Επιμετρήσεις Επιδόσεων Υπολογιστικού Πλέγματος

«Test-driving the Grid»

Μάριος Δ. Δικαιάκος

Τμήμα Πληροφορικής
Πανεπιστήμιο Κύπρου

Acknowledgments

- In collaboration with:
  - George Tsouloupas, UCY
- With funding from:

Choosing an automobile

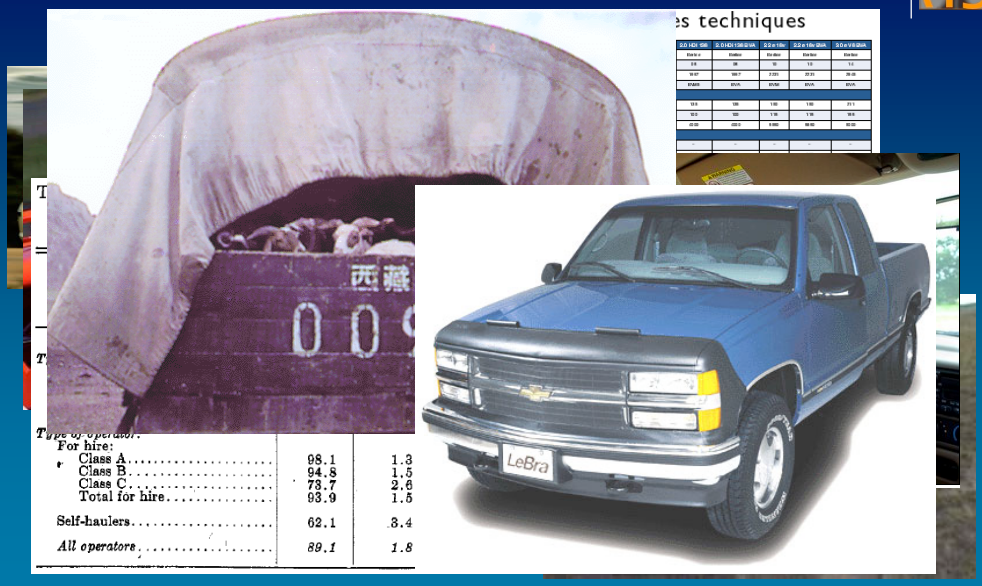

Type of operator		
For hire:		
Class A	98.1	1.3
Class B	94.8	1.5
Class C	73.7	2.6
Total for hire	93.9	1.5
Self-haulers	62.1	3.4
All operators	*89.1*	*1.8*

Benchmarks: definition

- Standardized programs (or detailed specifications) designed or chosen to investigate performance properties of computer systems:
  - Characterize performance capacity and behavior.
  - Compare different systems in a fair manner.
  - Guide the optimization and assessment of system designs and implementation.
  - Help researchers establish quantitative arguments in systems research.
- Complete applications, kernels, probes, or synthetic programs. E.g:
  - Whetstone, SPEC, Parkbench, NAS, TPC-C, Linpack, SPLASH, MediaBench…

Benchmarking process

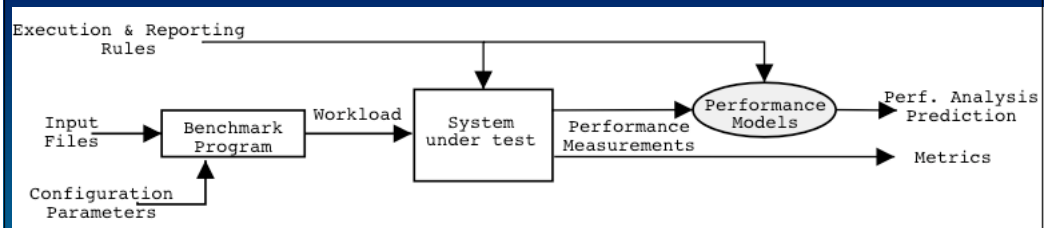

- Conditions of execution and measurement:
  - Well-defined and publicly available.
- Benchmarks are required to be:
  - Portable, fair, relevant, easy to measure, easy to explain.

Benchmarking: not a panacea

- Questions of fairness and relevance.
  - Relatively easy to optimize systems for selected benchmarks; what about real workloads?
- Questions of cost.
  - Esp. when benchmarking network-centric systems.
- Provide no direct insights about derived metrics.
  - The more complex a system is, the more difficult it becomes to combine individual measurements into a meaningful system evaluation.
- What about performance prediction based on benchmarking?
  - Benchmark results need to be interpreted in the context of models for hardware, architecture, available resources.

The Grid

- Middleware infrastructure that enables flexible, secure, coordinated resource sharing among dynamic collections of individuals and institutions (Foster, Kesselman, Tuecke).
- Enables communities ("Virtual Organizations") to share geographically distributed resources as they pursue common goals -- *assuming the absence of…*
  - Homogeneity
  - Central location
  - Central control
  - Existing trust relationships
- Enforces some level of resource virtualization.

