role Based Access Control for all operations.
- **Dynamic Customization** — Sites can create or modify record types on the fly enabling them to meet their custom accounting needs. Dynamic object creation allows sites to customize the types of accounting data they collect without modifying the code. This capability turns this system into a generalized information service. This capability is extremely powerful and can be used to manage all varieties of custom configuration data, to provide meta-scheduling resource mapping, or to function as a persistence interface for other components.
- **Multi-Site Exchange** — A traceback mechanism will allow all parties of a transaction (resource requestor and provider) to have a first-hand record of the resource utilization and to have a say as to whether or not the job should be permitted to run, based on their independent policies and priorities. A job will only run if all parties are agreeable to the idea that the target resources can be used in the manner and amount requested. Support for traceback debits will facilitate the establishment of trust and exchange relationships between administrative domains.
- **Web Interface** — Gold will implement a powerful dynamic web-based GUI for easy remote access for users, managers, and administrators.
- **Journaling** — Gold implements a journaling mechanism that preserves the indefinite historical state of all objects and records. This powerful mechanism allows historical bank statements to be generated, provides an undo/redo capability, and allows commands to be run as if it were any arbitrary time in

the past.

- **Open Source** — Being open source allows for site self-sufficiency, customizability, and promotes community development and interoperability.

Interfaces

Gold provides a variety of means of interaction, including command-line interfaces, graphical user interfaces, application programming interfaces, and communication protocols.

Command Line Clients

The command-line clients provided feature rich argument sets and built-in documentation. These commands allow scripting and are the preferred way to interact with Gold for basic usage and administration. Use the `--help` option for usage information or the `--man` option for a manual page on any command.

Example 1. Listing Users

```
glsuser
```

Interactive Control Program

The `goldsh` command uses a control language to issue object-oriented requests to the server and display the results. The commands may be included directly as command-line arguments or read from stdin. Use the `"ShowUsage:=True"` option after a valid Object Action combination for usage information on the command.

Example 2. Listing Users

```
goldsh User Query
```



The `goldsh` control program allows you to make powerful and sweeping modifications to Gold objects. Do not use this command unless you understand the syntax and the potential for unintended results.

Web-based Graphical User Interface

A powerful and easy-to-use web-based GUI is being developed for use by users, managers, and administrators. It sports two interface types:

- **Management Interface** — The management interface supports an interface that makes administration and interaction very safe and easy. It approaches things from a functional standpoint, aggregating results and protecting against accidental modifications.
- **Object Interface** — The object interface exposes you to the full power of the actions the server can perform on the objects. This interface allows actions to be performed on many objects in a single command and can impose arbitrary field conditions, field updates, and field selections to the query.

Example 3. Listing Users

Click on "Manage Users" -> "List Users"

Perl API

You can access the full Gold functionality via the Perl API. Use `perldoc` to obtain usage information for the Perl Gold modules.

Example 4. Listing Users

```
use Gold;  
  
my $request = new Gold::Request(object => "User", action => "Query");
```

```
my $response = $request->getResponse();
foreach my $datum ($response->getData())
{
    print $datum->toString(), "\n";
}
```

SSSRMAP Wire Protocol

It is also possible to interact with Gold by directly using the SSSRMAP Wire Protocol and Message Format over the network. Documentation for these protocols can be found at [SSS Resource Management and Accounting Documentation](#).

Example 5. Listing Users

```
POST /SSSRMAP HTTP/1.1
Content-Type: text/xml; charset="utf-8"
Transfer-Encoding: chunked

190
<?xml version="1.0" encoding="UTF-8"?>
<Envelope>
  <Body actor="scottmo" chunking="True">
    <Request action="Query" object="User"></Request>
  </Body>
  <Signature>
    <DigestValue>azu4obZswzBt890gATukBeLyt6Y=</DigestValue>
    <SignatureValue>YXE/C08XX3RX4PMU1bWju+5/E5M=</SignatureValue>
    <SecurityToken type="Symmetric" name="scottmo"></SecurityToken>
  </Signature>
</Envelope>
0
```

Installation

Gold uses the standard configure, make, and make install steps. However, there are a number of preparation, prerequisite, setup, and customization steps that need to be performed. This document provides general installation guidance and provides a number of sample steps referenced to a particular installation on a Linux platform using the bash shell. These steps indicate the userid in brackets performing the step. The exact commands to be performed and the user that issues them will vary based on the platform, shell, installation preferences, etc.

Preparation

To build and install Gold, you first need to unpack the archive and change directory into the top directory of the distribution. For security reasons, it is recommended that you install and run Gold under its own non-root userid.

```
[root]# useradd scottmo
[root]# passwd scottmo
[scottmo]$ mkdir ~/src
[scottmo]$ cd ~/src
[scottmo]$ gzip -cd gold-2.1.10.0.tar.gz | tar xvf -
[scottmo]$ cd gold-2.1.10.0
```

Select a Database

Gold makes use of a database for transactions and data persistence. Three databases have been tested for use with Gold thus far: PostgreSQL, MySQL, and SQLite. Postgres and MySQL are external databases which run in a distinct (possibly remote) process and communicate over sockets. These databases must be separately installed, configured, and started. SQLite is an embedded database bundled with the Gold source code with SQL queries being performed within the goldd process itself through library calls. The following information may help you make a choice of databases to use.

- **PostgreSQL** — PostgreSQL is an open source database. Gold requires Postgres 7.2 or higher (7.1 can probably be used but generates warnings from the DBD::Pg module). The PostgreSQL database has been thoroughly tested in production with Gold and all Gold functionality is available since it was developed using the PostgreSQL database. Postgres supports multiple connections so Gold is configured to be a forking server when using PostgreSQL.

PostgreSQL is recommended since it is an excellent database, has been more thoroughly tested than the others, and supports all Gold features.

- **MySQL** — MySQL is an open source database. Gold requires MySQL 4.0.6 or higher. (Prior versions did not support UNION which is used by Gold in time travel. It is possible to use 4.0 with a minor code tweak to the OFFSET line in Database.pm).

MySQL 4.1 is required in order to have support for the (undocumented) Transaction Undo and Redo functionality since subqueries were not supported until this version.

- **SQLite** — SQLite is an open source embedded database bundled with Gold. It does not require any configuration and reads and writes from a file. Initial testing has shown Gold to perform at least as fast as PostgreSQL for small databases.

Due to the lack of "ALTER TABLE" functionality, Gold objects cannot be customized after installation. It appears that this functionality is likely to be forthcoming in a future release of SQLite.

Since SQLite supports only a single connection, Gold is not configured to be a forking server when

using SQLite. This should probably not be an issue for small to medium sized clusters.

Due to a lack of support for multi-column IN clauses, the (undocumented) Transaction Undo and Redo functions are not available.

Install Prerequisites

You will first need to build, test and install the following prerequisites:

PostgreSQL database 7.2 or higher (or other tested database) [REQUIRED]

Gold makes use of a database for transactions and data persistence. Three databases have been tested for use with Gold thus far: PostgreSQL, MySQL and SQLite (see [Select a Database](#)). If you intend to use the PostgreSQL or the MySQL database, you will need to install it. PostgreSQL is recommended since it is an excellent database, has been more thoroughly tested than the others, and supports the most features. PostgreSQL is available at: <<http://www.postgresql.org/>>

```
[root]# cd /usr/local/src

[root]# wget
http://ftp7.us.postgresql.org/pub/postgresql//source/v8.3.3/postgresql-
8.3.3.tar.gz

[root]# gzip -cd postgresql-8.3.3.tar.gz | tar -xvf -

[root]# cd postgresql-8.3.3

[root]# ./configure

[root]# make

[root]# make install

[root]# adduser postgres

[root]# mkdir /usr/local/pgsql/data

[root]# chown postgres /usr/local/pgsql/data

[root]# touch /var/log/pgsql

[root]# chown postgres /var/log/pgsql
```

Or if you are using rpms, you will need the postgresql, postgresql-libs, postgresql-server, and postgresql-devel rpms appropriate for your architecture and operating system:

```
[root]# wget
ftp://rpmfind.speakeasy.net/linux/redhat/updates/9/en/os/i386/postgresq
7.3.2-3.i386.rpm

[root]# wget
ftp://rpmfind.speakeasy.net/linux/redhat/updates/9/en/os/i386/postgresq
libs-7.3.2-3.i386.rpm

[root]# wget
ftp://rpmfind.speakeasy.net/linux/redhat/updates/9/en/os/i386/postgresq
server-7.3.2-3.i386.rpm

[root]# wget
ftp://rpmfind.speakeasy.net/linux/redhat/updates/9/en/os/i386/postgresq
devel-7.3.2-3.i386.rpm

[root]# rpm -Uvh postgresql-7.3.2-3.i386.rpm postgresql-libs-7.3.2-
3.i386.rpm postgresql-server-7.3.2-3.i386.rpm postgresql-devel-7.3.2-
```

3.i386.rpm

Perl 5.6.1 or higher (with suidperl) [REQUIRED]

NOTE: Recent operating systems distributing Perl 5.12 or higher no longer provide suidperl. Since this is the only security promotion method available in Gold, you will need to upgrade to Moab Accounting Manager 7.1 or higher in order to use the new gauth security promotion method.

The Gold server and clients are written in Perl. Perl 5.6.1 or higher is required. The perl installation must include suidperl for proper client authentication. Use 'perl -v' to see what level of Perl is installed and 'suidperl -v' to see if suidperl is installed. Perl is available at: <<http://www.perl.com/>>

```
[root]# wget http://www.cpan.org/src/perl-5.10.0.tar.gz
[root]# gzip -cd perl-5.10.0.tar.gz | tar xvf -
[root]# cd perl-5.10.0
[root]# sh Configure -Dd_dosuid -de
[root]# make
[root]# make test
[root]# make install
[root]# (cd /usr/include && /usr/local/bin/h2ph *.h sys/*.h)
```

Or if you are using rpms, you will need the perl and the perl-suidperl rpms appropriate for your architecture and operating system:

```
[root]# wget
ftp://rpmfind.speakeasy.net/linux/redhat/updates/9/en/os/i386/perl-
5.8.3-18.1.i386.rpm

[root]# wget
ftp://rpmfind.speakeasy.net/linux/redhat/updates/9/en/os/i386/perl-
suidperl-5.8.3-18.1.i386.rpm

[root]# rpm -Uvh perl-5.8.3-18.1.i386.rpm perl-suidperl-5.8.3-
18.1.i386.rpm
```

libxml2 2.4.25 or higher [REQUIRED]

LibXML2 is needed by the XML::LibXML perl module to communicate via the SSSRMAP message format. LibXML2 is available at: <<http://www.xmlsoft.org/>>

```
[root]# cd /usr/local/src

[root]# wget -passive-ftp ftp://xmlsoft.org/libxml2/libxml2-
2.6.32.tar.gz

[root]# gzip -cd libxml2-2.6.32.tar.gz | tar xvf -
[root]# cd libxml2-2.6.32
[root]# ./configure
[root]# make
[root]# make install
```

Gnu readline 2.0 or higher [OPTIONAL]

The interactive control program (goldsh) can support command-line-editing capabilities if readline support is enabled. Most recent linux distributions come with the appropriate readline support. Gnu readline is available at: <<http://www.gnu.org/>>

```
[root]# cd /usr/local/src
[root]# wget http://ftp.gnu.org/gnu/readline/readline-5.0.tar.gz
[root]# gzip -cd readline-5.0.tar.gz | tar xvf -
[root]# cd readline-5.0
[root]# ./configure
[root]# make
[root]# make install
```

Apache Httpd Server 2.0 or higher [OPTIONAL]

Gold provides a web based GUI so that managers, users, and administrators can interact with the accounting and allocation system. The web interface utilizes Perl CGI and SSL and needs to have an httpd server (preferably apache) installed. Apache httpd is available at: <<http://httpd.apache.org/>>

```
[root]# cd /usr/local/src
wget http://rpm.emsl.pnl.gov/3.0AW/en/os/i386-U4/RedHat/RPMS/httpd-
2.0.46-44.ent.i386.rpm
[root]# rpm -Uvh httpd-2.0.46-44.ent.i386.rpm
```

OpenSSL 0.9.5a or higher [OPTIONAL]

If you are installing the GUI you will need SSL (preferably OpenSSL). OpenSSL is a command line toolkit for using secure socket layer encryption on a server. OpenSSL is available at: <<http://www.openssl.org/>>

```
[root]# cd /usr/local/src
wget http://rpm.emsl.pnl.gov/3.0AW/en/os/i386-U4/RedHat/RPMS/openssl-
0.9.7a-33.12.i386.rpm
[root]# rpm -Uvh openssl-0.9.7a-33.12.i386.rpm
```

mod_ssl 2.26 or higher [OPTIONAL]

If you are installing the GUI you will need an apache interface to OpenSSL (preferably mod_ssl). There are other alternatives to mod_ssl (one of which is apache-ssl from which the mod_ssl code was forked), however mod_ssl has become the defacto standard and is the most widely adopted. mod_ssl is available at: <<http://www.modssl.org/>>

```
[root]# cd /usr/local/src
wget http://rpm.emsl.pnl.gov/3.0AW/en/os/i386-U4/RedHat/RPMS/mod_ssl-
2.0.46-44.ent.i386.rpm
[root]# rpm -Uvh mod_ssl-2.0.46-44.ent.i386.rpm
```

Configuration

To configure Gold, run the "configure" script provided with the distribution.

To see the list of options:

- -h, --help display the list of options
- Use prefix to tell it where Gold should be installed (defaults to /opt/gold)
--prefix=PREFIX install architecture-independent files in PREFIX
- Use with-db to specify the database you intend to use with Gold. Currently only PostgreSQL (Pg), MySQL (mysql), and SQLite (SQLite) have been tested for use with Gold. Postgres and MySQL are external databases which runs in a distinct (possibly remote) process and communicates over sockets while SQLite is an embedded database bundled with Gold with SQL queries being performed within the gold process itself through library calls. Initial testing has shown SQLite to be at least as fast as PostgreSQL for small installations. The default is to use PostgreSQL.
--with-db=DATABASE database to be used { Pg, mysql, SQLite } [Pg]
- Use without-readline if you do not want to use the gnu readline library
--without-readline Don't use readline in interactive control program
- Use with-user to specify the userid that Gold will run under (defaults to the user running the configure command).
--with-user=USER user id under which the Gold server will run
- Use with-log-dir to specify the directory to which logs will be written (defaults to PREFIX/log).
--with-log-dir=PATH directory for log files [PREFIX/log]
- Use with-perl-libs to indicate whether you want to install the required perl modules in a local Gold directory (PREFIX/lib) or in the default system site-perl directory (triggered by running make deps).
--with-perl-libs=local|site install policy for prerequisite perl libs [local]
- Use with-gold-libs to indicate whether you want to install the Gold modules in a local Gold directory (PREFIX/lib) or in the default system site-perl directory (defaults to local).
--with-gold-libs=local|site install policy for Gold perl libs [local]
- If you will intend to use the Gold web GUI, use with-cgi-bin to specify the directory where you want the Gold CGI files to reside (defaults to /var/www/cgi-bin/gold).
--with-cgi-bin=DIR directory to install cgi-bin files if using web GUI [/var/www/cgi-bin/gold]

The PERL environment variable helps the install process find the desired (5.6) perl interpreter if it is not in your path or not found first in a path search.

PERL full pathname of the Perl interpreter

Some other influential environment variables are:

CC C compiler command

CFLAGS C compiler flags

LDFLAGS linker flags, e.g. -L<lib dir> if you have libraries in a nonstandard directory <lib dir>

CPPFLAGS C/C++ preprocessor flags, e.g. -I<include dir> if you have headers in a nonstandard directory <include dir>

Example 1:

```
[scottmo]$ cd gold-2.2.0
[scottmo]$ ./configure
```


Compilation

To compile the program, type make:

```
[scottmo]$ make
```

If you would like to install the web GUI, type make gui:

```
[scottmo]$ make gui
```

Perl Module Dependencies

Gold requires the use of a number of Perl modules. These modules are included in tarball form in the Gold distribution and they can be installed by typing 'make deps':

```
[root]# make deps
```

This will install the following Perl modules as necessary. By default, these will be installed under gold's lib/perl5 directory. To install these in the system site-perl directory, use the configure parameter with-perl-libs as described in the configuration section.

- CGI.pm
- CGI::Session
- Compress::Zlib
- Crypt::CBC
- Crypt::DES
- Crypt::DES_EDE3
- Data::Properties
- Date::Manip
- DBI
- DBD::Pg or DBD::SQLite
- Digest
- Digest::HMAC
- Digest::MD5
- Digest::SHA1
- Error
- Log::Dispatch
- Log::Dispatch::FileRotate
- Log::Log4perl
- MIME::Base64
- Module::Build
- Params::Validate
- SOAP
- Term::ReadLine::Gnu
- Time::HiRes
- XML::SAX
- XML::LibXML::Common
- XML::LibXML
- XML::NamespaceSupport

If you would prefer to do so, you could install these modules via other sources, such as from rpm, or from CPAN using 'perl -MCPAN -e shell'.

