

CrossGrid User Guide

GridBench

WP2.3 Grid Benchmarking

Document Filename:	GridBench_User_Guide.pdf
Workpackage:	WP2.3 Grid Benchmarking
Partner(s):	UCY
Lead Partner:	UCY
Config ID:	cg-gridbench-user
Document classification:	PUBLIC

Abstract: This document explains what you need to know as a user of the GridBench system. It describes the usage of the GridBench Graphic User Interface.

Delivery Slip

	Name	Partner	Date	Signature
From	George Tsouloupas, Marios Dikaiakos	UCY	November 2004	
Verified By				
Approved By				

Document Log

Version	Date	Summary of changes	Author
1-0	30 Nov 2004	First public draft	George Tsouloupas, Marios Dikaiakos

Contents

CopyrightNotice	4
1 INTRODUCTION	5
1.1 Abbreviations and Acronyms	5
1.2 References and Source Code	5
2 PRODUCT USAGE	6
2.1 RUNNING THE PRODUCT	6
2.2 OPERATION	7
3 TROUBLESHOOTING Q&A	14
4 EDG License Agreement	15

Copyright Notice

Copyright (c) 2005 by **George Tsouloupas, Marios Dikaiakos**. All rights reserved.

Use of this product is subject to the terms and licenses stated in the EDG license agreement. Please refer to Section 4 for details.

This research is partly funded by the European Commission IST-2001-32243 Project CrossGrid.

1 INTRODUCTION

GridBench is a tool for benchmarking grids. It has two main objectives:

1. Generate metrics that characterize the performance capacity of resources belonging to a Virtual Organization and *spanning across multiple Grid nodes*, in terms of computational power, file-transfer speed, inter-process communication bandwidth, application-kernel performance, scalability etc.
2. Provide a *tool* for researchers that wish to investigate various aspects of Grid performance, using well-understood kernels that are representative of more complex applications deployed on the Grid. Having access to a corpus of such kernels and being able to easily specify and dispatch parameterized runs of these kernels on Grids, facilitates the characterization of factors that affect application and infrastructure performance, the quantitative comparison of different middleware solutions, algorithms for scheduling, resource allocation, etc.

This guide focuses on the usage of the GridBench GUI.

1.1 Abbreviations and Acronyms

API	Application Programming Interface
EDG	European DataGrid
JIMS	Jiro-base Infrastructure Monitoring System
GUI	Graphical User Interface
SANTA-G	Grid-enabled System Area Networks Trace Analysis
SQL	Structured Query Language

1.2 References and Source Code

The full source code for the Gridbench system can be found in the FZK CVS repository. Please refer to <http://gridportal.fzk.de/?group=cg-wp2.3> for details on how to obtain it.

The source code can also be browsed at the following address:

http://savannah.fzk.de/cgi-bin/viewcvs.cgi/crossgrid/crossgrid/wp2/wp2_3-bench/

2 PRODUCT USAGE

The GridBench User Interface provides a virtual "workbench" for performing benchmarking experiments. It provides tools for the definition of benchmarks and their execution as well as the analysis of results. Performing basic benchmarking experiments is as simple as a drag-and-drop.

2.1 RUNNING THE PRODUCT

The GridBench GUI is available via the CrossGrid Migrating Desktop. (Stand-alone operation is also possible and will be described in a future version of this guide)

2.1.1 Operating Requirements

Local Software Requirements

(Requirements from the Migrating Desktop are inherited)

Grid Infrastructure Requirements

The GridBench services must be installed and running as specified in the configuration file described in the following sub-section.

2.1.2 Step-by-Step User Setup

The configuration file for GridBench is found under:

```
/opt/cg/etc/gridbench.conf
```

This configuration file is common for the client (the GridBench GUI) and the GridBench services. Here we will explain the parts of the configuration file that concern the client.