Grid operation

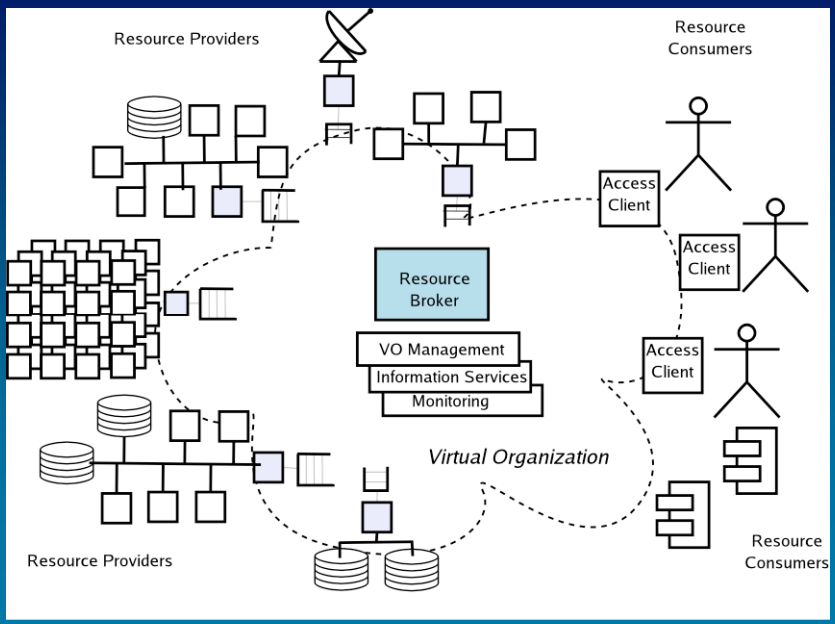

A Grid (interactive) Application

Surgical Planning

- Problem: vascular diseases
- Solution: placement of a bypass by a surgeon
- Planning for intervention is based on 3D images obtained from MRI or CT scans.
- The attainable improvement in blood flow should determine which possibility is the best for a particular patient.
- A 3D arterial model is built on the basis of the images, and presented to the surgeon in an inmersive environment.

Source: Univ. of Amsterdam
CrossGrid project

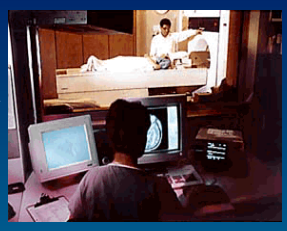

Observation

Stenosis
(narrowing of an artery)

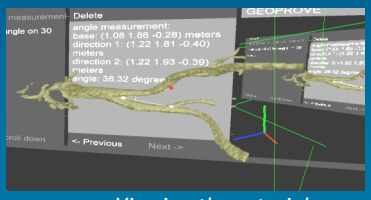

Viewing the arterial
structure in an immersive 3D environment

http://grid.ucy.ac.cy

Surgical Planning

- Goal: Simulate vascular reconstruction
- Method:
  - Interactive Virtual Reality Environment to
    - View scanned data
    - Define proposed interventions
    - View simulation results
  - Advanced fluid code to simulate flows

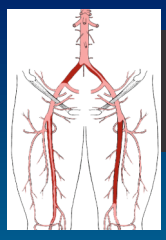

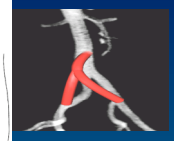

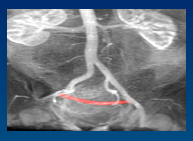

Arterial structures from scans with proposed bypasses

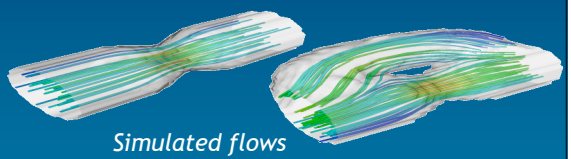

Simulated flows

- Need Grid in interactive mode (the surgeon should not wait long…)
  - Access distributed computational resources for flow simulation and visualization, so get a high performance environment at low cost
  - Distribute simulations for different bypass configurations

http://grid.ucy.ac.cy

How about resource selection?