Installation

Use `make install` to install Gold. You may need to do this as root if any of the installation or log directories do not already have write permission as the Gold admin user.

```
[root]# make install
```

If you would like to install the web GUI, type make install-gui (as root).

```
[root]# make install-gui
```

The standard installation process will copy the binaries and perl scripts to /usr/local/bin, install the server in /usr/local/sbin, put the libs in /usr/local/lib, the config files in /usr/local/etc, and the man pages in /usr/local/man. You can customize the directories either through the configuration process or by making the necessary changes in the Makefile.

To delete the files created by the Gold installation, you can use 'make uninstall'.

You will also need to generate a secret key which enables secure communication between clients and server. This key is a pass-phrase consisting of up to 80 characters and can include spaces and the regular visible ASCII characters. Note that if you are using Gold with the Maui Scheduler, they will need both need to use a shared secret key.

```
[root]# make auth_key
```

Enter your secret key (up to 80 characters and can include spaces): sss

General Setup

Edit the Gold configuration files.

```
[scottmo]$ vi /opt/gold/etc/goldd.conf
```

```
[scottmo]$ vi /opt/gold/etc/gold.conf
```

Database Setup

If you have chosen to use PostgreSQL, you will need to configure the database to support Gold connections and schema. No setup is needed if you are using SQLite.

Initialize the database (if you installed from tarball).

```
[postgres]$ /usr/local/pgsql/bin/initdb -D /usr/local/pgsql/data
```

Add the IP ADDRESS of the host where the Gold server will run (even if it is the same host as the database server).

```
[postgres]$ echo "host all all 192.168.1.1 255.255.255.255 trust"
>>/usr/local/pgsql/data/pg_hba.conf
```

Startup postgres with the -i option to allow internet domain sockets

```
[postgres]$ /usr/local/pgsql/bin/postmaster -i -D
/usr/local/pgsql/data >/var/log/pgsql 2>&1 &
```

Add the "gold" user as a database administrator

```
[postgres]$ /usr/local/pgsql/bin/createuser gold
Shall the new user be allowed to create databases? y
Shall the new user be allowed to create more new users? n
```

Create the Gold database

```
[scottmo]$ /usr/local/pgsql/bin/createdb gold
```

Edit the Gold configuration files.

```
[scottmo]$ vi /opt/gold/etc/goldd.conf
[scottmo]$ vi /opt/gold/etc/gold.conf
```

Web Server Setup

If you want to use the Gold web GUI, you will need to configure your Httpd server to use SSL. For RedHat Linux systems, a good guide on this is "Building a Secure RedHat Apache Server HOWTO" at <http://www.faqs.org/docs/Linux-HOWTO/SSL-RedHat-HOWTO.html>.

The following shows an example configuration that involves making some modifications to the httpd configuration to support the use of cgi-bin and SSL connections as well as the creation of a private key and a self-signed certificate.

Edit the httpd.conf file under /etc/httpd/conf:

```
[root]# cd /etc/httpd/conf
[root]# cp httpd.conf httpd.conf.orig
[root]# vi httpd.conf
```

Edit your cgi-bin Directory to agree with the cgi-bin directory you configured Gold to use and ensure it has the following properties:

```
<Directory "/var/www/cgi-bin">
    Options ExecCGI
    AddHandler cgi-script .cgi .pl
</Directory>
```

Add a virtual host definition and edit as appropriate for your environment:

```
<VirtualHost 192.168.72.24:443>
    DocumentRoot /var/www/cgi-bin/gold
    ServerName gold-server.whatever.org
    ServerAdmin Your.Email@whatever.org
    ErrorLog logs/gold-error_log
    TransferLog logs/gold-access_log
    SSLEngine on
    SSLCertificateFile /etc/httpd/conf/ssl.crt/gold-server.crt
    SSLCertificateKeyFile /etc/httpd/conf/ssl.key/gold-server.key
    SetEnvIf User-Agent ".*MSIE.*" nokeepalive ssl-unclean-shutdown
</VirtualHost>
```



If you are installing your cgi-bin files directly under /var/www/cgi-bin, use /var/www/cgi-bin as your DocumentRoot. If you are installing your cgi-bin files under a subdirectory such as /var/www/cgi-bin/gold, you may want to use /var/www/cgi-bin/gold as your DocumentRoot. You could specify /var/www/cgi-bin here, but then you'll need to use an extra gold subdirectory in your URL when accessing the Gold GUI from your browser.

Create an Alias for cgi-bin pointing to your cgi-bin directory. You may need to callout your specific cgi-bin subdirectory if your web server configuration interferes with your cgi-bin alias. You may also need to comment out any conflicting ScriptAlias definition:

```
#ScriptAlias /cgi-bin/ "/var/www/cgi-bin/"
Alias /cgi-bin/gold "/var/www/cgi-bin/gold"
```

Create a Private Key for Gold

```
[root]# mkdir ssl.key
[root]# openssl genrsa -out ssl.key/gold-server.key 1024
```

Create a Self-Signed Certificate

```
[root]# openssl req -new -key ssl.key/gold-server.key -x509 -out  
ssl.crt/gold-server.crt
```

Startup or restart httpd.

```
[root]# /usr/sbin/apachectl restart
```



In order to use the web GUI, users will have to generate passwords for themselves using the gchpasswd client command.

```
[scottmo]# gchpasswd
```

To access the web GUI, open a browser with url: `https://$server/gold.cgi`

```
[scottmo]# mozilla https://gold-server/gold.cgi
```

Bootstrap

You will need to populate the Gold database with an sql dump that defines the objects, actions, and attributes necessary to function as an Accounting and Allocation Manager.

If you are using PostgreSQL:

```
[scottmo]$ /usr/local/pgsql/bin/psql gold < bank.sql
```

If you are using SQLite:

```
[scottmo]$ /opt/gold/sbin/sqlite /opt/gold/data/gold.db < bank.sql
```


Startup

Start the Gold server daemon. It is located in the PREFIX/sbin directory.

```
[scottmo]$ /opt/gold/sbin/goldd
```

Alternatively, if you are on linux system that supports init.d scripts, you can add an add Gold as a system startup service by copying etc/gold.d to /etc/init.d/gold, giving it execute permission, and then start Gold by issuing:

```
[root]# service gold start
```

Initialization

You are now ready to define users, projects, machines, accounts etc. as necessary for your site. The next chapter (Getting Started) provides a useful primer for this phase of the Gold setup.

Getting Started

In order to prepare Gold for use as an allocation and accounting manager, you will need to perform some initial steps to define users, machines, and projects, and then make deposits, etc. This chapter proceeds by offering a number of examples in performing these steps. These steps may be used as a guide, substituting values and options appropriate for your system.

It is assumed that you have already installed and bootstrapped Gold as an allocation and accounting manager and started the Gold server before performing the steps suggested in this section.



You will need to be a Gold System Administrator to perform the tasks in this chapter.

Define Users

First, you will need to define the users that will use, manage or administer the resources (see [Creating Users](#)).

Example 1. Add the users amy, bob, and dave.

```
$ gmkuser -n "Wilkes, Amy" -E "amy@western.edu" amy
Successfully created 1 User
$ gmkuser -n "Smith, Robert F." -E "bob@western.edu" bob
Successfully created 1 User
$ gmkuser -n "Miller, David" -E "dave@western.edu" dave
Successfully created 1 User
$ glsuser
```

Name	Active	CommonName	PhoneNumber	EmailAddress
DefaultProject		Description		
gold	True			
Gold		Admin		
amy	True	Wilkes, Amy		
amy@western.edu				
bob	True	Smith, Robert F.		
bob@western.edu				
dave	True	Miller, David		
dave@western.edu				

Define Machines

You will also need to add the names of the machines that provide resources (see [Creating Machines](#)).

Example 2. Define machines called colony and blue.

```
$ gmkmachine -d "Linux Cluster" colony
Successfully created 1 Machine
$ gmkmachine -d "IBM SP2" blue
Successfully created 1 Machine
$ glsmachine
```

Name	Active	Architecture	OperatingSystem
Description			
colony	True		Linux
Cluster			
blue	True		IBSP2

Define Projects

Next you should create the projects that will use the resources (see [Creating Projects](#)).



In these examples, assume that the `account.autogen` configuration parameter is set to automatically create a default account for each project (see [Server Configuration](#)).

Example 3. Define the projects **biology** and **chemistry**.

```
$ gmkproject -d "Biology Department" biology
Successfully created 1 Project
Auto-generated Account 1

$ gmkproject -d "Chemistry Department" chemistry
Successfully created 1 Project
Auto-generated Account 2

$ glsproject
```

Name	Active	Users	Machines	Description
biology	True			BiologyDepartment
chemistry	True			ChemistryDepartment

Add Users to the Projects

Although this could have been done at the project creation step, you can now assign users to be members of your projects (see [Modifying Projects](#)).

Example 4. Adding users to your projects.

```
$ gchproject -addUsers amy,bob biology
Successfully created 1 ProjectUser
Successfully created 1 ProjectUser

$ gchproject -addUsers amy,bob,dave chemistry
Successfully created 1 ProjectUser
Successfully created 1 ProjectUser
Successfully created 1 ProjectUser

$ glsproject
```

Name	Active	Users	Machines	Description
biology Department	True	amy,bob		Biology
chemistry Department	True	amy,dave,bob		Chemistry

Make Deposits

Now you can make some deposits (see [Making Deposits](#)).

Example 5. Add 360000000 credits to each project and cause them both to be valid just for the fiscal year 2005.

```
$ gdeposit -s 2005-01-01 -e 2006-01-01 -z 360000000 -p biology
Successfully deposited 360000000 credits into account 1
$ gdeposit -s 2005-01-01 -e 2006-01-01 -z 360000000 -p chemistry
Successfully deposited 360000000 credits into account 2
```

Use **glsalloc** to examine allocations:

```
$ glsalloc
```

Id	Account	StartTime	EndTime	Amount	
	CreditLimit	Deposited	Description		
1	1	2005-01-01	2006-01-01	360000000	0
2	2	2005-01-01	2006-01-01	360000000	0

Check The Balance

You can verify the resulting balance (see [Querying The Balance](#)).

Example 6. View amy's balance

```
$ gbalance -u amy
```

Id	Name	Amount	Reserved	Balance	CreditLimit
1	biology	360000000	0	360000000	0
2	chemistry	360000000	0	360000000	0

Example 7. You may just want the total balance for a certain project and machine

```
$ gbalance -u amy -p chemistry -m colony -total
```

Balance
360000000
The account balance is 360000000 credits

Integrate Gold with your Resource Management System

Now you are ready to run some jobs. Before doing so you will need to integrate Gold with your Resource Management System (see [Integrating with the Resource Management System](#)).

Although the quotation, reservation, and charge steps will most likely be invoked automatically by your resource management system, it is useful to understand their effects by invoking them manually.

Now we'll simulate the lifecycle of a job.

Example 8. Assume the job has the following characteristics:

```
Job Id:           PBS.1234.0
Job Name:         heavywater
User Name:        amy
Project Name:     chemistry
Machine Name:     colony
Requested Processors: 16
Estimated WallClock: 3600 seconds
Actual WallClock: 1234 seconds
```

Obtain A Job Quote

When a job is submitted, it is useful to check that the user's account has enough funds to run the job. This will be verified when the job starts, but by that point the job may have waited some time in the queue only to find out it never could have run in the first place. The job quotation step (see [Obtaining Job Quotes](#)) can fill this function. Additionally, the quote can be used to determine the cheapest place to run, and to guarantee the current rates will be used when the job is charged.

Example 9. See how much it will cost to run the job.

```
$ gquote -p chemistry -u amy -m colony -P 16 -t 3600  
Successfully quoted 57600 credits
```

Make A Job Reservation

When a job starts, the resource management system creates a reservation (or pending charge) against the appropriate allocations based on the estimated wallclock limit specified for the job (see [Making a Job Reservation](#)).

Example 10. Make a reservation for the job.

```
$ greserve -J PBS.1234.0 -p chemistry -u amy -m colony -P 16 -t 3600
Successfully reserved 57600 credits for job PBS.1234.0

$ glsres
```

Id		Account	Amount	Name	Job	User	Project
Machine		EndTime				Type	Description

1	2	57600	PBS.1234.0	1	amy	chemistry	
colony	2005-08-03	15:29:30-07			Normal		

This reservation will decrease the balance by the amount reserved.

```
$ gbalance -p chemistry -total -quiet
359942400
```

Although the allocation has not changed.

```
$ glsalloc -p chemistry
```

Id		Account	StartTime	EndTime	Amount	
CreditLimit		Deposited	Description			

2	2	2005-01-01	2006-01-01	360000000	0	
360000000						

This is best illustrated by the detailed balance listing:

```
$ gbalance -p chemistry
```

Id		Name	Amount	Reserved	Balance	CreditLimit
Available						

2	chemistry	360000000	57600	359942400	0	
359942400						

Charge for a Job

After a job completes, any associated reservations are removed and a charge is issued against the appropriate allocations based on the actual wallclock time used by the job (see [Charging Jobs](#)).

Example 11. Issue the charge for the job.

```
$ gcharge -J PBS.1234.0 -u amy -p chemistry -m colony -P 16 -t 1234 -X WallDuration=1234

Successfully charged job PBS.1234.0 for 19744 credits
1 reservations were removed
```

Your allocation will now have gone down by the amount of the charge.

```
$ glsalloc -p chemistry

Id  Account      StartTime      EndTime      Amount
CreditLimit    Deposited      Description
-----
2    2            2005-01-01    2006-01-01    359980256    0
360000000
```

However, your balance actually goes up (because the reservation that was removed was larger than the actual charge).

```
$ gbalance -p chemistry -total

Balance
-----
359980256
The account balance is 359980256 credits
```

A job record was created for the job as a side-effect of the charge (see [Querying Jobs](#)).

```
$ glsjob

Id  JobId      User  Project      Machine  Charge  Class
Type      Stage      QualityOfService      Nodes  Processors
Executable Application      StartTime      EndTime      WallDuration
QuoteId    Description
-----
-----
1    PBS.1234.0    amy    chemistry    colony    19744
Normal    Charge
1234      1
```

Refund a Job

Since this was an imaginary job, refund the user's account (see [Issuing Job Refunds](#)).

Example 12. Issue a refund for the job.

```
$ grefund -J PBS.1234.0  
Successfully refunded 19744 credits for job PBS.1234.0
```

The balance is back as it was before the job ran.

```
$ gbalance -p chemistry -total  
Balance  
-----  
360000000  
The account balance is 360000000 credits
```

The allocation, of course, is likewise restored.

```
$ glsalloc -p chemistry  
Id  Account      StartTime      EndTime      Amount  
CreditLimit    Deposited      Description  
-----  
2   2             2005-01-01    2006-01-01    360000000    0  
360000000
```

Notice that the job charge is now zero because the job has been fully refunded.

```
$ glsjob  
Id  JobId      User  Project      Machine  Charge  Class  
Type  Stage      QualityOfService  Nodes  Processors  
Executable  Application      StartTime      EndTime      WallDuration  
QuoteId      Description  
-----  
-----  
-----  
1   PBS.1234.0    amy   chemistry    colony    0  
Normal      Charge  
1234         1  
16
```

List Transactions

You can now check the resulting transaction records (see [Querying Transactions](#)).

Example 13. List all the job transactions

```
$ glstxn -O Job -
show="RequestId,TransactionId,Object,Action,JobId,Project,User,Machine,
```

RequestId	TransactionId	Object	Action	JobId
Project	User	Machine	Amount	
298	299	Job	Create	
298	303	Job	Quote	
chemistry	amy	colony	57600	
299	304	Job	Modify	
299	307	Job	Reserve	PBS.1234.0
chemistry	amy	colony	57600	
300	311	Job	Charge	PBS.1234.0
chemistry	amy	colony	19744	
300	312	Job	Modify	
301	314	Job	Refund	PBS.1234.0
301	315	Job	Modify	

Example 14. It may also be illustrative to examine what transactions actually composed the charge request.

```
$ glstxn -R 655 -
show="Id,Object,Action,Name,JobId,Amount,Account,Delta"
```

Id	Object	Action	Name	JobId	Amount
Account	Delta				
308	Usage	Create			
309	Reservation	Delete	PBS.1234.0		
310	Allocation	Modify	2		
311	Job	Charge	1	PBS.1234.0	19744
2					
312	Job	Modify	1		

Examine Account Statement

Finally, you can examine the account statement for the activities (see [Obtaining an Account Statement](#)).

Example 15. You can request an itemized account statement over all time for the chemistry project (account 2)

```
$ gstatement -p chemistry

#####

#
# Statement for account 2 (chemistry) generated on Tue Aug 3
# 16:06:15 2005.
#
# Reporting account activity from -infinity to now.
#
#####

Beginning Balance:          0
-----
Total Credits:              360019744
Total Debits:               -19744
-----
Ending Balance:             360000000

##### Credit Detail
#####

Object      Action      JobId      Amount      Time
-----
Account      Deposit      360000000  2005-08-03
16:01:15-07
Job          Refund      PBS.1234.0  19744       2005-08-03
```

Examine Project Usage

An additional report examines the charge totals for each user that completed jobs (see [Project Usage Summary](#)).