A sample configuration for the Gridbench GUI is found below:

```
WSHOST=ui001.grid.ucy.ac.cy  
WSPORT=8080
```

```
information_index=ii01.lip.pt  
information_index_port=2170
```

```
information_index2=ic.fzk.de  
information_index_port2=2170
```

```
information_index3=ce101.grid.ucy.ac.cy  
information_index_port3=2135
```

```
wrapper_path=/opt/cg/gridbench/bin/gb_wrapper.sh  
plugins=globus edg
```

WSHOST and WSPORT indicate the host and port where the GridBench web-services can be found. An arbitrary number of information indices (MDS) can be specified as in the example above. Additionally, individual CE's can be included as shown in: (note the port number)

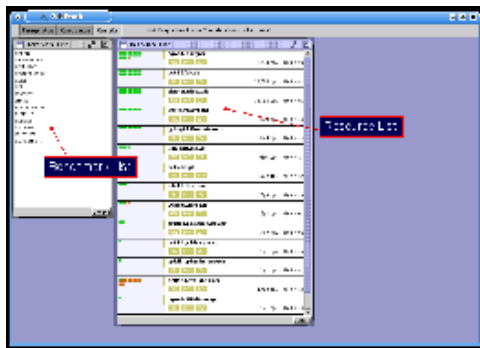
```
information_index3=ce101.grid.ucy.ac.cy
information_index_port3=2135
```

Do not modify the wrapper_path and plugins variables.

2.2 OPERATION

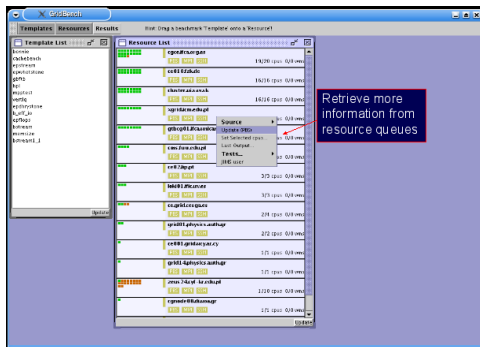
2.2.1 The GUI

The GridBench GUI is started from within the Migrating Desktop by clicking on the GridBench button in the tool bar or by selecting it from the plugins list.

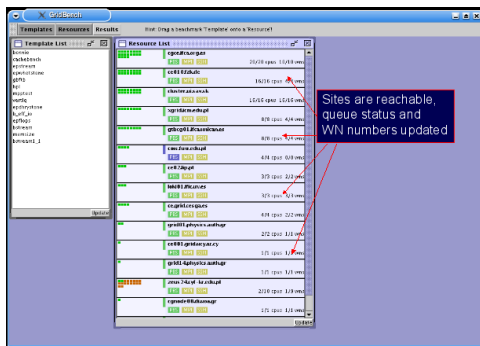


When the GridBench GUI starts the user is presented with the GridBench "workbench" containing two lists: *i*) the benchmark templates list and *ii*) the resources list. The benchmark templates list is a list of currently installed benchmarks. The Resource list is a list of resources (CE's) obtained by querying the MDS servers specified in the configuration file

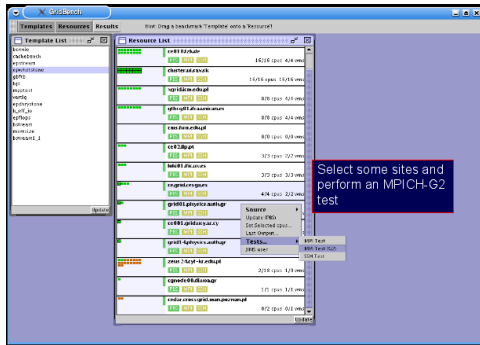
2.2.2 Performing tests.