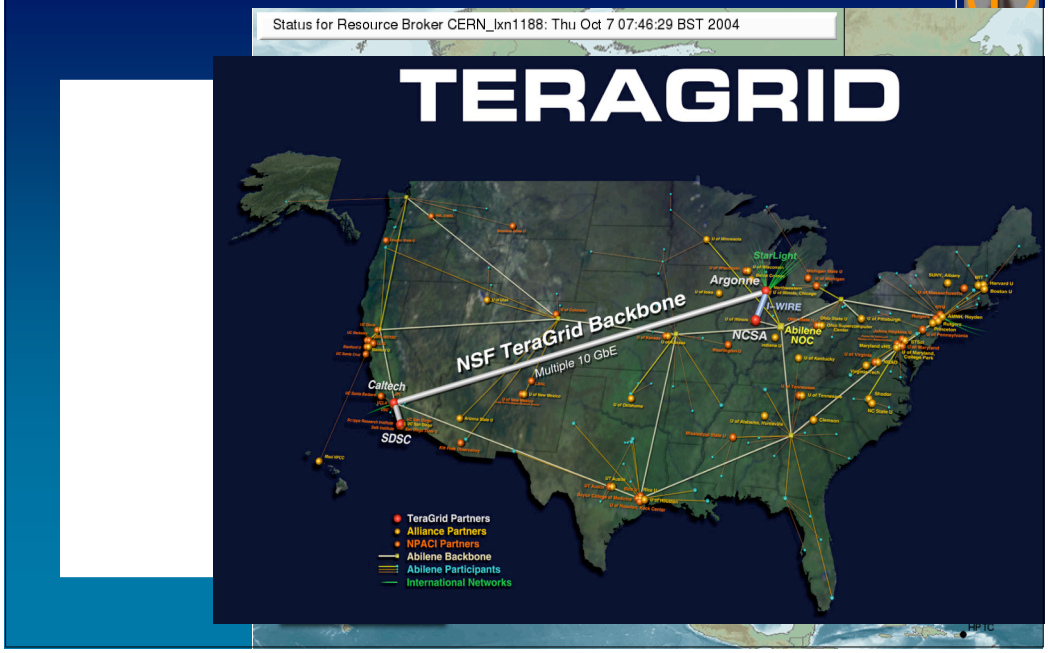

Ask a Grid Information Service!

```
dn: GlueSubClusterUniqueID=cgce.ifca.org.es,
       GlueClusterUniqueID=cgce.ifca.org.es,

dn: GlueCEUniqueID=cgce.ifca.org.es:2119/jobmanager-pbs-short,
       Mds-Vo-name=ifcapro,mds-vo-name=local,o=grid

GlueChunkKey: GlueClusterUniqueID=cgce.ifca.org.es
GlueHostApplicationSoftwareRunTimeEnvironment: CG2_0_4

GlueCEUniqueID: cgce.ifca.org.es:2119/jobmanager-pbs-short
GlueCEInfoGatekeeperPort: 2119
GlueCEInfoHostName: cgce.ifca.org.es
GlueCEInfoLRMSType: pbs
GlueCEInfoLRMSVersion: OpenPBS_2.4
GlueCEInfoTotalCPUs: 20
GlueCEStateEstimatedResponseTime: 0
GlueCEStateFreeCPUs: 20
GlueCEStateRunningJobs: 0
GlueHostMainMemoryVirtualSize: 1144

GlueForeignKey: GlueClusterUniqueID=cgce.ifca.org.es
```

Motivation and Focus

- How can we characterize the performance of Grid resources?
  - ⇨ Support more advanced criteria for choosing resources: performance, cost, functionality, reliability, robustness…
  - ⇨ Drive the design and configuration of Grid infrastructures.
  - ⇨ Open marketplaces based on performance negotiation.
  - ⇨ Developing models for performance prediction.

- ***"We have no real idea how the Grid and Grid applications could be characterized from the point of view of performance"*** *(APART Working Group on Automatic Performance Analysis, Rackeve Workshop, 11/2003)*

GridBench

- The GridBench Framework:
  - A software tool (workbench)for characterizing the performance of Grids and Grid resources quantitatively, using benchmarks.

- GridBench Suite of Benchmarks:
  - A hierarchical suite of benchmarks deployed on a Grid testbed.
  - Geared towards high-performance and high-throughput computing needs.

Key challenges for benchmarking Grids

- Measuring a moving and fuzzy target
- Trusting our measurements
- Describing performance (metrics)
- Managing cost
- In search of relevance

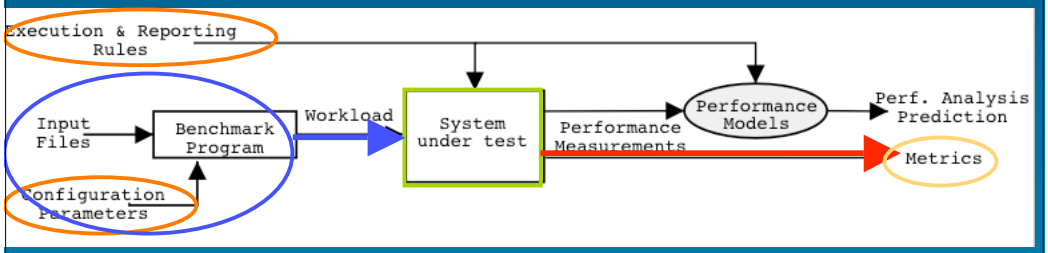

Measuring a moving target

- Grid jobs run upon heterogeneous resources which are:
  - Assembled dynamically and subject to change.
  - Described inaccurately or inadequately by GIS.
  - Often not properly operable, due to configuration errors, operator faults etc.
- Hence:
  - We need to capture the real set of resources we try to characterize.
  - Functionality benchmarking is equally important.

Trusting measurements

- A benchmarking job submitted to the Grid will run upon resources for which:
  - Exclusivity is not guaranteed.
  - Often, a benchmark runs in co-location with other jobs.
- Hence:
  - We cannot trust all of our measurements.
  - Need to filter out polluted measurements.

Describing performance (metrics)

- Performance capacity of Grid infrastructure is defined by:
  - The performance of a hierarchical collection of measurable entities (CPUs, computers, clusters, collections of clusters..)
- Thus:
  - Small sets of metrics not adequate for Grids.
  - Definition, organization, storage, and interpretation requires advanced, open data models, amenable to post-processing (statistical, data mining, AI).
- ...what about interpretation of metrics?

Managing cost

- Grid benchmarking seeks to capture a complete and valid "performance snapshot" of a Grid infrastructure. But:
  - Not all resources are available and running properly at a given time.
  - Jobs are susceptible to partial failure and degraded performance.
- Hence, the derivation of a complete snapshot requires:
  - Series of experiments, measurements, analyses.
  - Integration of metrics from multiple sites and runs.
  - Job submission to the Grid takes effort!

