RiskAMP **CS**

RiskAMP Simulation Software for Windows HPC Server 2008

High-Performance Computing

RiskAMP CS uses the power of Windows HPC Server 2008 to run simulations faster & more efficiently.

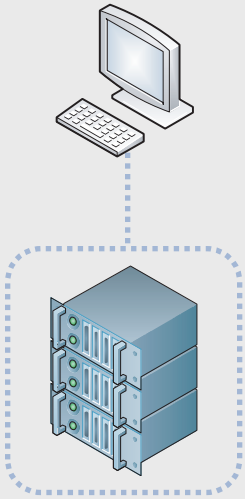

Simulations that take hours on a single workstation can be calculated in a fraction of the time on an HPC cluster.

Run simulations on the cluster using the same interface as the standard RiskAMP Excel desktop software.

RiskAMP CS

All the features of RiskAMP plus the power of High-Performance Computing

RiskAMP CS software for Windows HPC Server 2008 extends the capabilities of RiskAMP desktop software to support calculation on a high-performance Windows HPC Cluster.

- **Better Productivity** RiskAMP CS frees up your desktop to keep working while the cluster handles intensive calculations.
- **Faster Decision-making** The faster you get results, the better informed you'll be in making time-sensitive decisions.
- **Greater Certainty** With RiskAMP CS you can run larger and more complex simulations in less time.

From within Excel, you can submit simulation tasks to a cluster for calculation, dramatically reducing the time required to perform complex simulations. RiskAMP CS includes all the standard probability distributions and statistical analysis functions, supports Latin Hypercube Sampling (LHS), and supports all Excel functions and third-party add-ins.

On compute nodes, intensive computation can take advantage of multiple processors, and jobs are processed according to resource availability and priority. Simulations can be run asynchronously - meaning you can submit a simulation for processing and then go on with other work. When the simulation is complete, you can retrieve the results for review or share the data with other users and groups.

Structured Data, LLC
133 E 17th Street
New York, NY, 10003

1 (800) 573 6559
1 (917) 546 6939 fax
sales@riskamp.com

Using RiskAMP CS

From the user perspective, RiskAMP CS software works just like the standard RiskAMP Excel add-in for running Monte Carlo and probability simulations. In fact, you can run a standard desktop simulation at any time.

Running a simulation on the cluster works the same way. To send a job to the cluster, just click the "Send to Cluster" button on the ribbon or menu. You'll need to enter some parameters - the cluster address, the number of trials, and so on. This information is cached so you only need to enter it once.

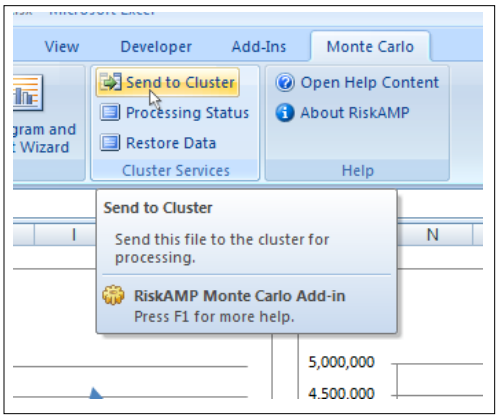

Figure 1: Send a job to the cluster

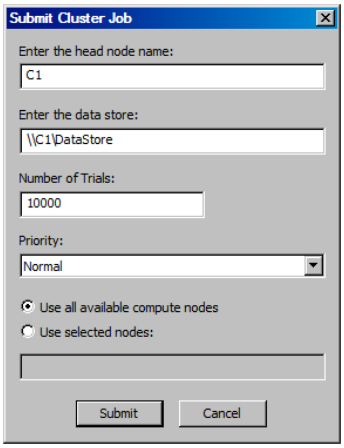

Once you have the configuration set up, click "Submit" and the job will go to the cluster for processing. Depending on the priority you set, the job may be completed immediately or it might wait for other pending jobs. Administrators can limit priority so critical jobs always run first.

When a job has been submitted, it will start calculating on the cluster as soon as resources are available. This doesn't tie up your desktop, so you can go on to other work until the simulation is complete.

Figure 2: Set the job parameters

When the job is complete, your spreadsheet will update with the new results. Any charts or graphs in your spreadsheet will be updated, and you can work with the "live" data just as if you had run the simulation on the desktop.

Structured Data, LLC
133 E 17th Street
New York, NY, 10003

1 (800) 573 6559
1 (917) 546 6939 fax
sales@riskamp.com

Cluster Operation

To the user, running a simulation on the cluster looks the same as running on the desktop - except the user's desktop isn't tied up with processing. Behind the scenes, a number of steps take place to prepare and then run the simulation.