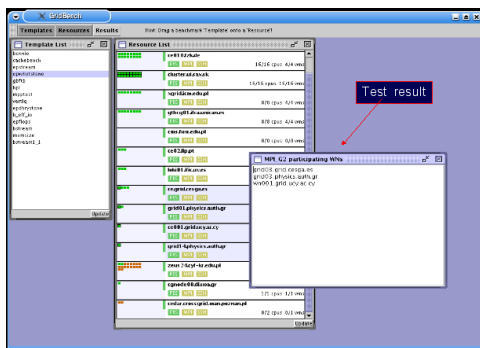
The information contained in the Resources list is derived from MDS. The information may be "stale" so the user can get an up-to-date status of the resources by querying the local PBS queues. Select all the resources (press ctrl-A while in the Resources list window) then right-click, then "update PBS". This will also test whether PBS is functional.



When a test is invoked (as in the "update PBS") the status of the test turns blue, until it completes successfully and turns green. If the test fails it turns red. When the test finishes, the actual number of Worker Nodes (not CPU's) available in a resource is also updated (this information cannot be found in the MDS).

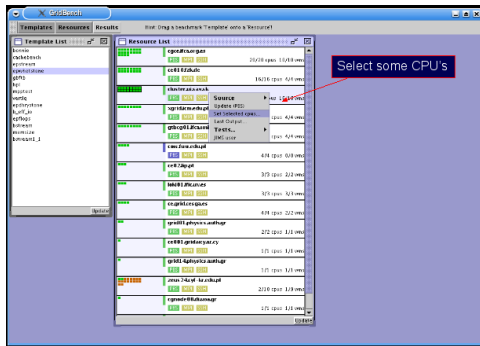


Performing an MPICH-G2 test: In order to test whether a set of resources can participate in an MPICH-G2 job, select the set of resources, then right-click, then “Tests...”, then “MPI Test (G2)”.

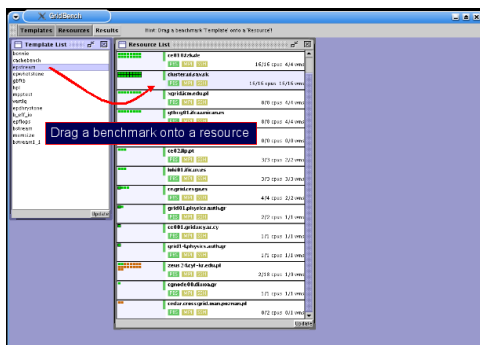


After a while the result of the test will be presented. In this case a problem is detected: only three sites reply. This could be due to incorrect configuration, or a transient error, or sometimes to unavailable resources.

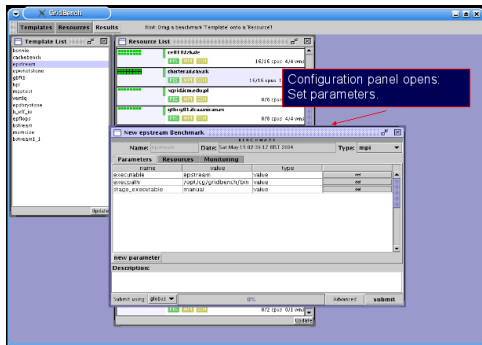
2.2.3 Configuring a benchmark.



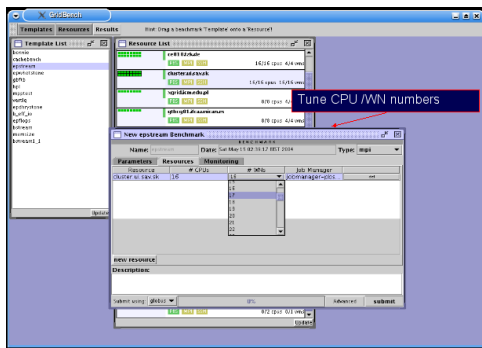
Select the number of participating CPU’s by right-clicking on a resource and “Set selected CPU’s”.