⇨ The cost of benchmarking increases substantially.

In search of relevance

- Relevant benchmarks are the ones producing realistic workloads, i.e., workloads representative of:
  - The prevalent programming model.
  - "Killer" Grid applications.
- However, the field:
  - Is still not mature; a prevalent programming model has yet to arise.
  - Production-quality Grid infrastructures are just beginning to emerge; no clear clues about typical workloads and applications.

Elements of GridBench

- A hierarchical model for the Grid architecture.
- A hierarchical suite of benchmarks characterizing the performance of abstract-model elements [HPC, HTC, MPI].
- A platform-independent language (GBDL) for specifying the configuration and for representing the conditions and results of benchmarking experiments.
- **GridBench**: a virtual workbench for administering Grid benchmarks, archiving, publishing and browsing results.

Outline

- **A simple reference architecture for the Grid.**
- The GridBench suite of benchmarks.
- GridBench description language.
- Filtering polluted measurements.
- The GridBench software.
- Using GridBench to characterize Grid sites.
- Conclusions and Future Work.

A simple reference architecture

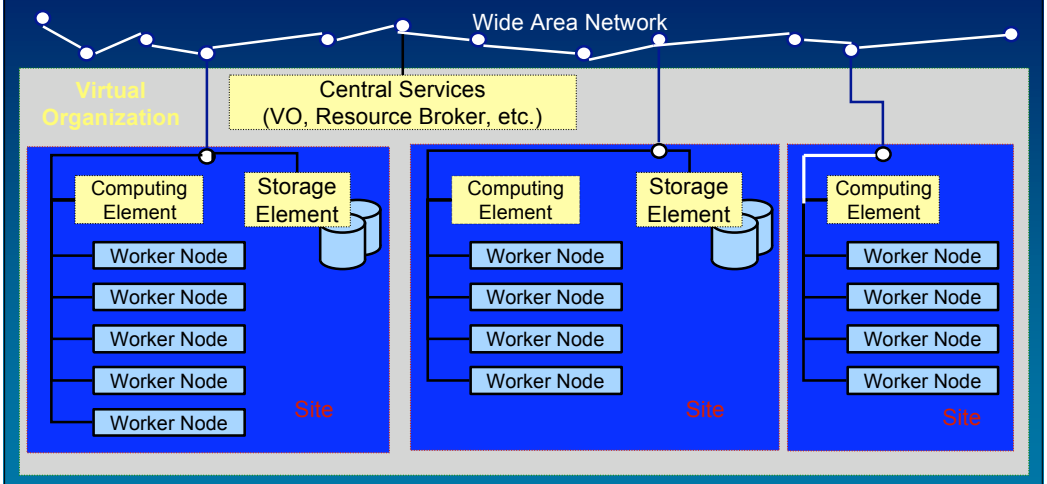

- ✓ Inspired by the DataGrid/CrossGrid/LCG architecture (Globus 2-based)
- ✓ Represented by the GLUE Schema.
- ✓ Necessary for defining benchmarking-targets and interpreting metrics.

GridBench: a hierarchical approach

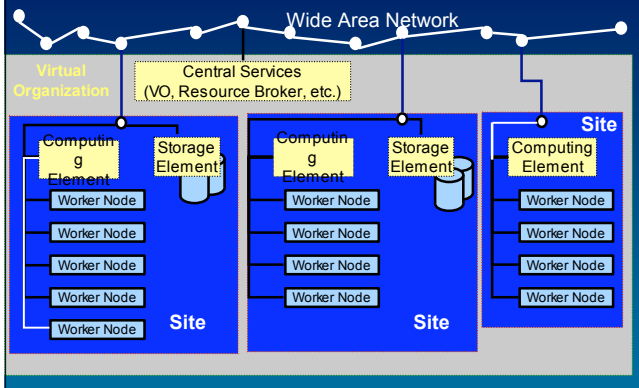

- Performance measurements at the different levels of the Grid architecture.

A hierarchical approach

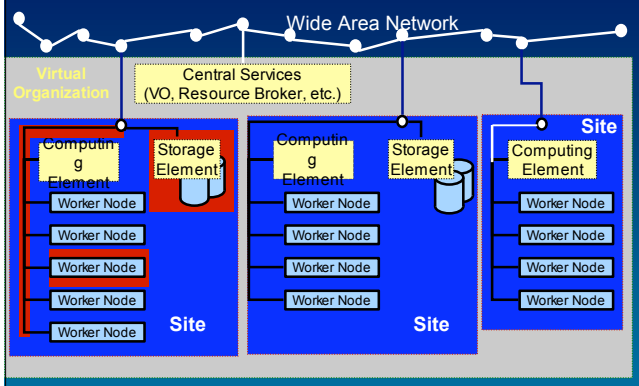

Individual Resources
(cluster nodes, mass storage)

- Performance measurements at the different levels of the Grid architecture.