Example 16. Display usage by user for the chemistry project

```
$ gusage -p chemistry

#####

#
# Usage Summary for project chemistry
# Generated on Tue Feb  8 11:05:06 2005.
# Reporting user charges from 2006-07-01 to 2006-10-01
#
#####

User      Amount
-----
amy       19744
```


Getting More Advanced

In the previous chapter, a view of the system was presented that largely ignored the presence of accounts and other advanced features in Gold. This chapter will touch on the additional versatility derived from explicit use of accounts and other advanced features.



You need to be a Gold System Administrator to perform the tasks in this chapter.

Define Projects

Assume that you have created users and machines as before in the Getting Started chapter (see [Define Users](#) and [Define Machines](#)). Again you will create some projects.



In these examples, assume that the `account.autogen` configuration parameter is NOT set to automatically create a default account for each project (see [Server Configuration](#)).

Example 1. Define the project members at the same time.

For the biology project, define a set of users and a default set of machines for the project. The specified default machine will be honored within accounts associated with this project that specify MEMBERS in the machine list.

```
$ gmkproject -d "Biology Department" -u amy,bob -m blue biology
Successfully created 1 Project
```

For the chemistry projects, define a set of member users.

```
$ gmkproject -d "Chemistry Department" -u amy,bob,dave chemistry
Successfully created 1 Project
```

Use **glsproject** to see your projects.

```
$ glsproject
```

Name	Active	Users	Machines	Description
biology Department	True	amy,bob	blue	Biology
chemistry Department	True	amy,dave,bob		Chemistry



Note that accounts were not auto-generated this time because the `account.autogen` feature is set to false.

Define Accounts

Next, you can create your accounts (see [Creating Accounts](#)). Think of your accounts as bank accounts to which you can associate the users, projects, and machines that can use them.

Example 2. Create some accounts for use by the biology and chemistry projects.

```
$ gmkaccount -p biology -u MEMBERS -m MEMBERS -n "biology"
Successfully created Account 1

$ gmkaccount -p chemistry -u MEMBERS -m colony -n "chemistry on
colony"
Successfully created Account 2

$ gmkaccount -p chemistry -u amy -n "chemistry for amy"
Successfully created Account 3

$ gmkaccount -p chemistry -u MEMBERS,-amy -n "chemistry not amy"
Successfully created Account 4

$ glsaccount
```

Id	Name	Amount	Projects	Users
Machines	Description			
1	biology		biology	MEMBERS
MEMBERS				
2	chemistry on colony		chemistry	MEMBERS
colony				
3	chemistry for amy		chemistry	amy
ANY				

This shows that there is:

- a single account for biology available to all of its defined members and able to be used only on the blue machine (since blue is its only member machine).
- an account usable toward the chemistry project on the colony machine only.
- an account usable anywhere for chemistry by amy only.
- an account usable anywhere for chemistry by any member except for amy.

Make Deposits

Now you can make some deposits (see [Making Deposits](#)).

Example 3. Deposit 100 million credits for use by the biology project. Establish a use-it-or-lose-it policy in which one fourth of the credits expire each quarter. Since there is only one account for the biology project, you can specify the project name in the deposit.

```
$ gdeposit -s 2005-01-01 -e 2005-04-01 -z 25000000 -p biology
Successfully deposited 25000000 credits into account 1
$ gdeposit -s 2005-04-01 -e 2005-07-01 -z 25000000 -p biology
Successfully deposited 25000000 credits into account 1
$ gdeposit -s 2005-07-01 -e 2005-10-01 -z 25000000 -p biology
Successfully deposited 25000000 credits into account 1
$ gdeposit -s 2005-10-01 -e 2006-01-01 -z 25000000 -p biology
Successfully deposited 25000000 credits into account 1
```

Example 4. Next, make some deposits valid toward the chemistry project for the entire year. Since there are multiple accounts for the chemistry project, you must specify the appropriate account id in the deposit.

First, dedicate 50 million credits for use on colony.

```
$ gdeposit -s 2005-01-01 -e 2006-01-01 -z 50000000 -a 2
Successfully deposited 50000000 credits into account 2
```

Then give amy special access to 10 million credits that she can use anywhere — with 9 million credits prepaid and a million credits of overdraft.

```
$ gdeposit -s 2005-01-01 -e 2006-01-01 -z 9000000 -L 1000000 -a 3
Successfully deposited 9000000 credits into account 3
```

Finally, give all the other members except amy access to the remaining 40 million credits.

```
$ gdeposit -s 2005-01-01 -e 2006-01-01 -z 40000000 -a 4
Successfully deposited 40000000 credits into account 4
```

Example 5. Take a closer look at the accounts and the allocations that have been created.

```
$ glsaccount
```

Id	Name	Description	Amount	Projects	Users
1	biology		25000000	biology	MEMBERS
2	chemistry on colony		50000000	chemistry	MEMBERS

colony				
3	chemistry for amy	9000000	chemistry	amy
ANY				
4	chemistry not amy	40000000	chemistry	MEMBERS,-
amy	ANY			

Examine the allocations with the time period information.

```
$ glsalloc
```

Id	Account	StartTime	EndTime	Amount	CreditLimit
Deposited	Description				
1	1	2005-01-01	2005-04-01	25000000	0
25000000					
2	1	2005-04-01	2005-07-01	25000000	0
25000000					
3	1	2005-07-01	2005-10-01	25000000	0
25000000					
4	1	2005-10-01	2006-01-01	25000000	0
25000000					
5	2	2005-01-01	2006-01-01	50000000	0
50000000					
6	3	2005-01-01	2006-01-01	9000000	1000000
9000000					
7	4	2005-01-01	2006-01-01	40000000	0
40000000					

Check The Balance

You can examine the resulting balance (see [Querying The Balance](#)).

Example 6. View amy's balance

```
$ gbalance -u amy
```

Id Name		Amount	Reserved	Balance
CreditLimit		Available		

1	biology	25000000	0	25000000
0		25000000		
2	chemistry on colony	50000000	0	50000000
0		50000000		
3	chemistry for amy	9000000	0	9000000
1000000		10000000		

We see that amy's total balance is composed of some 25000000 credits useable toward the biology project, 50000000 for chemistry on colony, and another 10000000 which can be used for chemistry on any machine. Notice that the 10000000 credits available for use in account 3 is composed of a 9000000 balance plus an overdraft limit of 1000000 (meaning your account can go negative by that amount).

Example 7. Retrieve amy's balance for chemistry on colony.

```
$ gbalance -u amy -p chemistry -m colony -total

Balance
-----
59000000
The account balance is 60000000 credits
```

Example 8. Get the total that can be used by amy for chemistry on colony. This includes amy's available credit.

```
$ gbalance -u amy -p chemistry -m colony -total -available

Balance
-----
60000000
The account balance is 60000000 credits
```

Define Charge Rates

Gold allows you to define how much you will charge for your resources (see [Creating Charge Rates](#)).

In the Getting Started chapter, you relied on the fact that the default Gold installation predefines a Processors charge rate for you. This means that the total charge for a job will be calculated by taking the number of processors used in the job multiplied by the Processors charge rate which is then multiplied by the wallclock limit. For example:

```
( ( 16 [Processors] * 1 [ChargeRate{Resource}{Processors}] ) ) * 1234
[WallDuration] = 19744.
```

Example 9. Examine the predefined charge rates.

```
$ goldsh ChargeRate Query
```

Type	Name	Instance	Rate	Description
VBR	Processors		1	

Now you can create some of your own.

Example 10. Charge for memory used

```
$ goldsh ChargeRate Create Type=VBR Name=Memory Rate=0.001
Successfully created 1 ChargeRate
```

Example 11. You might want a quality of service multiplier

```
$ goldsh ChargeRate Create Type=NBM Name=QualityOfService
Instance=BottomFeeder Rate=0.5
Successfully created 1 ChargeRate
```

Example 12. Creating another quality-based charge multiplier

```
$ goldsh ChargeRate Create Type=NBM Name=QualityOfService
Instance=Premium Rate=2
Successfully created 1 ChargeRate
```

Example 13. View current charge rates.

```
$ goldsh ChargeRate Query
```

Type	Name	Instance	Rate
VBR	Processors		1
VBR	Memory		0.001
NBM	QualityOfService	BottomFeeder	0.5
NBM	QualityOfService	Premium	2

Obtain A Guaranteed Job Quote

This time, use the job quote to guarantee the charge rates (this may be useful in the case of fluxuating rates like market based rates).

Example 14. Request a guaranteed charge quote that reflects the memory and quality of service you expect to use.

```
$ gquote -p chemistry -u amy -m colony -P 16 -M 2048 -t 3600 -Q
Premium --guarantee

Successfully quoted 129946 credits with quote id 1
```

This time it actually created a persistent quote ...

```
$ glsquote 1

Id Amount      Job Project      User Machine      StartTime
EndTime                               WallDuration      CallType      Used
ChargeRates
Description
-----
1 129946      1 chemistry amy colony      2005-02-16 12:06:25
2005-02-23 13:06:25 3600 Normal 0
NBM:QualityOfService:Premium:2,VBR:Processors::1,VBR:Memory::0.001
```

... and created a job entry.

```
$ glsjob -j 1

Id JobId User Project      Machine Queue      QualityOfService
Stage Charge Processors Nodes WallDuration      StartTime
EndTime      Description
-----
1 amy chemistry colony Premium
Quote 16
```

Make A Quoted Job Reservation

If the quote id is specified when making the reservation, the reservation will use the quoted amounts in calculating the amount to reserve and it will connect to the existing job entry.

Example 15. Make a reservation for the job that reflects the resource and quality preferences while specifying the quote id.

```
$ greserve -J PBS.1234.0 -p chemistry -u amy -m colony -P 16 -M 2048
-t 3600 -Q Premium -q 1

Successfully reserved 129946 credits for job PBS.1234.0

$ glsres
```

Id	Name	Amount	StartTime	EndTime	
Job	User	Project	Machine	Accounts	Description
1	PBS.1234.0	129946	2005-02-16 12:35:13	2005-02-16 13:35:13	
		3	amy	chemistry	colony 3

The reservation modifies the job entry to take on the new JobId and to change its stage from Quote to Reserve.

```
$ glsjob -j 1
```

Id	JobId	User	Project	Machine	Queue	Nodes
QualityOfService	Stage	Charge	Processors			
WallDuration	StartTime	EndTime	Description			
1	PBS.1234.0	amy	chemistry	colony	Premium	
Reserve		16				

As before, the reservation will decrease the balance by the amount reserved.

```
$ gbalance -u amy -p chemistry -m colony
```

Id	Name	Amount	Reserved	Balance
CreditLimit	Available			
2	chemistry on colony	50000000	0	50000000
0	50000000			
3	chemistry for amy	8960512	129946	8830566
1000000	9830566			

Gold has two accounts to choose from. Gold will debit allocations in the order of earliest expiring and most specific first. Specifically, precedence is considered in the following order of highest to lowest: hierarchical relation, expiration time, generality of the project, generality of the user, and generality of the machine. Here, Gold considers the account that is exclusively for amy to be more specific (and of hence of higher precedence) than the account that is exclusively for the colony machine. This ordering will ensure that allocations that will expire the soonest will be used up first and that accounts with more specific access restrictions will be used in favor of accounts that have more general access (for example - amy will use up an account just for amy before she begins using a shared account).

Charge for a Quoted Job

Even if the charge rates change between submission and completion of a job, a job tied to a quote will use the quoted charge rates in a prorated manner.

Example 16. Change a charge rate and issue the charge for the job and request that the quote be honored.

```
$ goldsh ChargeRate Modify Type==VBR Name==Memory Rate=.002
Successfully modified 1 ChargeRate

$ gcharge -J PBS.1234.0 -u amy -p chemistry -m colony -P 16 -M 2048 -
t 1234 -Q Premium -X WallDuration=1234 -q 1

Successfully charged job PBS.1234.0 for 44542 credits
1 reservations were removed
```

The charge modifies the job entry with the actual usage, charges, and wallduration while changing its stage from Reserve to Charge.

```
$ glsjob -j 1
```

Id	JobId	User	Project	Machine	Queue	Nodes
QualityOfService			Stage	Charge	Processors	
WallDuration		StartTime	EndTime	Description		
3	PBS.1234.0	amy	chemistry	colony	Premium	
Charge	44542	16		1234		

The detail charge information for the job can be extracted from the transaction log.

```
$ glstxn -A Charge -J PBS.1234.0 --show Details

Details
-----
-----
-----
-----
-----
WallDuration=1234,QuoteId=1,QualityOfService=Premium,Processors=16,Item
( 16 [Processors] * 1 [ChargeRate{VBR}{Processors}] ) + ( 2048
[Memory] * 0.001 [ChargeRate{VBR}{Memory}] ) ) * 1234 [WallDuration]
* 2 [ChargeRate{QualityOfService}{Premium}] = 44542.464
```

Notice from the Itemized Charges above that the quoted memory charge rate of .001 was used instead of the current rate of .002. Notice also that the amounts have been prorated according to actual resources used and actual wallclock duration.

Partially Refund a Job

Example 17. Suppose you want to issue a partial refund.

```
$ grefund -j 1 -z 10000
Successfully refunded 10000 credits for job PBS.1234.0
```

Notice that the Job Charge is now 10000 credits lower as a result. Gold will not let your refunds total more than the total charge for the job.

```
$ glsjob 1
```

Id	JobId	User	Project	Machine	Queue	Nodes
QualityOfService	Stage	Charge	Processors			
WallDuration	StartTime	EndTime	Description			

3	PBS.1234.0	amy	chemistry	colony		Premium
Charge	34542	16		1234		

Examine Account Statement

You can get request account statement for activities as they apply to a particular account.

Example 18. You can request an itemized account statement over all time for account 3 (chemistry for amy)

```
$ gstatement -a 3

#####

#
# Statement for account 3 (chemistry for amy)
# Generated on Wed Feb 16 15:16:04 2005.
# Reporting account activity from -infinity to now.
#
#####

Beginning Balance:           0
-----
Total Credits:               9010000
Total Debits:                -44542
-----
Ending Balance:              8965458

##### Credit Detail
#####

Object      Action      JobId      Amount      Time
-----
-----
Account     Deposit      9000000    2005-02-16 15:10:44
Job         Refund       10000      2005-02-16 15:15:36

##### Debit Detail
```

Managing Users

A user is a person authorized to submit jobs to run on a high performance computing resource. User properties include the common name, phone number, email, organization, and default project for that person. A user can be created, queried, modified, and deleted.

Creating Users

To create a new user, use the command **gmkuser**:

gmkuser [-A | -I] [-n *common_name*] [-F *phone_number*] [-E *email_address*] [-p *default_project*] [-d *description*] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] {[-u] *user_name*}



It is possible to have users be created automatically when first encountered in a job function (charge, reserve or quote) by setting the user.autogen configuration parameter to true (see [Server Configuration](#)). However, bear in mind that users must be defined in order to assign them as members of a project. It is also possible to establish a system default user to be used in job functions (charge, reserve, quote) when the user is unspecified (user.default parameter).

Example 1. Creating a user

```
$ gmkuser -n "Smith, Robert F." -E "bob@western.edu" -F "(509) 555-1234" bob
```

```
Successfully created 1 User
```

Querying Users

To display user information, use the command **glsuser**:

glsuser [-A | -I] [--show *attribute_name*[,*attribute_name*...]...] [--showHidden] [--showSpecial] [-l | --long] [-w | --wide] [--raw] [--debug] [-? | --help] [--man] [--quiet] [[-u] *user_pattern*]

Example 2. Listing all info about active users

```
$ glsuser -A
Name      Active      CommonName      PhoneNumber
EmailAddress      DefaultProject      Description
-----
amy       True        Wilkes, Amy      (509) 555-8765
amy@western.edu
bob       True        Smith, Robert F. (509) 555-1234
bob@western.edu
```

Example 3. Displaying bob's phone number

```
$ glsuser --show PhoneNumber bob --quiet
(509) 555-1234
```

Example 4. Listing all user names without the header

```
$ glsuser --show Name --quiet
amy
bob
```

Example 5. Listing a user's projects

```
$ glsuser --show Projects amy -l
Projects
-----
chemistry
biology
```

Modifying Users

To modify a user, use the command **gchuser**:

gchuser [-A | -I] [-n *common_name*] [-F *phone_number*] [-E *email_address*] [-p *default_project*] [-d *description*] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] {[-u] *user_name*}

Example 6. Activating a user

```
$ gchuser -A bob
Successfully modified 1 User
```

Example 7. Changing a user's email address

```
$ gchuser -E "rsmith@cs.univ.edu" bob
Successfully modified 1 User
```

Deleting Users

To delete a user, use the command **grmuser**:

grmuser [**--debug**] [**-?** | **--help**] [**--man**] [**--quiet**] [**-v** | **--verbose**] {**[-u]** *user_name*}

Example 8. Deleting a user

```
$ grmuser bob  
Successfully deleted 1 User
```

Managing Machines

A machine is a resource that can run jobs such as a cluster or an SMP box. Machine properties include the description and whether it is active. A machine can be created, queried, modified, and deleted.