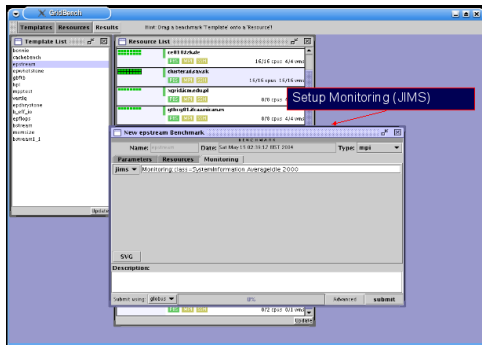
Drag a benchmark from the benchmark templates list and drop it onto a resource.



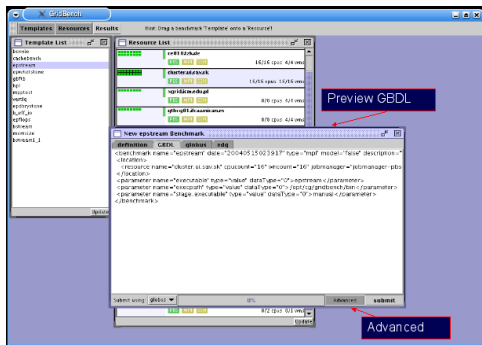
A default benchmark configuration is presented in a new benchmark configuration panel. The default values for the parameters are usually the best values to use.



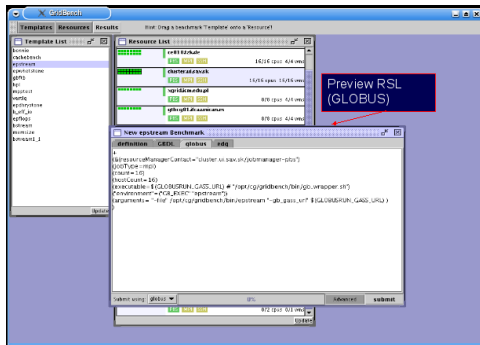
The “resources” is preconfigured based on which resource the benchmark was dragged onto. If the resource had “selected” CPU’s the that value is used, otherwise the default is the number of free CPU’s. The user is provided with the option to specify how the CPU assignment will be performed (e.g. 8 CPU’s on 8 worker-nodes, or 8 CPU’s on 4 (dual) worker-nodes).



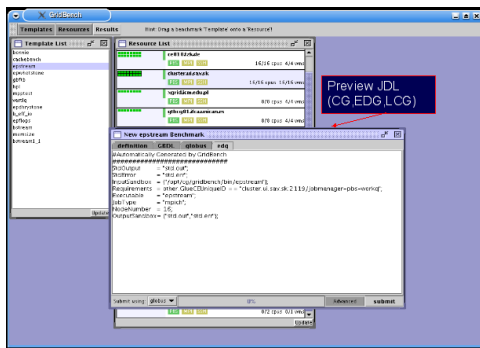
The user can request that the resource on which the benchmark will run will be monitored using JIMS monitoring. The user specifies three parameters: the MonitoringClass (e.g. SystemInformation, the attribute (e.g. AverageIdle CPU time) and the frequency of collection (e.g 2000ms).



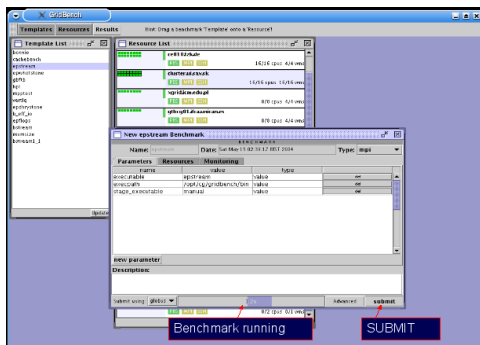
By clicking on the “Advanced” button the user can review the underlying GridBench Definition Language, which is an XML definition of the benchmark that will be executed. This is the definition that will be sent to the GridBench services that will execute the benchmark.