A hierarchical approach

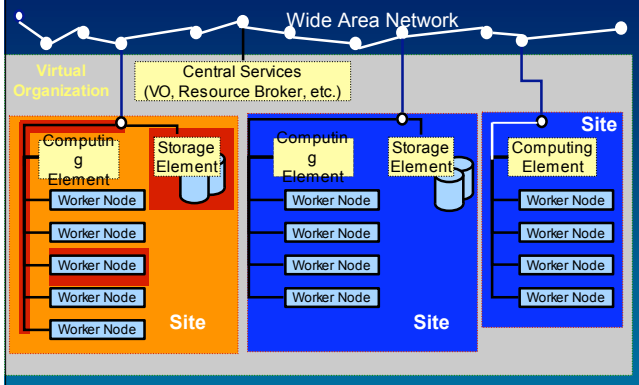

Individual Resources
(cluster nodes, mass storage)

Sites
(clusters, SMPs)

- Performance measurements at the different levels of the Grid architecture.

http://grid.ucy.ac.cy

A hierarchical approach

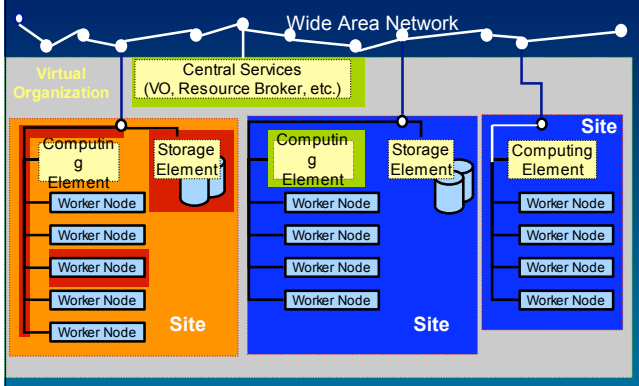

Individual Resources
(cluster nodes, mass storage)
Sites
(clusters, SMPs)
Middleware
(middleware layer providing access to shared resources)

- Performance measurements at the different levels of the Grid architecture.

A hierarchical approach

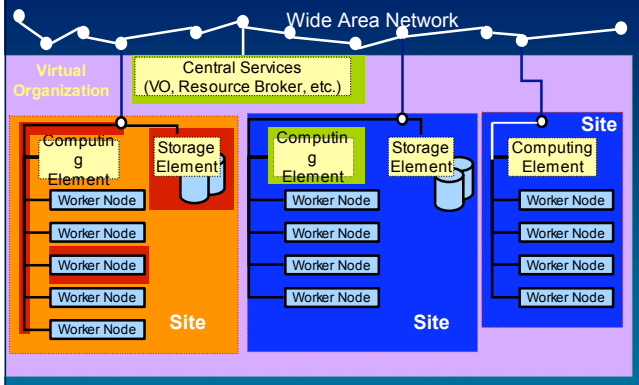

Individual Resources
(cluster nodes, mass storage)
Sites
(clusters, SMPs)
Middleware
(middleware layer providing access to shared resources)
Virtual Organization
(multiple sites, VO)

- Performance measurements at the different levels of the Grid architecture.

http://grid.ucy.ac.cy

The GridBench suite of benchmarks

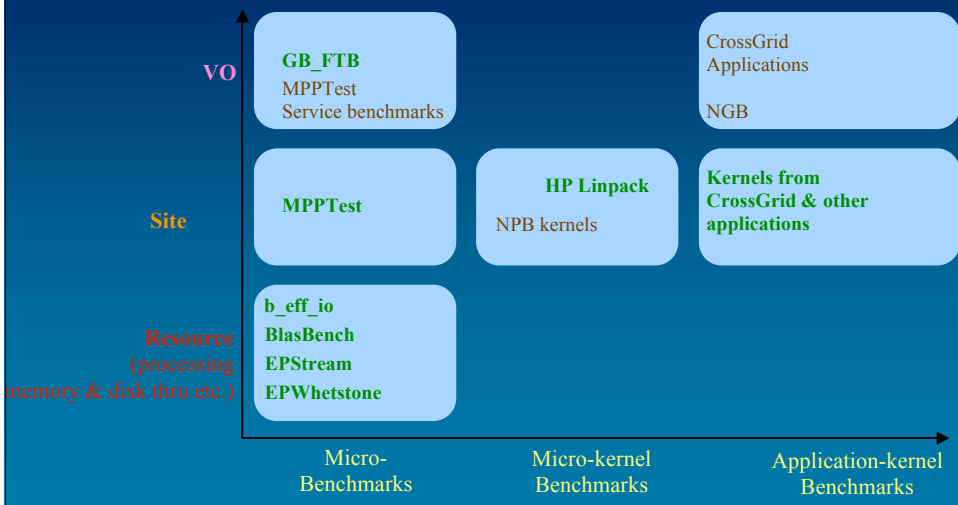

http://grid.ucy.ac.cy

The GridBench Description Language

- An XML-schema specification representing benchmarking metadata.
- Allows easy transformation to different job description formats.
- Benchmark definition co-exists with benchmarking configuration and results.
- GBDL elements include metadata about benchmark components:
  - Parameters
  - Location
  - Co-requisites
  - Prerequisites
  - Constraints
  - Metrics
  - Metrics
  - Archive
  - Monitoring