Creating Machines

To create a new machine, use the command **gmkmachine**:

gmkmachine [-A | -I] [--arch *architecture*] [--opsys *operating_system*] [-d *description*] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] {[-m] *machine_name*}



It is possible to have machines be created automatically when first encountered in a job function (charge, reserve or quote) by setting the machine.autogen configuration parameter to true (see [Server Configuration](#)). However, bear in mind that machines must be defined in order to assign them as members of a project. It is also possible to establish a system default machine to be used in job functions (charge reserve, quote) when the machine is unspecified (machine.default parameter).

Example 1. Creating a machine

```
$ gmkmachine -d "Linux Cluster" colony
Successfully created 1 Machine
```


Querying Machines

To display machine information, use the command **glsmachine**:

glsmachine [-A | -I] [--show *attribute_name*[,*attribute_name*...]...] [--showHidden] [--showSpecial] [--raw] [--debug] [-? | --help] [--man] [--quiet] [[-m] *machine_pattern*]

Example 2. Listing all inactive machine names and descriptions

```
$ glsmachine -I --show Name,Description
Name      Description
-----
inert      This machine is unusable
```

Modifying Machines

To modify a machine, use the command **gchmachine**:

gchmachine [-A | -I] [--arch *architecture*] [--opsys *operating_system*] [-d *description*] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] {[-m] *machine_name*}

Example 3. Deactivating a machine

```
$ gchmachine -I colony  
Successfully modified 1 Machine
```

Deleting Machines

To delete a machine, use the command **grmmachine**:

grmmachine [**--debug**] [**-?** | **--help**] [**--man**] [**--quiet**] [**-v** | **--verbose**] **{** [**-m**] *machine_name* **}**

Example 4. Deleting a machine

```
$ grmmachine colony  
Successfully deleted 1 Machine
```

Managing Projects

A project is a research interest or activity requiring the use of computational resources for a common purpose. Users may be designated as members of a project and allowed to share its allocations. The project user list will be honored within accounts including the project that specify MEMBERS in the user list. Machines may also be designated as members of a project as a default resource pool. The project machine list will be honored within accounts including the project that specify MEMBERS in the machine list.

Creating Projects

To create a new project, use the command **gmkproject**:

gmkproject [-A | -I] [-u [+ | -]user_name [, [+ | -]user_name...]] [-m [+ | -]machine_name [, [+ | -]machine_name...]] [-d description] [--createAccount=True|False] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] {[-p] project_name}



If the account.autogen configuration parameter is set to true (see [Server Configuration](#)), an account will be automatically created for the project (unless overridden with the `--createAccount` option). The auto-generated account will be associated with the new project, the user MEMBERS of the project and ANY machine.



It is possible to have projects be created automatically when first encountered in a job function (charge, reserve, or quote) by setting the project.autogen configuration parameter to true (see [Server Configuration](#)). It is also possible to establish a system default project (project.default) to be used in job functions (charge, reserve, quote) when the project is unspecified and the user does not have a default project.

Example 1. Creating a project

```
$ gmkproject -d "Chemistry Department" chemistry
Successfully created 1 Project
```

Example 2. Creating a project and specifying user members at the same time

```
$ gmkproject -d "Chemistry Department" -u amy,bob,dave chemistry
Successfully created 1 Project
```

Querying Projects

To display project information, use the command **glsproject**:

glsproject [-A | -I] [--show *attribute_name* [,*attribute_name*...]...] [--showHidden] [--showSpecial] [-l | -long] [-w | -wide] [--raw] [--debug] [-? | -help] [--man] [--quiet] [[-p] *project_pattern*]

Example 3. Listing all info about all projects

```
$ glsproject
```

Name	Active	Users	Machines	Description
biology	True	amy,bob	colony	Biology
Department				
chemistry	True	amy,dave,bob		Chemistry
Department				

Example 4. Displaying the name and user members of a project in long format

```
$ glsproject --show Name,Users -l chemistry
```

Name	Users
chemistry	bob
	dave
	amy

Example 5. Listing all project names

```
$ glsproject --show Name --quiet
```

biology
chemistry

Modifying Projects

To modify a project, use the command **gchproject**:

```
gchproject [-A | -I] [-d description] [--addUser(s) [+ | -]user_name [, [+ | -]user_name...]] [--addMachines(s) [+ | -]machine_name [, [+ | -]machine_name...]] [--delUser(s) user_name [,user_name...]] [--delMachines(s) machine_name [,machine_name...]] [--actUser(s) user_name [,user_name...]] [--actMachines(s) machine_name [,machine_name...]] [--deactUser(s) user_name [,user_name...]] [--deactMachines(s) machine_name [,machine_name...]] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] {[-p] project_name}
```

Example 6. Deactivating a project

```
$ gchproject -I chemistry
Successfully modified 1 Project
```

Example 7. Adding users as members of a project

```
$ gchproject --addUsers jsmith,barney chemistry
Successfully created 2 ProjectUsers
```

Example 8. Adding machines as members of a project

```
$ gchproject --addMachines colony chemistry
Successfully created 1 ProjectMachines
```

Deleting Projects

To delete a project, use the command **grmproject**:

grmproject [**--debug**] [**-?** | **--help**] [**--man**] [**--quiet**] [**-v** | **--verbose**] **{[-p] *project_name*}**

Example 9. Deleting a project

```
$ grmproject chemistry
Successfully deleted 1 Project
```

Project Usage Summary

To generate a project usage summary broken down by user, use the command **gusage**. This report lists the total charges by each of the active users during the specified time frame.

gusage [-s *start_time*] [-e *end_time*] [-h | --hours] [--debug] [-? | --help] [--man] {[-p] *project_name*}

Example 10. Displaying a usage summary for the chemistry project during the third quarter of 2006

```
$ gusage -p chemistry -s 2006-07-01 -e 2006-10-01

#####
#
# Usage for project chemistry
# Generated on Tue Feb  8 11:05:06 2005.
# Reporting user charges from 2006-07-01 to 2006-10-01
#
#####

User      Amount
-----
amy       19744
bob       36078
```


Managing Accounts

An account is a container for time-bounded resource credits valid toward a specific set of projects, users, and machines. Much like with a bank, an account is a repository for resource credits. Each account has a set of access control lists designating which users, projects, and machines may access the account. An account may restrict the projects that can charge to it. Normally an account will be tied to a single project but it may be tied to an arbitrary set of projects or ANY project. An account may restrict the users that can charge to it. It will frequently be tied to the the user MEMBERS of the associated project(s) but it may be tied to an arbitrary set of users or ANY user. An account may restrict the machines that can charge to it. It may be tied to an arbitrary set of machines, just the machine MEMBERS of the associated project(s) or ANY machine.

When resource credits are deposited into an account, they are associated with a time period within which they are valid. These time-bounded pools of credits are known as allocations. (An allocation is a pool of resource credits associated with an account for use during a particular time period.) By using multiple allocations that expire in regular intervals it is possible to implement a use-it-or-lose-it policy and establish a project cycle.

Accounts may be nested. Hierarchically nested accounts may be useful for the delegation of management roles and responsibilities. Deposit shares may be established that assist to automate a trickle-down effect for funds deposited at higher level accounts. Additionally, an optional overflow feature allows charges against lower level accounts to trickle up the hierarchy.

Operations include creating, querying, modifying, and deleting accounts as well as making deposits, withdrawals, transfers, and balance queries.

Creating Accounts

gmkaaccount is used to create a new account. A new id is automatically generated for the account.

gmkaaccount [-n *account_name*] [-p [+ | -]*project_name* [, [+ | -]*project_name*...]] [-u [+ | -]*user_name* [, [+ | -]*user_name*...]] [-m [+ | -]*machine_name* [, [+ | -]*machine_name*...]] [-d *description*] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose]



When creating an account, it is important to specify at least one user, machine, and project designation. If omitted, these will default to ANY.



It is possible to have accounts be created automatically when projects are created by setting the `account.autogen` configuration parameter to true (see [Server Configuration](#)). The auto-generated account will be associated with the new project, the user MEMBERS of the project and ANY machine.

Example 1. Creating an account

```
$ gmkaaccount -p chemistry -u MEMBERS -m ANY -n "Chemistry"

Successfully created 1 Account
Successfully created 1 AccountProject
Successfully created 1 AccountUser
Successfully created 1 AccountMachine
```

Example 2. Creating a wide-open account

```
$ gmkaaccount -p ANY -u ANY -m ANY -n "Cornucopia"

Successfully created 1 Account
Successfully created 1 AccountProject
Successfully created 1 AccountUser
Successfully created 1 AccountMachine
```

Example 3. Creating an account valid toward all biology project members except for dave and all machines except for blue

```
$ gmkaccount -p biology -u MEMBERS,-dave -m ANY,-blue -n "Not Dave"

Successfully created 1 Account
Successfully created 1 AccountProject
Successfully created 1 AccountUser
Successfully created 1 AccountUser
Successfully created 1 AccountMachine
Successfully created 1 AccountMachine
```

Querying Accounts

To display account information, use the command **glsaccount**:

glsaccount [-A | -I] [-n *account_name*] [-p *project_name*] [-u *user_name*] [-m *machine_name*] [-s *start_time*] [-e *end_time*] [--exact-match] [--show *attribute_name* [,*attribute_name*...]...] [--showHidden] [-l | -long] [-w | -wide] [-raw] [-h | -hours] [--debug] [-? | -help] [--man] [--quiet] [[-a] *account_id*]

Example 4. Listing all info about all accounts with multi-valued fields displayed in a multi-line format

```
$ glsaccount -long
```

Id	Name	Amount	Projects	Users	Machines
Description					
1	Biology	360000000	biology	MEMBERS	blue
2	Chemistry	360000000	chemistry	MEMBERS	ANY
3	Cornucopia	0	ANY	ANY	ANY
4	Not Dave	250000	biology	-dave	-blue

Example 5. Listing all info about all accounts useable by dave

```
$ glsaccount -u dave -long
```

Id	Name	Amount	Projects	Users	Machines
Description					
2	Chemistry	360000000	chemistry	MEMBERS	ANY
3	Cornucopia	0	ANY	ANY	ANY

Modifying Accounts

To modify an account, use the command **gchaccount**:

```
gchaccount [-n account_name] [-d description] [--addProject(s) [+ | -]project_name [, [+ | -]  
] project_name...] [--addUser(s) [+ | -]user_name [, [+ | -]user_name...] [--addMachine(s) [+ | -]  
] machine_name [, [+ | -]machine_name...] [--delProject(s) project_name [, project_name...] [--delUser(s)  
user_name [, user_name...] [--delMachine(s) machine_name [, machine_name...] [--debug] [-? | --help] [--  
man] [--quiet] [-v | --verbose] {[-a] account_id}
```

Example 6. Adding a user to the list of users that share the account

```
$ gchaccount --addUser dave 1  
  
Successfully created 1 AccountUser
```

Making Deposits

gdeposit is used to deposit time-bounded resource credits into accounts resulting in the creation or enlargement of an allocation. (See [Allocations](#) for managing allocations). The start time will default to -infinity and the end time will default to infinity if not specified. Accounts must first be created using **gmkaaccount** (unless auto-generated).

gdeposit {-a *account_id* | -p *project_name*} [-i *allocation_id*] [-s *start_time*] [-e *end_time*] [[-z] *amount*] [-L *credit_limit*] [-d *description*] [-h | -hours] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose]

Example 7. Making a deposit

```
$ gdeposit -s 2003-10-01 -e 2004-10-01 -z 360000000 -a 1
Successfully deposited 360000000 credits into account 1
```

Example 8. Making a deposit "into" a project

If a project has a single account then a deposit can be made against the project.

```
$ gdeposit -s 2003-10-01 -e 2004-10-01 -z 360000000 -p chemistry
Successfully deposited 360000000 credits into account 2
```

Example 9. Creating a credit allocation

```
$ gdeposit -L 10000000000 -a 3
Successfully deposited 0 credits into account 3
```

Querying The Balance

To display balance information, use the command **gbalance**:

gbalance [-p *project_name*] [-u *user_name*] [-m *machine_name*] [--total] [--available] [--raw] [-h | -hours] [--debug] [--? | --help] [--man] [--quiet]

Example 10. Querying the project balance detail broken down by account

```
$ gbalance -p chemistry

Id  Name          Amount      Reserved      Balance
CreditLimit Available
-----
1   Chemistry     360000000    0              360000000    0
360000000
2   Cornucopia    0            0              0
1000000000000  1000000000000
```

Example 11. Querying the total balance for a particular user in a particular project on a particular machine

```
$ gbalance -u bob -m colony -p chemistry --total

Balance
-----
360000000
The account balance is 360000000 credits
```

Example 12. List the projects and available balance amy can charge to

```
$ gbalance -u amy --show Project,Balance

Project      Balance
-----
biology      360000000
chemistry    360000000
```

Personal Balance

The **mybalance** has been provided as a wrapper script to show users their personal balance. It provides a list of balances for the projects that they can charge to:

gbalance [-h | -hours] [-? | -help] [-man]

Example 13. List my (project) balances

```
$ mybalance
Project      Balance
-----
biology      324817276
chemistry    9999979350400
```

Example 14. List my balance in (Processor) hours

```
$ mybalance -h
Project      Balance
-----
biology      90227.02
chemistry    2777772041.77
```

Making Withdrawals

To issue a withdrawal, use the command **gwithdraw**:

gwithdraw {-a *account_id* | -p *project_name*} [-i *allocation_id*] {[-z] *amount*} [-d *description*] [-h | -hours] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose]

Example 15. Making a withdrawal

```
$ gwithdraw -z 12800 -a 1 -d "Grid Tax"
Successfully withdrew 12800 credits from account 1
```

Example 16. Making a withdrawal "from" a project

If a project has a single account then a withdrawal can be made against the project.

```
$ gwithdraw -z 12800 -p chemistry
Successfully withdrew 12800 credits from account 2
```


Making Transfers

To issue a transfer between accounts, use the command **gtransfer**. If the allocation id is specified, then only credits associated with the specified allocation will be transferred, otherwise, only active credits will be transferred. Account transfers preserve the allocation time periods associated with the resource credits from the source to the destination accounts. If a one-to-one mapping exists between project and account, then the fromProject/toProject options may be used in place of the fromAccount/toAccount options.

```
gtransfer {--fromAccount source_account_id | --fromProject source_project_name | -i allocation_id}  
{--toAccount destination_account_id | --toProject destination_project_name} [-d description] [-h | -  
hours] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] {[ -z] amount}
```

Example 17. Transferring credits between two accounts

```
$ gtransfer --fromAccount 1 --toAccount 2 10000  
Successfully transferred 10000 credits from account 1 to account 2
```

Example 18. Transferring credits between two single-account projects

```
$ gtransfer --fromProject biology --toProject chemistry 10000  
Successfully transferred 10000 credits from account 1 to account 2
```

Obtaining an Account Statement

To generate an account statement, use the command **gstatement**. For a specified time frame it displays the beginning and ending balances as well as the total credits and debits to the account over that period. This is followed by an itemized report of the debits and credits. Summaries of the debits and credits will be displayed instead of the itemized report if the `—summarize` option is specified. If a project, user or machine is specified instead of an account, then the statement will consist of information merged from all accounts valid toward the specified entities.

gstatement `[[-a] account_id] [-p project_name] [-u user_name] [-m machine_name] [-s start_time] [-e end_time] [—summarize] [-h | —hours] [—debug] [-? | —help] [—man]`

Example 19. Generating an account statement for the third quarter of 2006

```
$ gstatement -a 2 -s 2006-07-01 -e 2006-10-01

#####
#
# Statement for account 2 (chemistry) generated on Tue Aug 3
# 16:06:15 2005.
#
# Reporting account activity from -infinity to now.
#
#####

Beginning Balance:          0
-----
Total Credits:              360019744
Total Debits:               -19744
-----
Ending Balance:             360000000

##### Credit Detail
#####

Object      Action      JobId      Amount      Time
-----
Account      Deposit      360000000  2005-08-03
16:01:15-07
Job          Refund      PBS.1234.0  19744       2005-08-03
```

Deleting Accounts

To delete an account, use the command **grmaccount**:

grmaccount [**—debug**] [**-?** | **—help**] [**—man**] [**—quiet**] [**-v** | **—verbose**] {**[-a]** *account_id*}

Example 20. Deleting an account

```
$ grmaccount 2
Successfully deleted 1 Account
```

Managing Allocations

An allocation is a time-bounded pool of resource credits associated with an account. An account may have multiple allocations, each for use during a different time period. An allocation may also have a credit limit representing the amount by which it can go negative.

Operations include querying, modifying, and deleting allocations.

Creating Allocations

Allocations are created by making account deposits via the [gdeposit](#) command (See [Making Deposits](#)).