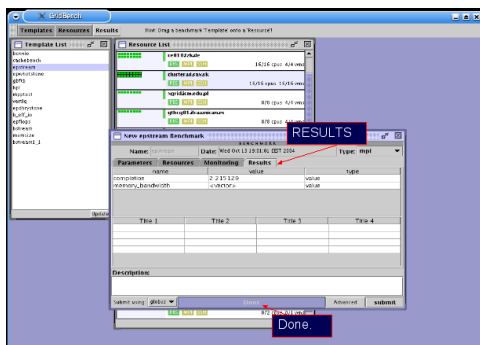
The “globus” tab shows the generated RSL. (Generated by the “globus” middleware plugin).



The “edg” tab shows the generated JDL. (Generated by the “edg” middleware plugin).

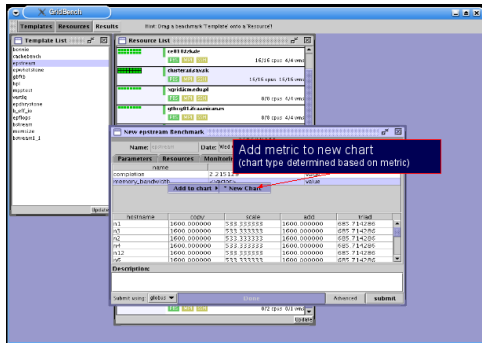


Once the defaults have been tuned, the user can go ahead and submit the benchmark by clicking on the “submit” button. While the benchmark is running the progress bar indicates the elapsed time until the benchmark finishes. The user can specify the mechanism by which to execute the benchmark by specifying which middleware plugin to use. This is done by selecting a plugin from the “Submit using” drop-down list (bottom-left of the benchmark configuration panel)

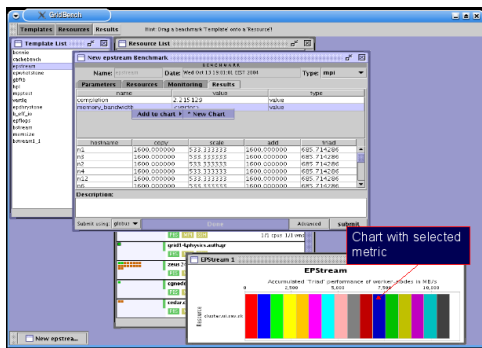


When the benchmark finishes a new tab is added to the benchmark configuration panel which includes the results and indicated that the benchmark executed successfully. (The result is automatically saved in the archive.)

2.2.4 Analysing Results.

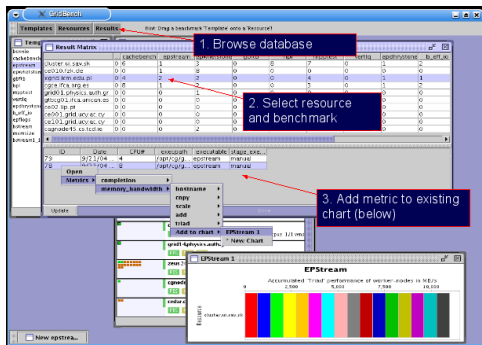


The results can be reviewed by inspecting the contents of the “results” tab. The user can generate a chart from the resulting metrics by selecting a metric, then right-click, then “Add to chart”, then “New Chart”.

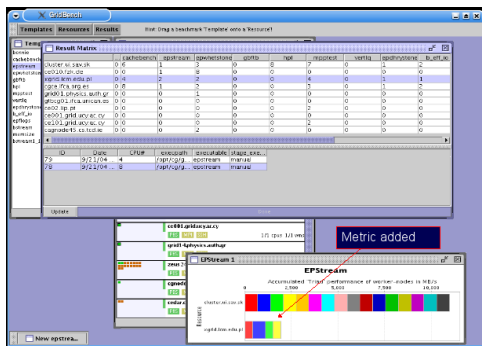


A new chart containing the metric is presented. In this case it is a stacked bar-chart showing the result of the benchmark just executed.