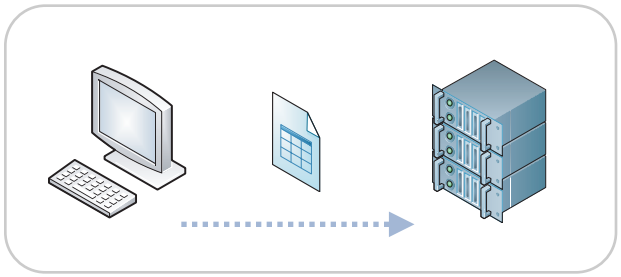

When the user sends the job to the cluster, first the file is saved & sent to a dedicated storage area. This ensures that all the simulation steps use a single, consistent version of the file.

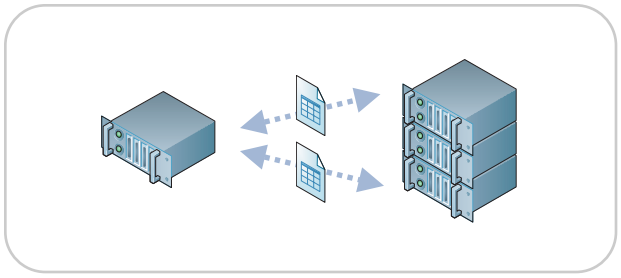

On the cluster, the head node divides the simulation job into a number of tasks. Each task contains some number of iterations. Individual tasks are sent to compute nodes when they have available resources. When a task is complete, the results are sent back to the head node.

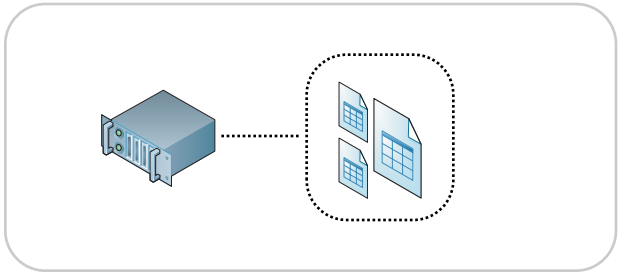

When all the simulation tasks are complete, a final tasks consolidates all the partial results, and generates the finished simulation worksheet.

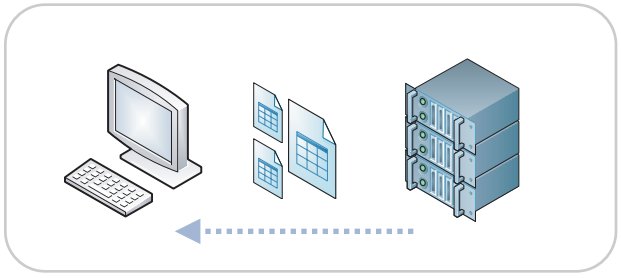

After the simulation is complete, the user's spreadsheet is automatically updated with the simulation results - including updating any charts or graphs in the spreadsheet.

Structured Data, LLC
133 E 17th Street
New York, NY, 10003

1 (800) 573 6559
1 (917) 546 6939 fax
sales@riskamp.com

System Requirements

To support RiskAMP CS, a typical cluster configuration includes a cluster head node, and one or more compute nodes. The standard cluster configuration is based on Windows HPC Server 2008 and Excel 2007. RiskAMP CS also supports Windows Server 2003 and the Windows Compute Cluster Pack.

At least one node within the cluster domain requires the RiskAMP CS Manager software, and network available disk space for temporary data storage. This is typically installed on the cluster head node, but can be installed on a separate host if desired. Compute nodes require Microsoft Excel 2007 or Excel 2003 and the RiskAMP CS Server add-in and control software.

On the desktop, RiskAMP CS requires Excel 2007 or Excel 2003 on Windows Vista or XP.

About RiskAMP Software

RiskAMP is desktop software for running Monte Carlo and probability simulations. The desktop software includes random distribution functions, statistical analysis functions, and an engine for performing simulations - recalculating a spreadsheet thousands of times to generate raw simulation data, then analyzing the data and generating statistics.

RiskAMP is an add-in for Microsoft Excel® supporting all Excel features, including VBA scripting, macros, and functions from other add-ins.

Structured Data, LLC
133 E 17th Street
New York, NY, 10003

1 (800) 573 6559
1 (917) 546 6939 fax
sales@riskamp.com