Querying Allocations

To display allocation information, use the command **glsalloc**:

glsalloc [-A | -I] [-a *account_id*] [-p *project_name*] [--show *attribute_name* [,*attribute_name*...]...] [--showHidden] [--raw] [-h | --hours] [--debug] [-? | --help] [--man] [--quiet] [--i *allocation_id*]

Example 1. Listing allocations for account 4

\$ glsalloc -a 4

Id	Account	StartTime	EndTime	Amount	CreditLimit
		Deposited	Description		
4	4	2005-01-01	2005-04-01	250000	0
5	4	2005-04-01	2005-07-01	250000	0
6	4	2005-07-01	2005-10-01	250000	0
7	4	2005-10-01	2006-01-01	250000	0

Modifying Allocations

To modify an allocation, use the command **gchalloc**:

gchalloc [-s *start_time*] [-e *end_time*] [-L *credit_limit*] [-d *description*] [-h | --hours] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] {[-i] *allocation_id*}

Example 2. Changing the end time for an allocation

```
$ gchalloc -e "2005-01-01" 4
Successfully modified 1 Allocation
```

Example 3. Changing the credit limit for an allocation

```
$ gchalloc -L 500000000000 -i 2
Successfully modified 1 Allocation
```

Deleting Allocations

To delete an allocation, use the command **grmalloc**:

grmalloc [**--debug**] [**-?** | **--help**] [**--man**] [**--quiet**] [**-v** | **--verbose**] **{-I | [-i] *allocation_id*}**

Example 4. Deleting an allocation

```
$ grmalloc 4  
Successfully deleted 1 Allocation
```

Example 5. Purging inactive allocations

```
$ grmalloc -I  
Successfully deleted 2 Allocations
```

Managing Reservations

A reservation is a hold placed against an account. Before a job runs, a reservation (or hold) is made against one or more of the requesting user's applicable account(s). Subsequent jobs will also post reservations while the available balance (active allocations minus reservations) allows. When a job completes, the reservation is removed and the actual charge is made to the account(s). This procedure ensures that jobs will only run so long as they have sufficient reserves.

Associated with a reservation is the name of the reservation (often the job id requiring the reservation), the user, project, and machine as applicable, an expiration time, and an amount. Operations include creating, querying, modifying, and deleting reservations.

Creating Reservations

Reservations are created by the resource management system with the [greserve](#) command (See [Making Job Reservations](#)).

Querying Reservations

To display reservation information, use the command **glsres**:

glsres [-A | -I] [-n *reservation_name* / *job_id_pattern*] [-p *project_name*] [-u *user_name*] [-m *machine_name*] [--show *attribute_name* [,*attribute_name*...]...] [--showHidden] [-l | --long] [-w | --wide] [--raw] [-h | --hours] [--debug] [-? | --help] [--man] [--quiet] [--r *reservation_id*]

Example 1. Listing all info about all reservations for bob

```
$ glsres -u bob
```

Id		Name	Amount		StartTime		Machine	
EndTime			Job	User	Project			
Accounts	Description							

1	Interactive.789654		3600		2005-01-13	16:48:15		
2005-01-13	17:48:15	1	bob	chemistry	blue	1		

Example 2. Listing all info about all reservations that impinge against amy's balance

```
$ glsres -u amy --option name=UseRules value=True
```

Id		Name	Amount		StartTime		Machine	
EndTime			Job	User	Project			
Accounts	Description							

1	Interactive.789654		3600		2005-01-13	16:48:15		
2005-01-13	17:48:15	1	bob	chemistry	blue	1		
2	PBS.1234.0		7200		2005-01-13	17:59:09		
2005-01-14	02:28:41	2	amy	chemistry	colony	2		

Modifying Reservations

To modify a reservation, use the command **gchres**:

gchres [-s *start_time*] [-e *end_time*] [-d *description*] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] {[-r] *reservation_id*}

Example 3. Changing the expiration time of a reservation

```
$ gchres -e "2004-08-07 14:43:02" 1  
Successfully modified 1 Reservation
```

Deleting Reservations

To delete a reservation, use the command **grmres**:

```
grmres [--debug] [-? | --help] [--man] [-q | --quiet] [-v | --verbose] {-I | -n reservation_name /  
job_id | [-r] reservation_id}
```

Example 4. Deleting a reservation by name (JobId)

```
$ grmres -n PBS.1234.0  
Successfully deleted 1 Reservation
```

Example 5. Deleting a reservation by ReservationId

```
$ grmres 1  
Successfully deleted 1 Reservation
```

Example 6. Purging stale reservations

```
$ grmres -I  
Successfully deleted 2 Reservations
```

Managing Quotations

A quotation provides a way to determine beforehand how much would be charged for a job. When a quotation is requested, the charge rates applicable to the job requesting the quote are saved and a quote id is returned. When the job makes a reservation and the final charge, the quote can be referenced to ensure that the saved charge rates are used instead of current values. A quotation has an expiration time after which it cannot be used. A quotation may also be used to verify that the given job has sufficient funds and meets the policies necessary for the charge to succeed.

Operations include querying, modifying, and deleting quotations.

Creating Quotations

Quotations are normally created by the resource management system with the `gquote` command (See [Making Job Quotations](#)).

Querying Quotations

To display quotation information, use the command **glsquote**:

glsquote [-A | -I] [-p *project_name*] [-u *user_name*] [-m *machine_name*] [--show *attribute_name* [,*attribute_name*...]...] [--showHidden] [-l | --long] [-w | --wide] [--raw] [-h | --hours] [--debug] [-? | --help] [--man] [--quiet] [--q] *quote_id*

Example 1. Listing all info about all quotes for user amy on machine colony

```
$ glsquote -u amy -m colony
```

Id	Amount	Job	Project	User	Machine	StartTime	EndTime	WallDuration	Type	Used
ChargeRates	Description									
1	57600	1	chemistry	amy	colony	2005-01-14 10:09:58	2005-09-10 15:27:07	3600	Normal	0

VBR:Processors:1

Modifying Quotations

To modify a quotation, use the command **gchquote**:

gchquote [-s *start_time*] [-e *expiration_time*] [-d *description*] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] {[-q] *quote_id*}

Example 2. Changing the expiration time of a quotation

```
$ gchquote -e "2005-03-01" 1
Successfully modified 1 Quotation
```

Deleting Quotations

To delete a quotation, use the command **grmquote**:

grmquote [**—debug**] [**—?** | **—help**] [**—man**] [**—quiet**] [**—v** | **—verbose**] {**—I** | [**—q**] *quote_id*}

Example 3. Deleting a quotation

```
$ grmquote 1
Successfully deleted 1 Quotation
```

Example 4. Purging stale quotations

```
$ grmquote -I
Successfully deleted 2 Quotations
```

Managing Jobs

Gold can track the jobs that run on your system, recording the charges and resources used for each job. Typically, a job record is created when the resource manager charges for a job. Job quotes, reservations, charges, and refunds can be issued.

Creating Jobs

In most cases, jobs will be created by the resource management system with the [greserve](#) command or the [gcharge](#) command.

However, it is also possible to create job records by hand using the **gmckjob** command:

```
gmckjob [-u user_name] [-p project_name] [-m machine_name] [-o organization] [-C queue_name] [-Q quality_of_service] [-P processors] [-N nodes] [-M memory] [-D disk] [-n job_name] [--application application] [--executable executable] [-t wallclock_duration] [-s start_time] [-e end_time] [-T job_type] [-d description] [-X | --extension property=value...] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] [-V | --version] [[-J] job_id]
```

Example 1. Creating a job record

```
$ gmckjob -u jsmith -p chem -m cluster -X Charge=2468 -P 2 -t 1234 -J  
PBS.1234.0  
  
Successfully created Job 102
```


Querying Jobs

To display job information, use the command **glsjob**:

glsjob *[(-J) job_id_pattern] [-p project_name] [-u user_name] [-m machine_name] [-C queue] [-T type]*
[--stage stage] [-s start_time] [-e end_time] [--show attribute_name[,attribute_name...]...] [-
showHidden] [--raw] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] [-V | --version] [(-j)
gold_job_id]

Example 2. Show specific info about jobs run by amy

```
$ glsjob --show=JobId,Project,Machine,Charge -u amy
```

JobId	Project	Machine	Charge
PBS.1234.0	chemistry	colony	0

Modifying Jobs

It is possible to modify a job record by using the command **gchjob**:

gchjob [-u *user_name*] [-p *project_name*] [-m *machine_name*] [-o *organization*] [-C *queue_name*] [-Q *quality_of_service*] [-P *processors*] [-N *nodes*] [-M *memory*] [-D *disk*] [-n *job_name*] [--application *application*] [--executable *executable*] [-t *wallclock_duration*] [-s *start_time*] [-e *end_time*] [-T *job_type*] [-d *description*] [-X | --extension *property=value...*] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] [-V | --version] [[-J] *job_id*]

Example 3. Changing a job

```
$ gchjob -Q HalfPrice --application=NwChem -X Charge=1234 -d  
"Benchmark" -J PBS.1234.0  
  
Successfully modified 1 Job
```

Deleting Jobs

To delete a job, use the command **grmjob**:

grmjob [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] [-V | --version] [[-J] *job_id*]

Example 4. Deleting a job

```
$ grmjob -J PBS.1234.0
Successfully deleted 1 Job
```

Obtaining Job Quotes

Job quotes can be used to determine how much it will cost to run a job. This step verifies that the submitter has sufficient funds for, and meets all the allocation policy requirements for running the job and can be used at job submission as an early filter to prevent jobs from getting in and waiting in the job queue just to be blocked from running later. If a guaranteed quote is requested, a quote id is returned and can be used in the subsequent charge to guarantee the rates that were used to form the original quote. A guaranteed quote has the side effect of creating a quotation record and a permanent job record.

To request a job quote, use the command **gquote**:

```
gquote [-u user_name] [-p project_name] [-m machine_name] [-o organization] [-C queue_name] [-Q quality_of_service] [-P processors] [-N nodes] [-M memory] [-D disk] [-n job_name] [--application application] [-t wallclock_duration] [-s start_time] [-e end_time] [-T job_type] [-d description] [-X | --extension property=value...] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] [-V | --version] [[-J] job_id]
```

Example 5. Requesting a quotation

```
$ gquote -p chemistry -u amy -m colony -P 2 -t 3600  
Successfully quoted 7200 credits
```

Example 6. Requesting a guaranteed quote

```
$ gquote -p chemistry -u amy -m colony -P 16 -t 3600 --guarantee  
Successfully quoted 57600 credits with quote id 1  
  
$ glsquote  


| Id          | Amount | Job | Project          | User                | Machine | StartTime  |
|-------------|--------|-----|------------------|---------------------|---------|------------|
| EndTime     |        |     |                  | WallDuration        | Type    | Used       |
| ChargeRates |        |     |                  | Description         |         |            |
| 1           | 57600  | 1   | chemistry        | amy                 | colony  | 2005-01-14 |
| 10:09:58    |        |     |                  | 2005-08-10 15:27:07 | 3600    |            |
| Normal      | 0      |     | VBR:Processors:1 |                     |         |            |


```



It is possible to establish a system default machine, project or user to be used in job functions (charge, reserve or quote) when left unspecified (see [Server Configuration](#)).

Making Job Reservations

A job reservation can be used to place a hold on the user's account before a job starts to ensure that the credits will be there when it completes.

To create a job reservation use the command **greserve**:

```
greserve [-u user_name] [-p project_name] [-m machine_name] [-o organization] [-C queue_name] [-Q quality_of_service] [-P processors] [-N nodes] [-M memory] [-D disk] [-n job_name] [--application application] [-t wallclock_duration] [-s start_time] [-e end_time] [-T job_type] [-d reservation_description] [-X | --extension property=value...] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] [-V | --version] [[-J] job_id]
```

Example 7. Creating a reservation

```
$ greserve -J PBS.1234.0 -p chemistry -u amy -m colony -P 2 -t 3600  
Successfully reserved 7200 credits for job PBS.1234.0
```



It is possible to establish a system default machine, project or user to be used in job functions (charge, reserve or quote) when left unspecified (see [Server Configuration](#)).

Charging Jobs

A job charge debits the appropriate allocations based on the user, project and machine associated with the job. The charge is calculated based on factors including the resources used, the job run time, and other quality-based factors (See [Managing Charge Rates](#)).

To charge for a job use the command **gcharge**:

```
gcharge [-u user_name] [-p project_name] [-m machine_name] [-o organization] [-C queue_name] [-Q quality_of_service] [-P processors] [-N nodes] [-M memory] [-D disk] [-S job_state] [-n job_name] [-T job_type] [--application application] [--executable executable] [-t charge_duration] [-s charge_start_time] [-e charge_end_time] [-d reservation_description] [-X | --extension property=value...] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] [-V | --version] [-q quote_id] [-r reservation_id] [[-j] gold_job_id] {-J job_id}
```

Example 8. Issuing a job charge

```
$ gcharge -J PBS.1234.0 -p chemistry -u amy -m colony -P 2 -t 1234 -X  
WallDuration=1234  
  
Successfully charged job PBS.1234.0 for 2468 credits  
1 reservations were removed
```



It is possible to establish a system default machine, project or user to be used in job functions (charge, reserve or quote) when left unspecified (see [Server Configuration](#)).

Issuing Job Refunds

A job can be refunded in part or in whole by issuing a job refund. This action attempts to lookup the referenced job to ensure that the refund does not exceed the original charge and so that the charge entry can be updated. If multiple matches are found (such as the case when job ids are non-unique), this command will return the list of matched jobs with unique ids so that the correct job can be specified for the refund.

To issue a refund for a job, use the command **grefund**:

grefund [-z *amount*] [-a *account_id*] [-d *description*] [-h | --hours] [--debug] [-? | --help] [--man] [--quiet] [-v | --verbose] [-V | --version] [[-J] *job_id* | [-j] *gold_job_id*]

Example 9. Issuing a job refund

```
$ grefund -J PBS.1234.0
Successfully refunded 19744 credits for job PBS.1234.0
```

Managing Charge Rates

Charge Rates establish how much to charge for usage. There are nine main types of charge rates: Value Based Resources, Name Based Resources, Value Based Usage, Name Based Usage, Value Based Multipliers, Name Based Multipliers, Value Based Fees, Name Based Fees and Multi-dimensional Value Based Resources.

- **Value Based Resource** — Value Based Resource (or Consumable Resource) Charge Rates define how much it costs per unit of time to use a consumable resource like processors, memory, telescope time, generic resources that have a count and are charged per time used, etc. These resource metrics must first be multiplied by the wallclock duration before being added to the total charge. Value Based Resource Charge Rates are of Type "VBR", with the Name being the resource name (such as Processors) and the given Rate (such as 1) being multiplied by the consumed resource value (such as 8).
- **Name Based Resource** — Name Based Resource Charge Rates define how much it costs per unit of time to use a named resource like license, etc. The cost for the named resource must first be multiplied by the wallclock duration before being added to the total charge. Name Based Resource Charge Rates are of Type "NBR", with the Name being the resource name (such as License), with the Instance being the resource value (such as matlab), and having the given Rate (such as 5).
- **Value Based Usage** — Value Based Usage Charge Rates define how much to charge for metrics of total resource usage such as cputime, power consumed, generic resources or licenses that are charged flat fees per use, etc. These usage metrics are added to the total charge without being multiplied by wall duration. Value Based Usage Charge Rates are of Type "VBU", with the Name being the resource name (such as Power) and the given Rate (such as .001) being multiplied by the consumed resource value (such as 40000).
- **Name Based Usage** — Name Based Usage Charge Rates define how much it costs to use a named attribute having a flat charge such as feature, etc. These usage metrics are added to the total charge without being multiplied by wall duration. Name Based Usage Charge Rates are of Type "NBU", with the Name being the resource name (such as Feature), with the instance being the usage value (such as GPU), and having the given flat Rate (such as 200).
- **Value Based Multiplier** — Value Based Multiplier Charge Rates are scaled multipliers which apply a multiplicative charge factor based on a numeric scaling factor. These incoming scaling factors are multiplied against the Value-Based Multiplier Rate and then are multiplied against the total of the resource and usage charges for the job. Value Based Multiplier Charge Rates are of Type "VBM", with the Name being the multiplier name (such as Discount) and the given Rate (such as 1) being multiplied with the scaling factor (such as .5) before being multiplied to the total job charge.
- **Name Based Multiplier** — Name Based Multiplier Charge Rates are quality based multipliers which apply a multiplicative charge factor based on a quality of the job such as quality of service, nodetype, queue, user, time of day, etc. These charge multipliers are determined by a hash or lookup table based on the value of the job attribute. These rates are multiplied against the total of the resource and usage charges for the job. Name Based Multiplier Charge Rates are of Type "NBM", with the Name being the quality name (such as QualityOfService), with the Instance being the quality instance (such as Premium), and having the given multiplier Rate (such as 2).
- **Value Based Fee** — Value Based Fee Charge Rates define how much to charge for scaled or enumerated fees such as setup fees, shipping charges, etc. which should be added after the multipliers are applied. These fees are added to the total charge. Value Based Fee Charge Rates are of Type "VBF", with the Name being the fee name (such as Shipping) and the given Rate (such as 25) being multiplied by the scaling or counted value (such as 4).
- **Name Based Fee** — Name Based Fee Charge Rates define how much it costs to use a named attribute having a flat charge such as feature, etc. which should be added after the multipliers are applied. These fees are added to the total charge. Name Based Fee Charge Rates are of Type "NBF", with the Name being the fee name (such as Zone), with the instance being the fee value (such as Asia), and having the given flat Rate (such as 100).
- **Multi-dimensional Value Based Resource** — Multi-dimensional Value Based Resource Charge Rates applies a consumable resource cost that varies depending on the value of a separate named job property. These resource metrics will first be multiplied by the wallclock duration before being added to the total charge. For example, using this capability you can apply different processor rates for different

users or machines, or different disk prices for different queues. Multi-dimensional Value Based Resource Charge Rates have the Type being the consumable resource (such as Processors), the Rate being the cost of this resource (such as 1.5), the Name being the name of the controlling job property (such as User), and the Instance being the value of the controlling job property (such as frank).