The user can compare this result by retrieving data from the archive. To open a browser of the Archive the user clicks on the “Results” button at the top. The user is presented with a “Result matrix” where all the results in the archive are given.

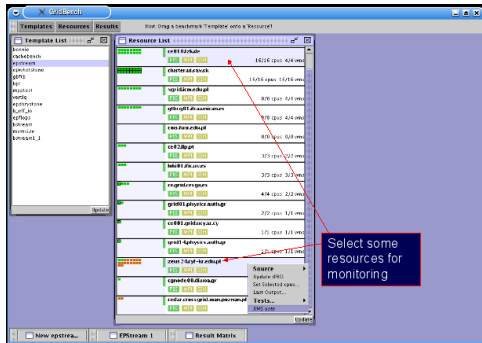


In the “Result Matrix” each column gives the number of available results for a given benchmark while results for a given resource are given on the corresponding row. For example there are 2 executions of *epstream* on *xgrid.icm.edu.pl*. By clicking on this cell in the matrix, the executions appear in the table below the result matrix. By right-clicking the selected execution the user can add it to the existing chart or to a new chart.



The metric is added to the existing chart and the user can use the generated chart to review and evaluate the results. The user can repeat the process an arbitrary number of times.

2.2.5 Live monitoring

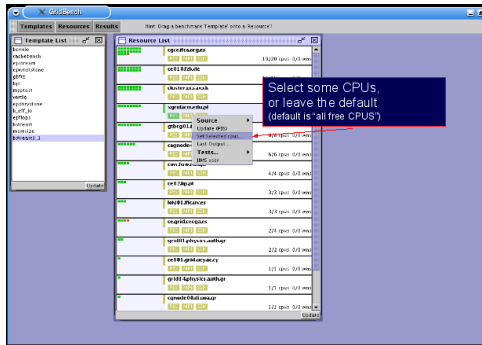


The user can have access to live monitoring data from JIMS. The user can select a set of resources the right-click then “JIMS user”. This will immediately start to collect and show data from the selected resources (at present it shows CPU usage).

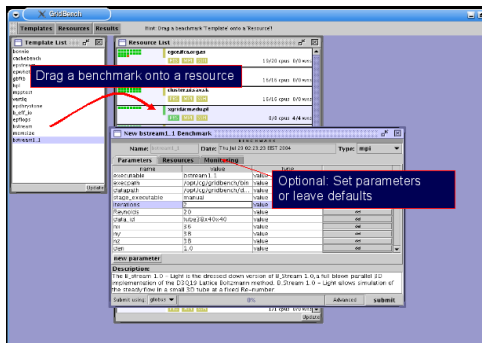


Data from each resource is given in a separate window. In the example presented on the left it can be seen that the nodes that appear red (busy) in the resources list *obviously* display a high CPU usage (the 200% is due to the dual-CPU nodes)

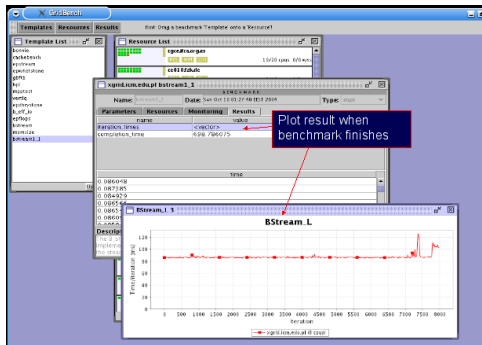
2.2.6 A simple application kernel-based benchmark experiment.