http://grid.ucy.ac.cy

GBDL structure

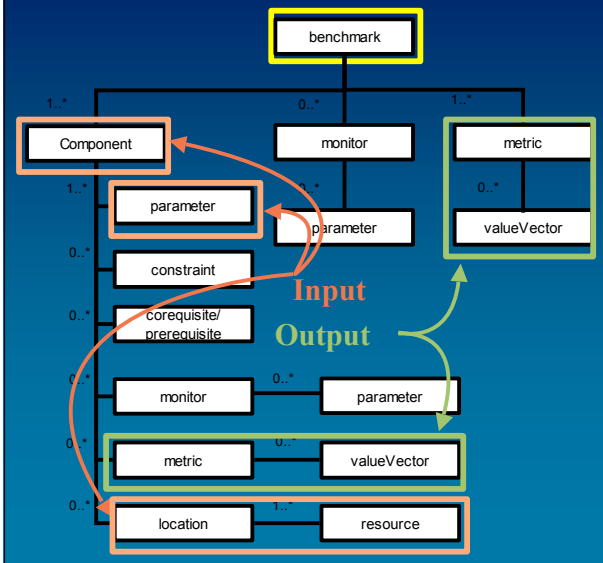

Metrics representation

- Benchmark-level metrics vs. component-level metrics:
  - e.g. benchmark completion time vs. component completion time.
- Single-value metrics

```
<metric name="epwhetstone_IPS" type="value">
    <vector unit="MIPS">53</vector>
</metric>
```

- Multi-value metrics

```
<metric name="xfer-rate" type="list">
    <vector unit="bps">
       2039430,2083930,1909830,2184750,...
    </vector>
    <vector unit="second" toffset="1055327287">
          0,10,20,30,...
    </vector>
</metric>
```

GBDL example

```
<benchmark date="031010121503" name="epwhetstone">
  <component id="A" name="epwhetstone" type="mpi">
      <location type="multi">
        <resource cpucount="2" name="apelatis.grid.ucy.ac.cy"/>
        <resource cpucount="2" name="ce010.fzk.de"/>
      </location>
      <parameter name="execpath" type="attribute">/opt/cg/gb/bin</parameter>
      <parameter name="stage_executable" type="attribute">manual</parameter>
      <parameter name="executable" type="attribute">epwhet</parameter>
      <parameter name="nloops" type="value">50000</parameter>
      <metric name="completion" unit="s">28.23</metric>
      <metric name="whetstone_mips" type="value">
       <vector name="hostname">wn001.grid.ucy.ac.cy
                              wn001.grid.ucy.ac.cy
                               n16.fzk.de
                               n15.fzk.de
       </vector>
       <vector name="mips">684.023 671.813 784.234 777.564 </vector>
      </metric>
  </component>    </benchmark>
```


Outline

- A simple reference architecture for the Grid.
- The GridBench suite of benchmarks.
- GridBench description language.
- **Filtering polluted measurements.**
- The GridBench software.
- Using GridBench to characterize Grid sites.
- Conclusions and Future Work.

Identifying polluted measurements

- Often a benchmark will run on a resource, in co-location with other jobs:
  - "Fellow passengers:" co-allocated by the Resource Broker.
  - "Free-riders:" unauthorized users, unknown to the VO.
  - "Runaways:" O/S processes, zombies, etc.
- These jobs may "pollute" our measurements and affect seriously the characterization accuracy.
- Their effect on metric accuracy can be:
  - Identified through monitoring.
  - Reduced through access control (for fellow passengers) and remote healing (for runaways).

Filtering polluted measurements

- GridBench retrieves monitoring information from existing Grid monitoring services, for a time window encapsulating a benchmarking experiment.
- The specification of the monitoring service and the data to be retrieved from it, is included in the GBDL document describing a specific experiment.

```
<component name="data-transfer" ID="xfer01">...</component>
<monitor type="RGMA" source="ccwp71.in2p3.fr:3306"
  query="select * from NetworkTCPThroughput
         where NMIdSource='adc0003.cern.ch'
           and NMIdDestination='ccwp7.in2p3.fr'
  <parameter name="begin">comp-begin="xfer01"</parameter>
  <parameter name="end">comp-end="xfer01"</parameter>
</monitor>
```

- Retrieval from monitoring services is conducted via monitoring-client plug-ins called by GridBench.