By default, job charges are calculated according to the following formula: For each Value Based Resource Charge Rate applicable to a given job, a value-based resource charge is calculated by multiplying the amount of the resource used by the amount of time it was used, multiplied by the charge rate for that resource. For each Name Based Resource Charge Rate applicable to a given job, a name-based resource charge is calculated by multiplying the charge rate for that named resource by the amount of time it was used. For each Value Based Usage Charge Type applicable to a given job, a value-based usage charge is calculated by multiplying the amount of the usage by the charge rate for that usage. For each Name Based Usage Charge Type applicable to a given job, a name-based usage charge is given by the charge rate for that usage. For each Multi-dimensional Value Based Resource Charge Rate applicable to a given job, a value-based resource charge is calculated by multiplying the amount of the resource used by the amount of time it was used, multiplied by the charge rate for that resource. These value-based, name-based and multi-dimensional value-based resource charges and the value-based and name-based usage charges are added together. Then, for each Value Based Multiplier Charge Rate applicable to the job, a value-based multiplier is calculated by multiplying the amount of the multiplier by the charge rate for that multiplier. For each Name Based Multiplier Charge Rate applicable to the job, a name-based multiplier is given by charge rate for that multiplier. The sum of the resource and usage charges is then multiplied by each of the applicable value-based and name-based multipliers. Next, for each Value Based Fee Charge Type applicable to a given job, a value-based fee charge is calculated by multiplying the amount of the fee by the charge rate for that fee. For each Name Based Fee Charge Type applicable to a given job, a name-based fee charge is given by the charge rate for that fee. Finally, these value-based and name-based fee charges are to the total job charge.

In short, the formula can be represented by
$$((((\sum(VBR \cdot value) + \sum(NBR) + \sum(MVBR \cdot value)) \cdot wall_duration) + (\sum(VBU \cdot value) + \sum(NBU))) \cdot \Pi(VBM \cdot value) \cdot \Pi(NBM)) + (\sum(VBF \cdot value) + \sum(NBF)))$$

Creating ChargeRates

To create a new charge rate, use the command **goldsh ChargeRate Create**:

```
goldsh ChargeRate Create Type=<Charge Rate Type> Name=<Charge Rate Name> [Instance=<Floating Point Multiplier>] Rate=<Floating Point Multiplier> [Description=<Description>] [ShowUsage:=True]
```

Example 1. Creating a couple of value-based resource charge rates

```
$ goldsh ChargeRate Create Type=VBR Name=Processors Rate=1
Successfully created 1 ChargeRate
```

```
$ goldsh ChargeRate Create Type=VBR Name=Processors Rate=0.001
Successfully created 1 ChargeRate
```

Example 2. Creating a name-based resource charge rate

```
$ goldsh ChargeRate Create Type=NBR Name=License Instance=Matlab
Rate=5
Successfully created 1 ChargeRate
```

Example 3. Creating a couple of value-based usage charge rates

```
$ goldsh ChargeRate Create Type=VBU Name=Power Rate=0.001
Successfully created 1 ChargeRate
```

```
$ goldsh ChargeRate Create Type=VBU Name=CpuTime Rate=1  
Successfully created 1 ChargeRate
```

Example 4. Creating a name-based usage charge rate

```
$ goldsh ChargeRate Create Type=NBU Name=Feature Instance=GPU Rate=200  
Successfully created 1 ChargeRate
```

Example 5. Creating a value-based multiplier charge rate

```
$ goldsh ChargeRate Create Type=VBM Name=Discount Rate=1  
Successfully created 1 ChargeRate
```

Example 6. Creating a couple of name-based multiplier charge rates

```
$ goldsh ChargeRate Create Type=NBM Name=QualityOfService  
Instance=Premium Rate=2  
Successfully created 1 ChargeRate
```

```
$ goldsh ChargeRate Create Type=NBM Name=QualityOfService  
Instance=BottomFeeder Rate=0.5  
Successfully created 1 ChargeRate
```

Example 7. Creating a value-based fee charge rate

```
$ goldsh ChargeRate Create Type=VBF Name=Shipping Rate=25  
Successfully created 1 ChargeRate
```

Example 8. Creating a name-based fee charge rate

```
$ goldsh ChargeRate Create Type=NBF Name=Zone Instance=Asia Rate=200  
Successfully created 1 ChargeRate
```

Example 9. Creating a couple of multi-dimensional value-based resource charge rates

```
$ goldsh ChargeRate Create Type=Disk Name=User Instance=dave Rate=0.02  
Successfully created 1 ChargeRate
```

```
$ goldsh ChargeRate Create Type=Disk Name=User Instance=michael  
Rate=0.05  
Successfully created 1 ChargeRate
```

Querying ChargeRates

To display charge rate information, use the command **goldsh ChargeRate Query**:

```
goldsh ChargeRate Query [Show:=<"Field1,Field2,...">] [Type==<Charge Rate Type>]
[Name==<Charge Rate Name>] [Instance==<Charge Rate Instance>] [Rate==<Floating Point
Multiplier>] [Description==<Description>] [ShowUsage:=True]
```

Example 8. Listing all charge rates

\$ goldsh ChargeRate Query

Type	Name	Instance	Rate	Description
VBR	Processors		1	
VBR	Memory		0.001	
NBR	License	Matlab	5	
VBV	Power		0.001	
VBV	CpuTime		1	
NBV	Feature	GPU	200	
VBM	Discount		1	
NBM	QualityOfService	Premium	2	
NBM	QualityOfService	BottomFeeder	0.5	
VBF	Shipping		25	
NBF	Zone	Asia	200	
Disk	User	dave	0.02	
Disk	User	michael	0.05	

Modifying Charge Rates

To modify a charge rate, use the command **goldsh ChargeRate Modify**:

goldsh ChargeRate Modify [Rate=<Floating Point Multiplier>] [Description=<Description>]
[Type==<Charge Rate Type>] [Name==<Charge Rate Name>] [Instance==<Charge Rate Instance>]
[Rate==<Floating Point Multiplier>] [ShowUsage:=True]



The goldsh control program allows you to make powerful and sweeping modifications to Gold objects. Misuse of this command could result in the inadvertent modification of all charge rates.

Example 9. Changing a charge rate

```
$ goldsh ChargeRate Modify Type==VBR Name==Memory Rate=0.05  
Successfully modified 1 ChargeRate
```

Deleting Charge Rates

To delete a charge rate, use the command **goldsh ChargeRate Delete**:

goldsh ChargeRate Delete [Type==<Charge Rate Type>] [Name==<Charge Rate Name>]
[Instance==<Charge Rate Instance>] [Rate==<Floating Point Multiplier>]



The goldsh control program allows you to make powerful and sweeping modifications to Gold objects. Misuse of this command could result in the inadvertent deletion of all charge rates.

Example 10. Deleting a charge rate

```
$ goldsh ChargeRate Delete Type==VBR Name==Memory  
Successfully deleted 1 ChargeRate
```

Managing Transactions

Gold logs all modifying transactions in a detailed transaction journal (queries are not recorded). Previous transactions can be queried but not modified or deleted.

Querying Transactions

To display transaction information, use the command **glstxn**:

```
glstxn [-O object] [-A action] [-n name_or_id] [-U actor] [-a account_id] [-i allocation_id] [-u user_name] [-p project_name] [-m machine_name] [-J job_id] [-s start_time] [-e end_time] [-T transaction_id] [-R request_id] [--show attribute_name[,attribute_name...]...] [--showHidden] [--raw] [--debug] [--? | --help] [--man] [--quiet]
```

Example 1. List all deposits made in 2004

```
$ glstxn -A Deposit -s 2004-01-01 -e 2005-01-01
```

Example 2. List everything done by amy since the beginning of 2004

```
$ glstxn -U amy -s 2004-01-01
```

Example 3. List all transactions affecting Job Id PBS.1234.0

```
$ glstxn -J PBS.1234.0
```

Example 4. List all transactions affecting charge rates

```
$ glstxn -O ChargeRate
```

Managing Roles

Gold uses instance-level role based access controls to determine what users can perform what functions. Named roles are created, privileges are associated with the roles, and users are assigned to these roles.

Querying Roles

To display the currently defined roles, use the command **goldsh Role Query**:

goldsh Role Query [Show:=<"*Field1,Field2,...*">] [Name==<*Role Name*>]
[Description==<*Description*>] [ShowUsage:=True]

Example 1. Listing all roles

```
$ goldsh Role Query

Name                Description
-----
SystemAdmin         Can update or view any object
Anonymous            Things that can be done by anybody
OVERRIDE            A custom authorization method will be invoked
ProjectAdmin         Can update or view a project they are admin for
UserServices         User Services
Scheduler            Scheduler relevant Transactions
```

Querying Role Users

To list what users can perform what roles, use the command **goldsh RoleUser Query**:

goldsh RoleUser Query [Show:=<"Field1,Field2,...">] [Role==<Role Name>] [Name==<User Name>]
[ShowUsage:=True]

Example 2. Listing all role users

```
$ goldsh RoleUser Query
```

Role	Name
-----	-----
SystemAdmin	gold
Anonymous	ANY
OVERRIDE	ANY
Scheduler	maui
SystemAdmin	root
UserServices	amy

Querying Role Actions

To list what actions can be performed by what roles, use the command **goldsh RoleAction Query**:

```
goldsh RoleAction Query [Show:=<"Field1,Field2,...">] [Role==<Role Name>] [Object==<Object Name>] [Name==<Action Name>] [Instance==<Instance Name>] [ShowUsage:=True]
```

Example 3. Listing all role actions

\$ goldsh RoleAction Query

Role	Object	Name	Instance
-----	-----	-----	-----
Anonymous	ANY	Query	ANY
Anonymous	Account	Balance	ANY
Anonymous	Password	ANY	SELF
OVERRIDE	Account	Balance	ANY
ProjectAdmin	Project	ANY	ADMIN
Scheduler	Job	Charge	ANY
Scheduler	Job	Quote	ANY
Scheduler	Job	Reserve	ANY
SystemAdmin	ANY	ANY	ANY
UserServices	Job	Refund	ANY
UserServices	Machine	ANY	ANY
UserServices	Project	ANY	ANY
UserServices	ProjectMachine	ANY	ANY
UserServices	ProjectUser	ANY	ANY
UserServices	User	ANY	ANY

Creating Roles

To create a new role, use the command **goldsh Role Create**:

goldsh Role Create Name=<*Role Name*> [Description=<*Description*>] [ShowUsage:=True]

Example 4. Creating a Manager role

```
$ goldsh Role Create Name=Manager Description="Manages Roles and Responsibilities"
```

Name	Description
------	-------------

Manager	Manages Roles and Responsibilities
---------	------------------------------------

Successfully created 1 Role

Associating an Action with a Role

To add an action to a role, use the command **goldsh RoleAction Create**:

goldsh RoleAction Create Role=<Role Name> Object=<Object Name> Name=<Action Name>
[Instance=<Instance Name>] [ShowUsage:=True]

The Instance indicates which specific instances of the object the action(s) can be performed on. Instances are interpreted as the value of the solitary primary key for an object. Unless otherwise specified, the instance will default to a value of ANY.

Valid values for Instance include:

ANY Any or all of the object instances

NONE No object instances

SELF Only objects identified with myself (like my own username)

ADMIN Only object instances that I am an admin for

<specific> A specific named instance

For example, the Role Action:

Role	Object	Name	Instance
-----	-----	-----	-----
ChemistryAdmin	Project	Modify	Chemistry

allows users having the ChemistryAdmin role to modify the Chemistry Project.

Example 5. Allow the Manager to change role responsibilities

```
$ goldsh RoleAction Create Role=Manager Object=RoleAction Name=ANY
```

Role	Object	Name	Instance
-----	-----	-----	-----
Manager	RoleAction	ANY	ANY

```
Successfully created 1 RoleAction
```

Adding a Role to a User

To associate a user with a role, use the command **goldsh RoleUser Create**:

goldsh RoleUser Create Role=<Role Name> Name=<User Name> [ShowUsage:=True]

Example 6. Adding a user to the Manager role

```
$ goldsh RoleUser Create Role=Manager Name=dave  
  
Role      Name  
-----  
Manager   dave  
  
Successfully created 1 RoleUser
```



A user must first be defined to Gold before they can be added to a role (see [Creating Users](#)).

Removing an Action from a Role

To disassociate an action from a role, use the command **goldsh RoleAction Delete**:

goldsh RoleAction Delete [Role==<Role Name>] [Object==<Object Name>] [Name==<Action Name>]
[Instance==<Instance Name>] [ShowUsage:=True]



The goldsh control program allows you to make powerful and sweeping modifications to Gold objects. Misuse of this command could result in the inadvertent deletion of all role actions.

Example 7. Don't let UserServices Create or Update Projects

```
$ goldsh RoleAction Delete Role==UserServices Object==Project
Name==ANY

Role          Object      Name      Instance
-----
UserServices  Project    ANY       ANY

Successfully deleted 1 RoleActions
```

Removing a Role from a User

To disassociate a user and a role, use the command **goldsh RoleUser Delete**:

goldsh RoleUser Delete [Role==<Role Name>] [Name==<User Name>] [ShowUsage:=True]



The goldsh control program allows you to make powerful and sweeping modifications to Gold objects. Misuse of this command could result in the inadvertent deletion of all role users.

Example 8. Removing dave as a Manager

```
$ goldsh RoleUser Delete Role==Manager Name==dave

Role      Name
-----
Manager   dave

Successfully deleted 1 RoleUser
```

Deleting Roles

To delete a role, use the command **goldsh Role Delete**:

goldsh Role Delete [Name==<Role Name>] [Description==<Description>] [ShowUsage:=True]



The goldsh control program allows you to make powerful and sweeping modifications to Gold objects. Misuse of this command could result in the inadvertent modification of all roles.

Example 9. Deleting the Manager role

```
$ goldsh Role Delete Name==Manager

Name      Description
-----
Manager    Manages Roles and Responsibilities

Successfully deleted 1 Roles and 3 associations
```

Managing Passwords

Passwords must be established for each user who wishes to use the web-based GUI. Passwords must be at least eight characters and are stored in encrypted form. Valid operations on passwords include creating, modifying, and deleting passwords.

Creating Passwords

To create a new password, use the command **goldsh Password Create**:

goldsh Password Create User=<User Name> Password=<Encrypted Password> [ShowUsage:=True]

Example 1. Creating a password

```
$ goldsh Password Create User=amy Password=mysecret
User      Password
-----
amy       Nn0NaSpwELQ+FKa36og9l6EczO+kUEoN
Successfully created 1 Password
```


Querying Passwords

To display password information, use the command **goldsh Password Query**:

goldsh Password Query [Show:=<"*Field1,Field2,...*">] [User==<*User Name*>] [ShowUsage:=True]

Example 2. List the users who have set passwords

```
$ goldsh Password Query Show:=User
```

```
User
```

```
-----
```

```
amy
```

```
gold
```

Modifying Passwords

To change a password, use the command **goldsh Password Modify**:

goldsh Password Modify [Password=<Encrypted Password>] [Name==<User Name>] [ShowUsage:=True]



The goldsh control program allows you to make powerful and sweeping modifications to Gold objects. Misuse of this command could result in the inadvertent modification of all passwords.

Example 3. Changing amy's password

```
$ goldsh Password Modify User==amy Password=changeme

User      Password
-----
amy       HZYzwD20o1XIE/gxRYyFKP2sumkCluHm

Successfully modified 1 Passwords
```

Deleting Passwords

To delete a password, use the command **goldsh Password Delete**:

goldsh Password Delete [Name==<User Name>]



The goldsh control program allows you to make powerful and sweeping modifications to Gold objects. Misuse of this command could result in the inadvertent deletion of all passwords.

Example 4. Deleting a password

```
$ goldsh Password Delete User==amy

User      Password
-----
amy       HZYzwD20o1XIE/gxRYyFKP2sumkCluHm

Successfully deleted 1 Passwords
```

Using the Gold Shell (goldsh)

goldsh is an interactive control program that can access all of the advanced functionality in Gold.



The goldsh control program allows you to make powerful and sweeping modifications to many objects with a single command. Inadvertant mistakes could result in modifications that are very difficult to reverse.