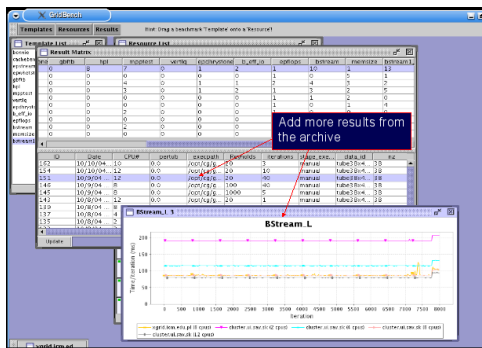
Optional: Pick a resource and set its selected CPU's.



Drag the Bstream benchmark onto a resource. Leave the default parameters and click submit.



When the benchmark finishes plot the resulting “iteration_times” metric.



Open the “Results” and add some older measurements to the chart. The user can now evaluate the performance of a resource based on the results from this kernel.

3 TROUBLESHOOTING Q&A

Q: On start-up the “Template List” appears empty.

A: Click “Update” on the “Template List”. If it still fails, make sure the correct WSHOST and WSPORT are given, the services are running and they are accessible (firewalls etc.).

Q: On start-up the “Resource List” appears empty.

A: Make sure the correct MDS index servers are given in the configuration file, and they are accessible (firewalls etc.).

4 EDG License Agreement

Copyright (c) 2005 CrossGrid. All rights reserved.

This software includes voluntary contributions made to the CrossGrid Project. For more information on CrossGrid, please see <http://www.eu-crossgrid.org>.

Installation, use, reproduction, display, modification and redistribution of this software, with or without modification, in source and binary forms, are permitted. Any exercise of rights under this license by you or your sub-licensees is subject to the following conditions:

1. Redistributions of this software, with or without modification, must reproduce the above copyright notice and the above license statement as well as this list of conditions, in the software, the user documentation and any other materials provided with the software.

2. The user documentation, if any, included with a redistribution, must include the following notice:

his product includes software developed by the CrossGrid Project (<http://www.eu-crossgrid.org>).

Alternatively, if that is where third-party acknowledgments normally appear, this acknowledgment must be reproduced in the software itself.

3. The names CrossGrid and Gmay not be used to endorse or promote software, or products derived therefrom, except with prior written permission by cgooffice@cyfronet.krakow.pl.

4. You are under no obligation to provide anyone with any bug fixes, patches, upgrades or other modifications, enhancements or derivatives of the features, functionality or performance of this software that you may develop. However, if you publish or distribute your modifications, enhancements or derivative works without contemporaneously requiring users to enter into a separate written license agreement, then you are deemed to have granted participants in the CrossGrid Project a worldwide, non-exclusive, royalty-free, perpetual license to install, use, reproduce, display, modify, redistribute and sub-license your modifications, enhancements or derivative works, whether in binary or source code form, under the license conditions stated in this list of conditions.

5. DISCLAIMER

THIS SOFTWARE IS PROVIDED BY THE CROSSGRID PROJECT AND CONTRIBUTORS AS IS AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, OF SATISFACTORY QUALITY, AND FITNESS FOR A PARTICULAR PURPOSE OR USE ARE DISCLAIMED. THE CROSSGRID PROJECT AND CONTRIBUTORS MAKE NO REPRESENTATION THAT THE SOFTWARE, MODIFICATIONS, ENHANCEMENTS OR DERIVATIVE WORKS THEREOF, WILL NOT INFRINGE ANY PATENT, COPYRIGHT, TRADE SECRET OR OTHER PROPRIETARY RIGHT.

6. LIMITATION OF LIABILITY

THE CROSSGRID PROJECT AND CONTRIBUTORS SHALL HAVE NO LIABILITY TO LICENSEE OR OTHER PERSONS FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL, EXEMPLARY, OR PUNITIVE DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, LOSS OF USE, DATA OR PROFITS, OR BUSINESS INTERRUPTION, HOWEVER CAUSED AND ON ANY THEORY OF CONTRACT, WARRANTY, TORT (INCLUDING NEGLIGENCE), PRODUCT LIABILITY OR OTHERWISE, ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.