Monitoring a benchmark

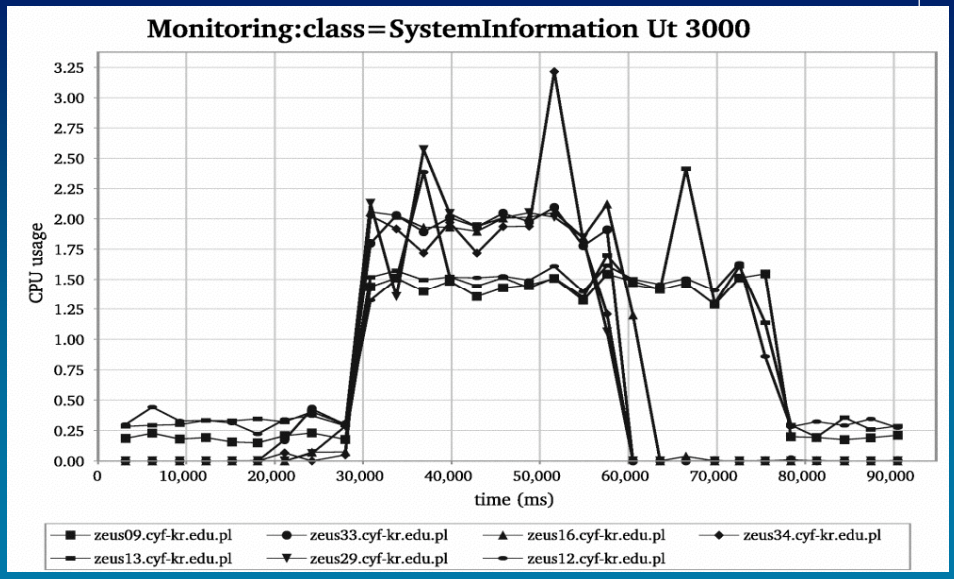

Filtering polluted measurements

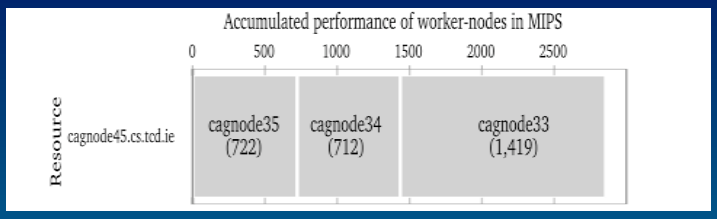

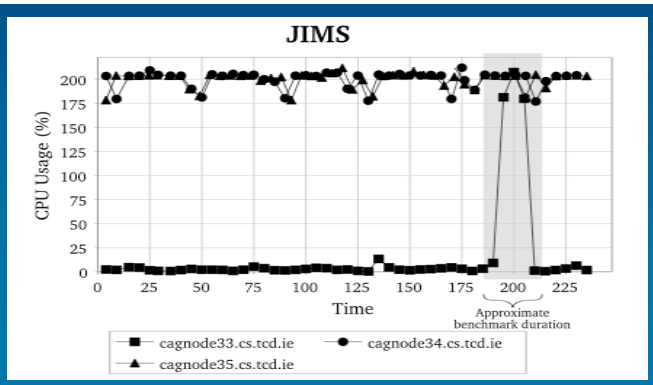

Outline

- A simple reference architecture for the Grid.
- The GridBench suite of benchmarks.
- GridBench description language.
- Filtering polluted measurements.
- **The GridBench software.**
- Using GridBench to characterize Grid sites.
- Conclusions and Future Work.

GridBench: requirements & functionality

- Supports:
  - Specification and execution benchmarks on a Grid.
  - Collection and archival of results.
- Archives together:
  - Benchmark specifications and measurements for publication and further analysis.
  - Monitoring information to help with result interpretation.
- Supports:
  - Retrieval and graphical representation of metrics.

GridBench Context: CrossGrid architecture

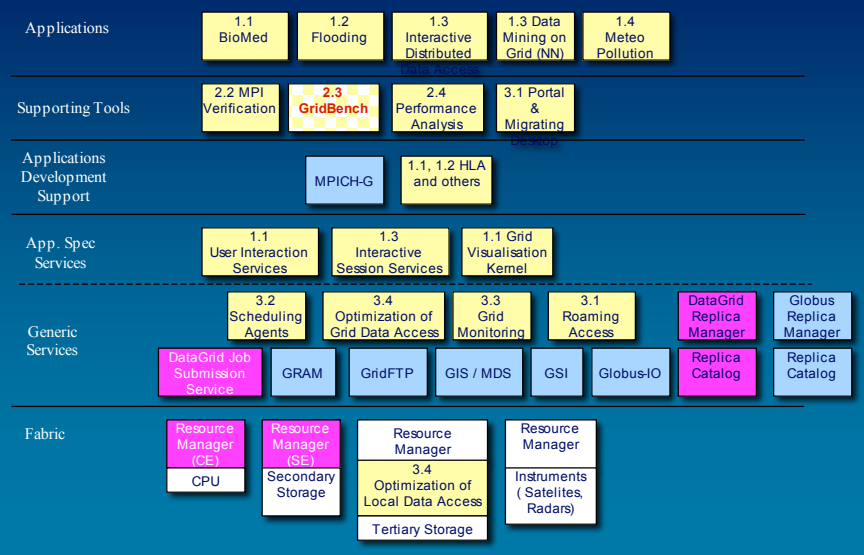

http://grid.ucy.ac.cy

GridBench front-end

- GBDL Translator
  - XML benchmark description (GBDL) to "job description language"
  - Supports JDL (EDG, Condor) and RSL (Globus).
- Benchmark Definition UI
  - GUI for defining and executing benchmarks.
- Benchmark Browser
  - GUI for browsing and analyzing results.
- Information Provider
  - Publishes results to Metacomputing Directory Service (MDS).

http://grid.ucy.ac.cy

GridBench back-end

- Orchestrator
  - Manager of execution and result collection.
  - Web Service (it must submit the job and wait for the output).
  - Co-located with work-load management client of CrossGrid.
- Archiver & Database
  - Stores benchmark definition, results, and monitoring.
  - Web Service.
  - Requires network connectivity to the host running the apache database.
- Benchmark Components
  - Benchmark executable code.
- Monitoring Component
  - Collects information using e.g. R-GMA or OCM-G.