Usage

Gold commands can be invoked directly from the command line as arguments, or read from stdin (interactively or redirected from a file).

goldsh [**—debug**] [**-?** | **—help**] [**--man**] [**—raw**] [**—quiet**] [**-v** | **—verbose**] [**<Command>**]

Example 1. Specifying the command as direct arguments

```
$ goldsh System Query
```

Name	Version	Organization	Description
Gold	2.0.b1.0		BetaRelease

Example 2. Using the interactive prompt

```
$ goldsh
```

```
gold> System Query
```

Name	Version	Organization	Description
Gold	2.0.b1.0		BetaRelease

```
gold> quit
```

Example 3. Reading commands from a file

```
$ cat >commands.gold <<EOF
```

```
System Query
```

```
quit
```

```
EOF
```

```
$ goldsh <commands.gold
```

Name	Version	Organization	Description
Gold	2.0.b1.0		BetaRelease

Command Syntax

Gold commands are of the form:

```
<Object> [, <Object> ...] <Action> [ [<Conjunction>] [<Open_Parenthesis> ...] [<Object> .] <Name>  
<Operator> [<Object> .] <Value> [<Close_Parenthesis> ...] ...]
```

The basic form of a command is <Object> <Action> [<Name><Operator><Value>]*. When an action is performed on more than one object, such as in a multi-object query, the objects are specified in a comma-separated list. Commands may accept zero or more predicates which may function as fields to return, conditions, update values, processing options, etc. Predicates, in their simplest form, are expressed as Name, Operator, Value tuples. Predicates may be combined via conjunctions with grouping specified with parentheses. When performing multi-object queries, names and values may need to be associated with their respective objects.

Valid conjunctions include:

&&

and

||

or

&!

and not

!|

or not

Open parentheses may be any number of literal open parentheses '('.

Name is the name of the condition, assignment, or option. When performing a multi-object query, a name may need to be prepended by its associated object separated by a period.

Valid operators include:

==

equals

<

less than

>

greater than

<=

less than or equal to

>=

greater than or equal to

!=

not equal to

~

matches

=

is assigned

+=

is incremented by

-=

is decremented by

:=

option

:!

not option

Value is the value of the selection list, condition, assignment, or option. When performing a multi-object query, a value may need to be prepended by its associated object (called the subject) separated by a period.

Close parentheses may be any number of literal closing parentheses ')'.

Valid Objects

To list the objects available for use in Gold commands, issue the Gold command: Object Query

Example 4. Listing all objects

```
gold> Object Query Show:="Sort (Name) "
```

```
Name
```

```
-----
```

```
ANY
```

```
Account
```

```
AccountAccount
```

```
AccountMachine
```

```
AccountOrganization
```

```
AccountProject
```

```
AccountUser
```

```
Action
```

```
Allocation
```

```
Attribute
```

```
ChargeRate
```

```
Job
```

```
Machine
```

```
NONE
```

```
Object
```

```
Organization
```

```
Password
```

```
Project
```

```
ProjectMachine
```

```
ProjectUser
```

```
Quotation
```

```
QuotationChargeRate
```

```
Reservation
```

```
Role
```

```
RoleAction
```

Valid Actions for an Object

To list the actions that can be performed on an object, use the Gold command: Action Query

Example 5. Listing all actions associated with the Account object

```
gold> Action Query Object==Account Show:="Sort(Name) "  
Name  
-----  
Balance  
Create  
Delete  
Deposit  
Modify  
Query  
Transfer  
Undelete  
Withdraw
```


Valid Predicates for an Object and Action

By appending the option "ShowUsage:=True" to a command, the syntax of the command is returned, expressed in SSSRMAP XML Message Format.

Example 6. Show the usage for Allocation Query

```
gold> Allocation Query ShowUsage:=True

<Request action="Query">
  <Object>Allocation</Object>
  [<Get name="Id" [op="Sort|Tros|Count|GroupBy|Max|Min"]></Get>]
  [<Get name="Account"
[op="Sort|Tros|Count|GroupBy|Max|Min"]></Get>]
  [<Get name="StartTime"
[op="Sort|Tros|Count|GroupBy|Max|Min"]></Get>]
  [<Get name="EndTime"
[op="Sort|Tros|Count|GroupBy|Max|Min"]></Get>]
  [<Get name="Amount"
[op="Sort|Tros|Count|GroupBy|Max|Min|Sum|Average"]></Get>]
  [<Get name="Deposited"
[op="Sort|Tros|Count|GroupBy|Max|Min|Sum|Average"]></Get>]
  [<Get name="Active" [op="Sort|Tros|Count|GroupBy"]></Get>]
  [<Get name="Description"
[op="Sort|Tros|Count|GroupBy|Max|Min"]></Get>]
  [<Where name="Id" [op="EQ|NE|GT|GE|LT|LE (EQ)"] [conj="And|Or
(And)"] [group="<Integer Number>Integer Number}</Where>]
  [<Where name="Account" [op="EQ|NE|GT|GE|LT|LE|Match (EQ)"]
[conj="And|Or (And)"] [group="<Integer Number>Account Name}</Where>]
  [<Where name="StartTime" [op="EQ|NE|GT|GE|LT|LE (EQ)"]
[conj="And|Or (And)"] [group="<Integer Number>YYYY-MM-DD
[hh:mm:ss]|-infinity|infinity|now</Where>]
  [<Where name="EndTime" [op="EQ|NE|GT|GE|LT|LE (EQ)"]
[conj="And|Or (And)"] [group="<Integer Number>YYYY-MM-DD
[hh:mm:ss]|-infinity|infinity|now</Where>]
  [<Where name="Amount" [op="EQ|NE|GT|GE|LT|LE (EQ)"]
```

Common Options

There are a number of options that may be specified for all commands. These options include: ShowUsage
ShowUsage

This option may be included with any command to cause the command to return a usage message in SSSRMAP XML Message Format.

Common Actions Available for most Objects

There are a number of actions that are available for most objects. These actions include Query, Create, Modify, Delete, and Undelete. Commands involving these actions inherit some common structure unique to the action type.

Query Action

The Query action is used to query objects. It accept predicates that describe the attributes (fields) to return (including aggregation operations on those attributes), conditions that select which objects to return the attributes for, and other options unique to queries.

Selections

Selections use the Show option to specify a list of the attributes to return for the selected object. If selections are not specified, a default set of attributes (those not marked as hidden) will be returned.

Name = Show

Op = :=

Value = "attribute1,attribute2,attribute3,..."

Aggregation operators may be applied to attributes by enclosing the target attribute in parenthesis and prepending the name of the desired operator. The aggregation operators that can be applied depend on the datatype of the attribute.

Valid selection operators include:

Sort Ascending sort

Tros Descending sort

Count Count

Max Maximum value

Min Minimum value

Average Average value

Sum Sum

GroupBy Group other aggregations by this attribute

For example: Allocation Query Show:="Sum(Amount),GroupBy(Account)"

Conditions

Conditions are used to select which objects the action is to be performed on.

Name = Name of the attribute to be tested

Op = conditional operator

Value = The object or value against which the attribute is tested

Valid condition operators include:

== Equal to

!= Not equal to

< Less than

> Greater than

<= Less than or equal to

>= Greater than or equal to

~ Matches

Matching uses the wildcards * and ? (equivalent to SQL % and _ respectively) in a manner similar to file globbing. * matches zero or more unspecified characters and ? matches exactly one unspecified character. For example mscf* matches objects having the specified attributes whose values start with the letters mscf, while mscf? matches objects having the specified attributes whose values start with mscf and have a total of exactly five characters.

Options

Options indicate processing options that affect the result.

Name = Name of the option

Op = :=

Value = Value of the option

Valid options for query actions include:

ShowHidden:=True|False (False) Includes hidden attributes in the result

Time:="YYYY-MM-DD [hh:mm:ss]" Run the command as if it were the specified time

Unique:=True|False (False) Display only unique results (like DISTINCT in SQL)

Limit:={Integer Number} Limit the results to the number of objects specified

Example 7. Return the number of inactive reservations

```
gold> Reservation Query EndTime<now Show:="Count(Id) "  
Id  
---  
8
```

Create Action

The Create action is used to create a new object. It accepts predicates that describe the values of the attributes to be set.

Assignments

Assignments specify values to be assigned to attributes in the new object.

Name = Name of the attribute being assigned a value

Op = = (is assigned)

Value = The new value being assigned to the attribute

Example 8. Add a new project member

```
gold> ProjectUser Create Project=chemistry Name=scottmo  


| Project   | Name    | Active | Admin |
|-----------|---------|--------|-------|
| chemistry | scottmo | True   | False |

  
Successfully created 1 ProjectUser
```

Modify Action

The Modify action is used to modify existing objects. It accepts predicates that select which objects will be modified and predicates that describe the values of the attributes to be set.

Assignments

Assignments specify values to be assigned to attributes in the selected objects.

Name = Name of the attribute being assigned a value

Op = assignment operators {=, +=, -=}

Value = The value being assigned to the attribute

Valid assignment operators include:

= is assigned

`+=` is incremented by
`-=` is decremented by

Conditions

Conditions are used to select which objects the action is to be performed on.

Name = Name of the attribute to be tested

Op = conditional operator

Value = The object or value against which the attribute is tested

Valid condition operators include:

`==` Equal to

`!=` Not equal to

`<` Less than

`>` Greater than

`<=` Less than or equal to

`>=` Greater than or equal to

`~` Matches

Matching uses the wildcards `*` and `?` (equivalent to SQL `%` and `_` respectively) in a manner similar to file globbing. `*` matches zero or more unspecified characters and `?` matches exactly one unspecified character. For example `m*scf*` matches objects having the specified attributes whose values start with the letters `m*scf`, while `m*scf?` matches objects having the specified attributes whose values start with `m*scf` and have a total of exactly five characters.

Example 9. Change/set scottmo phone number and email address

```
gold> User Modify Name==scottmo PhoneNumber="(509) 376-2204"
EmailAddress="Scott.Jackson@pnl.gov"

Name      Active   CommonName      PhoneNumber
EmailAddress      DefaultProject   Description
-----
scottmo    True      Jackson, Scott M.      (509) 376-2204
Scott.Jackson@pnl.gov

Successfully modified 1 Users
```

Example 10. Extend all reservations against project chemistry by 10 days

```
gold> Reservation Modify EndTime+="10 days" Project==chemistry

Id  Account   Amount   Name      Job  User   Project
Machine  EndTime
---
1    2         57600    PBS.1234.0  1    amy    chemistry
colony    2004-11-06 10:47:30

Successfully modified 1 Reservations
```

Delete Action

The Delete action is used to delete objects. It accepts predicates that select which objects are to be deleted.

Conditions

Conditions are used to select which objects the action is to be performed on.

Name = Name of the attribute to be tested
Op = conditional operator
Value = The object or value against which the attribute is tested

Valid condition operators include:

== Equal to
!= Not equal to
< Less than
> Greater than
<= Less than or equal to
>= Greater than or equal to
~ Matches

Matching uses the wildcards * and ? (equivalent to SQL % and _ respectively) in a manner similar to file globbing. * matches zero or more unspecified characters and ? matches exactly one unspecified character. For example mscf* matches objects having the specified attributes whose values start with the letters mscf, while mscf? matches objects having the specified attributes whose values start with mscf and have a total of exactly five characters.

Example 11. Get rid of the pesky Jacksons

```
gold> User Delete CommonName~"Jackson*"

Name      Active   CommonName      PhoneNumber
EmailAddress      DefaultProject   Description
-----
-----
scottmo    True      Jackson, Scott M.      (509) 376-2204
Scott.Jackson@pnl.gov

Successfully deleted 1 Users and 1 associations
```

Undelete Action

The Delete action is used to restore deleted objects. It accepts predicates that select which objects are to be undeleted.

Conditions

Conditions are used to select which objects the action is to be performed on.

Name = Name of the attribute to be tested
Op = conditional operator
Value = The object or value against which the attribute is tested

Valid condition operators include:

== Equal to
!= Not equal to
< Less than
> Greater than
<= Less than or equal to
>= Greater than or equal to
~ Matches

Matching uses the wildcards * and ? (equivalent to SQL % and _ respectively) in a manner similar to file globbing. * matches zero or more unspecified characters and ? matches exactly one unspecified character. For example mscf* matches objects having the specified attributes whose values start with the letters mscf, while mscf? matches objects having the specified attributes whose values start with mscf and have a total of exactly five characters.

Example 12. Resurrect the deleted users that were active

```
gold> User Undelete Active==True
```

Name	Active	CommonName	PhoneNumber
EmailAddress		DefaultProject	Description
-----	-----	-----	-----
-----	-----	-----	-----
scottmo	True	Jackson, Scott M.	(509) 376-2204
Scott.Jackson@pnl.gov			

```
Successfully undeleted 1 Users and 1 associations
```

Multi-Object Queries

Gold supports multi-object queries (table joins). Multiple objects are specified via a comma-separated list and attributes need to be prefixed by the associated object.

Example 13. Print the current and total allocation summed by project

```
gold> Allocation,AccountProject Query
Show:="GroupBy(AccountProject.Name),Sum(Allocation.Amount),Sum(Allocati
Allocation.Account==AccountProject.Account Allocation.Active==True

Name                Amount                Deposited
-----
biology             193651124             360000000
chemistry           296167659             360000000
```

Example 14. Show all active projects for amy or bob

```
gold> Project,ProjectUser Query Show:="Project.Name" (
ProjectUser.Name==bob || ProjectUser.Name==amy ) &&
Project.Name==ProjectUser.Project && Project.Active==True Unique:=True

Name
-----
biology
chemistry
```


Customizing Gold Objects

Gold provides the ability to dynamically create new objects or customize or delete existing objects through the Gold control program (goldsh).



The object customizations described in this chapter will be noticeable in subsequent goldsh queries (and in the web GUI after a fresh login). For installations with a database that supports multiple connections (e.g. PostgreSQL) these changes will be visible immediately while others (e.g. SQLite) will require the Gold server to be restarted. Client commands may need to be modified to properly interact with changed objects or attributes.



The goldsh control program allows you to make powerful and sweeping modifications to many objects with a single command. Inadvertent mistakes could result in modifications that are very difficult to reverse.

Removing an Attribute from an Object

To delete an attribute from an object, use the command **goldsh Attribute Delete**:

goldsh Attribute Delete Object==<Object Name> Name==<Attribute Name> [ShowUsage:=True]



The goldsh control program allows you to make powerful and sweeping modifications to Gold objects. Misuse of this command could result in the inadvertent deletion of all attributes.



When using Gold as an Allocation Manager, certain objects and attributes are assumed to exist. For example, a call to Job Charge would fail if you had deleted the Allocation Amount attribute. The Attribute Undelete command might come in useful in such a case.

Example 1. Removing the Organization attribute from Machine

```
$ goldsh Attribute Delete Object==Machine Name==Organization
Successfully deleted 1 Attribute
```

Example 2. Perhaps you don't care to track the Executable attribute in a Job

```
$ goldsh Attribute Delete Object==Job Name==Executable
Successfully deleted 1 Attribute
```

Adding an Attribute to an Object

To create a new attribute for an object, use the command **goldsh Attribute Create**:

```
goldsh Attribute Create Object=<Object Name> Name=<Attribute Name>  
[DataType=AutoGen|TimeStamp|Boolean|Float|Integer|Currency|(String)] [PrimaryKey=True|(False)]  
[Required=True|(False)] [Fixed=True|(False)] [Values=<Foreign Key or List of Values>]  
[DefaultValue=<Default Value>] [Sequence=<Integer Number>] [Hidden=<True|(False)>]  
[Description=<Description>] [ShowUsage:=True]
```

Example 3. Adding a Country Attribute to User

```
$ goldsh Attribute Create Object=User Name=Country  
Values="(Brazil,China,France,Russia,USA)" DefaultValue=USA  
  
Successfully created 1 Attribute
```

Example 4. Track submission time in jobs

```
$ goldsh Attribute Create Object=Job Name=SubmissionTime  
DataType=TimeStamp  
  
Successfully created 1 Attribute
```

Modifying an Attribute

To modify an attribute, use the command **goldsh Attribute Modify**:

goldsh Attribute Modify Object==<Object Name> Name==<Attribute Name> [Required=True|(False)]
[Fixed=True|(False)] [Values=<Foreign Key or List of Values>] [DefaultValue=<Default Value>]
[Sequence=<Integer Number>] [Hidden=<True/(False)>] [Description=<Description>]
[ShowUsage:=True]



The goldsh control program allows you to make powerful and sweeping modifications to Gold objects. Misuse of this command could result in the inadvertent modification of all attributes.

Example 5. Change User Organization values to not be restricted by foreign key

```
$ goldsh Attribute Modify Object==User Name==Organization Values=NULL  
Successfully modified 1 Attribute
```

Creating a Custom Object

To create a new object, use the command **goldsh Object Create**:

goldsh Object Create Name=<Object Name> [Association=True|False] [Parent=<Parent Object>]
[Child=<Child Object>] [Description=<Description>] [ShowUsage:=True]

Example 6. Creating a Node Object

```
$ goldsh Object Create Name=Node Description="Node Information"  
Successfully created 1 Object
```

Example 7. Track submission time in jobs

```
$ goldsh Attribute Create Object=Job Name=SubmissionTime  
DataType=TimeStamp  
Successfully created 1 Attribute
```

Adding an Action to an Object

To specify that an action is allowed for an object, use the command **goldsh Action Create**:

goldsh Action Create Object=<Object Name> Name=<Action Name> [Display=True|(False)]
[Description=<Description>] [ShowUsage:=True]

Example 8. Adding a Modify Action to Transaction

```
$ goldsh Action Create Object=Transaction Name=Modify  
Description=Modify  
  
Successfully created 1 Action
```

Examples Creating Custom Objects

Creating a custom object involves defining a new object, adding attributes to the object, and adding actions to the object.

