

Overview

Grid Computing is an effective way of utilizing all available processing power in order to significantly cut projection and valuation run times. Whether you are running one scenario or one thousand scenarios, the GridMaster software will take these large jobs and logically divide them into smaller and more manageable pieces. These job parts will each be assigned to an available processor to be worked on simultaneously. The more processors available, the more job parts that can be processed at any point in time.

The increase in the amount of performance is nearly linear to the total effective processing power used. For example, if a valuation takes 10 hours to run on a single processor machine with 3.0 GHz of processing power, and the same job were run using the GridMaster software on 5 single processors machines (with a total of 15.0 GHz), the job would complete nearly 5 times faster—or in 2 hours.

Hardware Components

The grid consists of a control server or host machine and any number of volunteer or client machines. The host houses all of your data along with the GridMaster software required to run the valuation, projection, or source of earnings (SOE) study. The host is responsible for controlling the entire grid job from beginning to end. This includes dividing the total work into small, independent job parts; allocating each job part to an available client machine; monitoring the progress of each part; receiving the completed work from each machine; and compiling all of the data into the final output.

Client machines can be any machine available for work. While some machines maybe dedicated for grid usage only, other machines that do not perform hardware intrusive tasks can be utilized as well. The grid job can be run at a lower priority on these machines and utilize unused cycles without affecting any current operations.

The host and clients communicate through a Windows[®] network via mailslots, named pipes, and drive shares. Each client machine runs a Windows[®] service identified by a grid name that is specified upon installation. Client machines with multiple CPU's should have a unique grid name assigned to each CPU. This Windows[®] service allows the client machines to essentially “log on” to the host for direct access to the data needed to complete the assigned job part. The data is never physically stored on the client machines.

GridMaster Software

GridMaster can easily be implemented within your current Poly software on the host machine. Once all the initial installation procedures are completed by an administrator or tech, all that is needed are four executables, four PERL-scripts, and a few coded *.TBL files. The executables are designed to partition the input file, launch the Grid Job Management Console, and execute the appropriate scripts when necessary. The scripts provide full control over the details of the grid process, including performing the initial checks; making job part assignments; launching the projection, valuation, or SOE study with the appropriate parameters; and analyzing

the grid results. You simply need to code the Grid Computing tables (located under the *Admin System* menu) in addition to your regular projection, valuation, or SOE coding.

The *Grid Wizard* provides an easy way of coding most of the other Grid Computing tables. In this wizard you define a name for your job, the location of the corresponding scripts, the number of job parts, the client machines to be used, and how often progress is checked on job parts before determining any failures. The *Grid Wizard* only needs to be coded once for any particular job as long as these parameters remain unchanged. *Run Grid Job* must also be coded and defines the host table directory (where *.TBL files are kept), the host work directory (where information from the grid run is written), and the location of your executables. *Run Grid Job* also launches the Grid Job Management Console and begins processing the projection, valuation, or SOE study.