Example 9. Creating a License object to track license usage and charges.

Invoke the Gold control program in interactive mode.

```
$ goldsh
```

Create the License Object.

```
gold> Object Create Name=License Description=License  
Successfully created 1 Object
```

Next, define its attributes. Give each record a unique id (so the record can be more easily modified), a license type that can be one of (Matlab,Mathematica,Compiler,AutoCAD,Oracle), the user who is using it, the start and end time, how many instances of the license were used, and how much was charged.

```
gold> Attribute Create Object=License Name=Id DataType=AutoGen  
PrimaryKey=True Description="Record Id"  
Successfully created 1 Attribute  
gold> Attribute Create Object=License Name=Type DataType=String  
Required=True Values="(Matlab,Mathematica,Compiler,AutoCAD,Oracle)"  
Fixed=True Description="License Type"  
Successfully created 1 Attribute  
gold> Attribute Create Object=License Name=User Required=True  
Values="@User" Description="User Name"  
Successfully created 1 Attribute  
gold> Attribute Create Object=License Name=StartTime  
DataType=TimeStamp Description="Start Time"  
Successfully created 1 Attribute  
gold> Attribute Create Object=License Name=EndTime  
DataType=TimeStamp Description="End Time"  
Successfully created 1 Attribute  
gold> Attribute Create Object=License Name=Count DataType=Integer  
Description="Number of Licenses Used"
```

Finally, designate the actions you want to allow on the object. The standard set of actions includes Create, Query, Delete, Modify, and Undelete. If you would like to manage licenses from the web GUI, set Display=True.

```
gold> Action Create Object=License Name=Create Display=True  
Description=Create  
Successfully created 1 Action
```

```
gold> Action Create Object=License Name=Query Display=True
Description=Query

Successfully created 1 Action

gold> Action Create Object=License Name=Modify Display=True
Description=Modify

Successfully created 1 Action

gold> Action Create Object=License Name=Delete Display=True
Description=Delete

Successfully created 1 Action

gold> Action Create Object=License Name=Undelete Display=True
Description=Undelete

Successfully created 1 Action
```

When done, exit the goldsh prompt.

```
gold> quit
```

That's about it. Licenses should now be able to be managed via the GUI and goldsh. The data source will need to use one of the methods of interacting with Gold (see [Methods of interacting with Gold](#)) in order to push license record usage info to Gold.

Apart from being used as an Allocation Manager, Gold can be used as a generalized information service. It can be used to manage just about any object oriented information over the web. For example, Gold could be used to provide meta-schedulers with machine/user mappings, or node/resource information.

Example 10. Using Gold as a Grid Map File.

Invoke the Gold control program in interactive mode.

```
$ goldsh
```

Create the GridMap Object.

```
gold> Object Create Name=GridMap Description="Online Grid Map File"

Successfully created 1 Object
```

Next, define its attributes. Each entry will consist of a userid (which will serve as the primary key) and a required public X.509 certificate.

```
gold> Attribute Create Object=GridMap Name=User PrimaryKey=True
Values=@User Description="User Name"

Successfully created 1 Attribute

gold> Attribute Create Object=GridMap Name=Certificate DataType=String
Required=True Description="X.509 Public Key"

Successfully created 1 Attribute
```

Finally, designate the actions you want to allow on the object. If you would like to manage certificates from the web GUI, set Display=True.

```
gold> Action Create Object=GridMap Name=Create Display=True
Description=Create

Successfully created 1 Action

gold> Action Create Object=GridMap Name=Query Display=True
Description=Query

Successfully created 1 Action

gold> Action Create Object=GridMap Name=Modify Display=True
Description=Modify

Successfully created 1 Action

gold> Action Create Object=GridMap Name=Delete Display=True
Description=Delete

Successfully created 1 Action

gold> Action Create Object=GridMap Name=Undelete Display=True
Description=Undelete

Successfully created 1 Action
```

Exit the goldsh prompt.

```
gold> quit
```

From this point, a peer service will need to use one of the methods of interacting with Gold (see [Methods of interacting with Gold](#)) in order to query the GridMap information.

Integration with the Resource Management System

Dynamic versus Delayed Accounting

Delayed Accounting

In the absence of a dynamic system, some sites enforce allocations by periodically (weekly or nightly) parsing resource manager job logs and then applying debits against the appropriate project accounts. Although Gold can easily support this type of system by the use of the qcharge command in post-processing scripts, this approach allows users or projects to use resources significantly beyond their designated allocation and generally suffers from stale accounting information.

Dynamic Accounting

Gold's design allows it to interact dynamically with your resource management system. Charges for resource utilization can be made immediately when the job finishes (or even incrementally throughout the job). Additionally, reservations can be issued at the start of a job to place a hold against the user's account, thereby ensuring that a job will only start if it has sufficient reserves to complete. The remainder of this document will describe the interactions for dynamic accounting.

Interaction Points

Job Quotation @ Job Submission Time [Optional — Recommended]

When a job is submitted to a grid scheduler or resource broker, it may be useful to determine how much it will cost to run on a particular resource by requesting a job quote. If the quote succeeds, it will return a quote id along with the quoted amount for the job. This quote id may be used later to guarantee that the same charge rates used to form the quote will also be used in the final job charge calculation.

Even when a job is exclusively scheduled locally, it is useful to obtain a quote at the time of submission to the local resource manager to ensure the user has sufficient funds to run the job and that it meets the access policies necessary for the charge to succeed. A warning can be issued if funds are low or the job might be rejected with an informative message in the case of insufficient funds or any other problems with the account. Without this interaction, the job might wait in the queue for days only to fail when it tries to start.

To make a job quotation with Gold at this phase requires that:

- the grid scheduler has built-in Gold allocation manager support {Silver}, or
- the resource manager supports a submit filter {LoadLeveler(SUBMIT_FILTER), LSF(esub)}, or
- a wrapper could be created for the submit command {PBS(qsub)}.

Job Reservation @ Job Start Time [Optional — Highly Recommended]

Just before a job starts, a hold (reservation) is made against the appropriate account(s), temporarily reducing the user's available balance by an amount based on the resources requested and the estimated wallclock limit. If this step is omitted, it would be possible for users to start more jobs than they have funds to support.

If the reservation succeeds, it will return a message indicating the amount reserved for the job. In the case where there are insufficient resources to run the job or some other problem with the reservation, the command will fail with an informative message. Depending on site policy, this may or may not prevent the job from starting.

To make a job reservation with Gold at this phase requires that:

- the scheduler or resource manager has built-in Gold allocation manager support {Maui(AMCFG)}, or
- the resource manager is able to run a script at job start time {LoadLeveler(prolog), PBS(prologue), LSF(pre_exec)}.

Job Charge @ Job End Time [Required]

When a job ends, a charge is made to the user's account(s). Any associated reservations are automatically removed as a side-effect. Depending on site policy, a charge can be elicited only in the case of a successful completion, or for all or specific failure cases as well. Ideally, this step will occur immediately after the job completes (dynamic accounting). This has the added benefit that job run times can often be reconstructed from Gold job reservation and charge timestamps in case the resource management job accounting data becomes corrupt.

If the charge succeeds, it will return a message indicating the amount charged for the job.

To make a job charge with Gold at this phase requires that:

- the scheduler or resource manager has built-in Gold allocation manager support {Maui(AMCFG)}, or
- the resource manager is able to run a script at job start time {LoadLeveler(epilog), PBS(epilogue),

LSF(post_exec)}, or

- the resource management system supports some kind of feedback or notification mechanism occurring at the end of a job (an email can be parsed by a mail filter).

Methods of interacting with Gold

There are essentially six ways of programatically interacting with Gold.

Configuring an application that already has hooks for Gold

The easiest way to use Gold is to use a resource management system with built-in support for Gold. For example, the Maui Scheduler and Silver Grid Scheduler can be configured to directly interact with Gold to perform the quotes, reservations, and charges by setting the appropriate parameters in their config files.

Example 1. Configuring Maui to use Gold

Add an appropriate AMCFG line into maui.cfg to tell Maui how to talk to Gold

```
$ vi /usr/local/maui/maui.cfg

AMCFG[bank] TYPE=GOLD HOST=control_node1 PORT=7112 SOCKETPROTOCOL=HTTP
WIREFPROTOCOL=XML CHARGEPOLICY=DEBITALLWC JOBFAILUREACTION=IGNORE
TIMEOUT=15
```

Add a CLIENTCFG line into maui-private.cfg to specify the shared secret key. This secret key will be the same secret key specified in the "make auth_key" step.

```
$ vi /usr/local/maui/maui-private.cfg

CLIENTCFG[AM:bank] KEY=sss AUTHTYPE=HMAC64
```

Gold will need to allow the the user id that maui runs under to perform scheduler related commands (Job Charge, Reserve, Quote, etc).

```
$ gmkuser -d "Maui Scheduler" maui

Successfully created 1 User

$ goldsh RoleUser Create Role=Scheduler Name=maui

Role          Name
-----
Scheduler maui

Successfully created 1 RoleUser
```

Using the appropriate command-line client

From inside a script, or by invoking a system command, you can use a command line client (one of the "g" commands in gold's bin directory).

Example 2. To issue a charge at the completion of a job, you would use gcharge:

```
gcharge -J PBS.1234.0 -p chemistry -u amy -m colony -P 2 -t 1234
```

Using the Gold control program

The Gold control program, goldsh, will issue a charge for a job expressed in xml (SSS Job Object).

Example 3. To issue a charge you must invoke the Charge action on the Job object:

```
goldsh
Data:="<Job><JobId>PBS.1234.0</JobId><ProjectId>chemistry</ProjectId>
<UserId>amy</UserId><MachineName>colony</MachineName>
<Processors>2</Processors><WallDuration>1234</WallDuration>"
```

Use the Perl API

If your resource management system is written in Perl or if it can invoke a Perl script, you can access the full Gold functionality via the Perl API.

Example 4. To make a charge via this interface you might do something like:

```
use Gold;

my $request = new Gold::Request(object => "Job", action => "Charge");
my $job = new Gold::Datum("Job");
$job->setValue("JobId", "PBS.1234.0");
$job->setValue("ProjectId", "chemistry");
$job->setValue("UserId", "amy");
$job->setValue("MachineName", "colony");
$job->setValue("Processors", "2");
$job->setValue("WallDuration", "1234");
$request->setDatum($job);
my $response = $request->getResponse();
print $response->getStatus(), ": ", $response->getMessage(), "\n";
```

Communicating via the SSSRMAP Protocol

Finally, it is possible to interact with Gold by directly using the SSSRMAP Wire Protocol and Message Format over the network (see [SSS Resource Management and Accounting Documentation](#)). This will entail building the request body in XML, appending an XML digital signature, combining these in an XML envelope framed in an HTTP POST, sending it to the server, and parsing the similarly formed response. The Maui Scheduler communicates with Gold via this method.

Example 5. The message might look something like:

```
POST /SSSRMAP HTTP/1.1
Content-Type: text/xml; charset="utf-8"
Transfer-Encoding: chunked

190
<?xml version="1.0" encoding="UTF-8"?>

<Envelope>
  <Body actor="scottmo" chunking="True">
    <Request action="Charge" object="Job">
      <Data>
        <Job>
          <JobId>PBS.1234.0</JobId>
          <ProjectId>chemistry</ProjectId>
          <UserId>amyh</UserId>
          <MachineName>colony</MachineName>
          <Processors>2</Processors>
          <WallDuration>1234</WallDuration>
        </Job>
      </Data>
    </Request>
  </Body>
  <Signature>
    <DigestValue>azu4obZswzBt89OgATukBeLyt6Y=</DigestValue>
```

```
<SignatureValue>YXE/C08XX3RX4PMU1bWju+5/E5M=</SignatureValue>  
  <SecurityToken type="Symmetric"></SecurityToken>  
</Signature>  
</Envelope>  
0
```

Configuration Files

Gold uses two configuration files: one for the server (goldd.conf) and one for the clients (gold.conf). For configuration parameters that have hard-coded defaults, the default value is specified within brackets.

Server Configuration

The following configuration parameters may be set in the server configuration file (goldd.conf).

- **account.autogen** [true] — If set to true, when a new project is created Gold will automatically create an associated default account. Additionally, if you try to make a deposit and no accounts match the specifications, an account will be created using the specified criteria and a deposit will be made into that account.
- **allocation.autogen** [true] — If set to true, when a new account is created Gold will automatically create an associated default allocation with zero credits.
- **database.datasource** [DBI:Pg:dbname=gold;host=localhost] — The Perl DBI data source name for the database you wish to connect to.
- **database.password** — The password to be used for the database connection (if any).
- **database.user** — The username to be used for the database connection (if any).
- **response.chunksize** [0] — Indicates the line length in the data response that will trigger message segmentation (or truncation). A value of 0 (zero) means unlimited, i.e. that the server will not truncate or segment large responses unless overridden by a chunksize specification in a client request. The response chunksize will be taken to be the smaller of the client and server chunksize settings.
- **currency.precision** [0] — Indicates the number of decimal places in the resource credit currency. For example, if you are will be dealing with processor-seconds of an integer resource unit, use 0 (which is the default). If you will be charging dollars and cents, then use 2. This parameter should be the same in the goldd.conf and gold.conf files.
- **log4perl.appender.Log.filename** — Used by log4perl to set the base name of the log file.
- **log4perl.appender.Log.max** — Used by log4perl to set the number of rolling backup logs.
- **log4perl.appender.Log.size** — Used by log4perl to set the size the log will grow to before it is rotated.
- **log4perl.appender.Log.Threshold** — Used by log4perl to set the debug level written to the log. The logging threshold can be one of TRACE, DEBUG, INFO, WARN, ERROR, and FATAL.
- **log4perl.appender.Screen.Threshold** — Used by log4perl to set the debug level written to the screen. The logging threshold can be one of TRACE, DEBUG, INFO, WARN, ERROR, and FATAL.
- **machine.autogen** [false] — If set to true, Gold will automatically create new machines when they are first encountered in a job function (charge, reserve, or quote). Additionally, a new machine will be automatically created if you try to add an undefined machine as a member of a project or account.
- **machine.default** [NONE] — If not set to NONE, Gold will use the specified default for the machine in a job function (charge, reserve, or quote) in which a machine was not specified.
- **project.autogen** [false] — If set to true, Gold will automatically create new projects when they are first encountered in a job function (charge, reserve, or quote). Additionally, a new project will be automatically created if you try to add an undefined project as a member of an account.
- **project.default** [NONE] — If not set to NONE, Gold will use the specified default for the project in a job function (charge, reserve, or quote) in which a project was not specified and no default project can be found for the user.
- **security.authentication** [true] — Indicates whether incoming message authentication is required.

- **security.encryption** [false] — Indicates whether incoming message encryption is required.
- **server.host** [localhost] — The hostname on which the Gold server runs.
- **server.port** [7112] — The port the Gold server listens on.
- **super.user** [root] — The primary Gold system admin which by default can perform all actions on all objects. The super user is sometimes used as the actor in cases where an action is invoked from within another action.
- **user.autogen** [false] — If set to true, Gold will automatically create new users when they are first encountered in a job function (charge, reserve, or quote). Additionally, a new user will be automatically created if you try to add an undefined user as a member of a project or account.
- **user.default** [NONE] — If not set to NONE, Gold will use the specified default for the user in a job function (charge, reserve, or quote) in which a user was not specified.

Client Configuration

The following configuration parameters may be set in the client configuration file (gold.conf).

- `log4perl.appender.Log.filename` — Used by log4perl to set the base name of the log file.
- `log4perl.appender.Log.max` — Used by log4perl to set the number of rolling backup logs.
- `log4perl.appender.Log.size` — Used by log4perl to set the size the log will grow to before it is rotated.
- `log4perl.appender.Log.Threshold` — Used by log4perl to set the debug level written to the log. The logging threshold can be one of TRACE, DEBUG, INFO, WARN, ERROR, and FATAL.
- `log4perl.appender.Screen.Threshold` — Used by log4perl to set the debug level written to the screen. The logging threshold can be one of TRACE, DEBUG, INFO, WARN, ERROR, and FATAL.
- `response.chunking` [true] — Indicates whether large responses should be chunked (segmented). If set to false, large responses will be truncated.
- `response.chunksize` [0] — Indicates the line length in the data response that will trigger message segmentation (or truncation). A value of 0 (zero) means unlimited, i.e. that the client will accept the chunksize set by the server. The response chunksize will be taken to be the smaller of the client and server chunksize settings.
- `currency.precision` [0] — Indicates the number of decimal places in the resource credit currency. For example, if you are will be dealing with processor-seconds of an integer resource unit, use 0 (which is the default). If you will be charging dollars and cents, then use 2. This parameter should be the same in the goldd.conf and gold.conf files.
- `security.authentication` [true] — Indicates whether outgoing message are signed.
- `security.encryption` [false] — Indicates whether outgoing messages are encrypted.
- `security.token.type` [Symmetric] — Indicates the default security token type to be used in both authentication and encryption.
- `server.host` [localhost] — The hostname on which the Gold server runs.
- `server.port` [7112] — The port the Gold server listens on.