GridBench at work

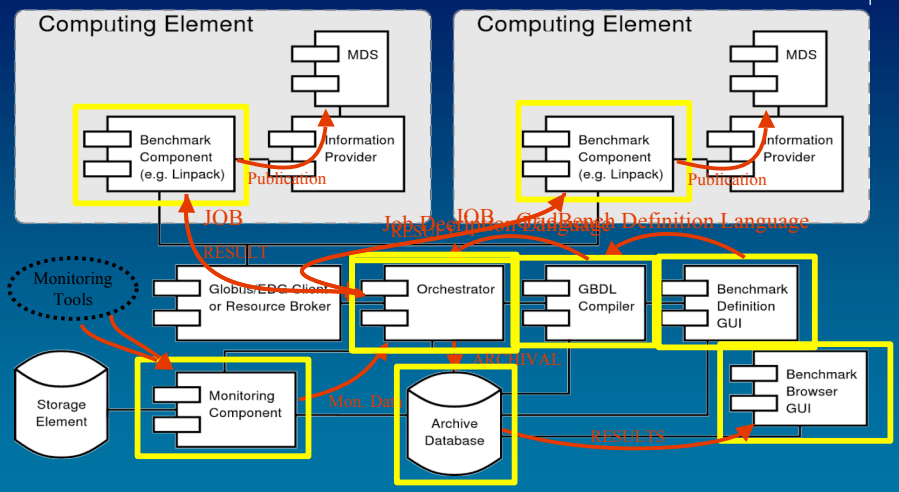

GridBench GUI

(Native XML database), collection of results

Outline

- A simple reference architecture for the Grid.
- The GridBench suite of benchmarks.
- GridBench description language.
- Filtering polluted measurements.
- The GridBench software.
- **Using GridBench to characterize Grid sites.**
- Conclusions and Future Work.

Characterizing Grid sites

- To characterize a site (cluster) and its resources, we are focusing on a small set of metrics, describing:
  - CPUs performance: OPS, FLOPS, INTOPS
  - Cache performance: MB/s
  - Main Memory performance: MB/s
  - Local Interconnect: latency, bandwidth
  - I/O performance: effective I/O bandwidth
  - Maximum Available Memory: available for dynamic allocation
- We are also looking at functionality aspects:
  - Are local queuing systems operating properly?
  - Is MPI installed properly and fully operable?
  - Is SSH working properly?

Metrics and Benchmarks

Factor	Metric	Delivered By
CPU	Operations per second (mixture of floating point and integer arithmetic)	EPWhetstone
CPU	Floating-Point operations per second	EPFlops
CPU	Integer operations per second	EPDhrystone
memory	sustainable memory bandwidth in MB/s (copy,add,multiply,triad)	EPStream
memory	Available physical memory in MB	EPMemsize
cache	memory bandwidth using different memory sizes in MB/s	CacheBench
Interconnect	latency, bandwidth and bisection bandwidth	MPPTest
I/O	Effective I/O bandwidth	b_eff_io

- We are using widely known micro-benchmarks to derive the required metrics.
- We run those mbenchs in parallel on all measured resources.

http://grid.ucy.ac.cy

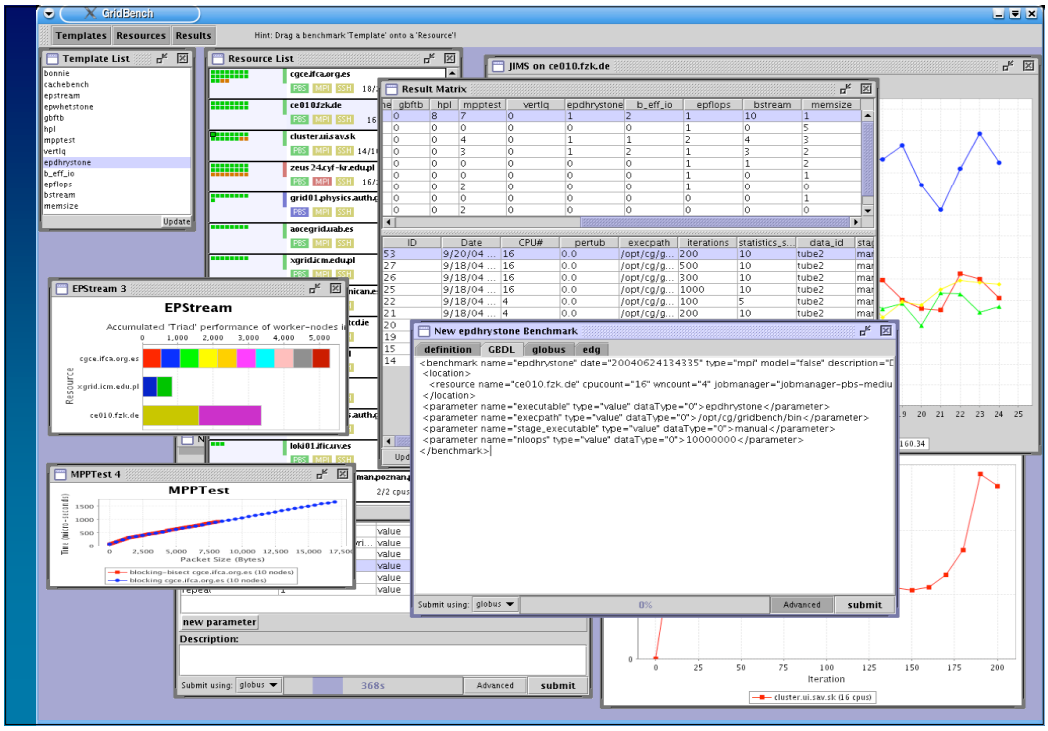

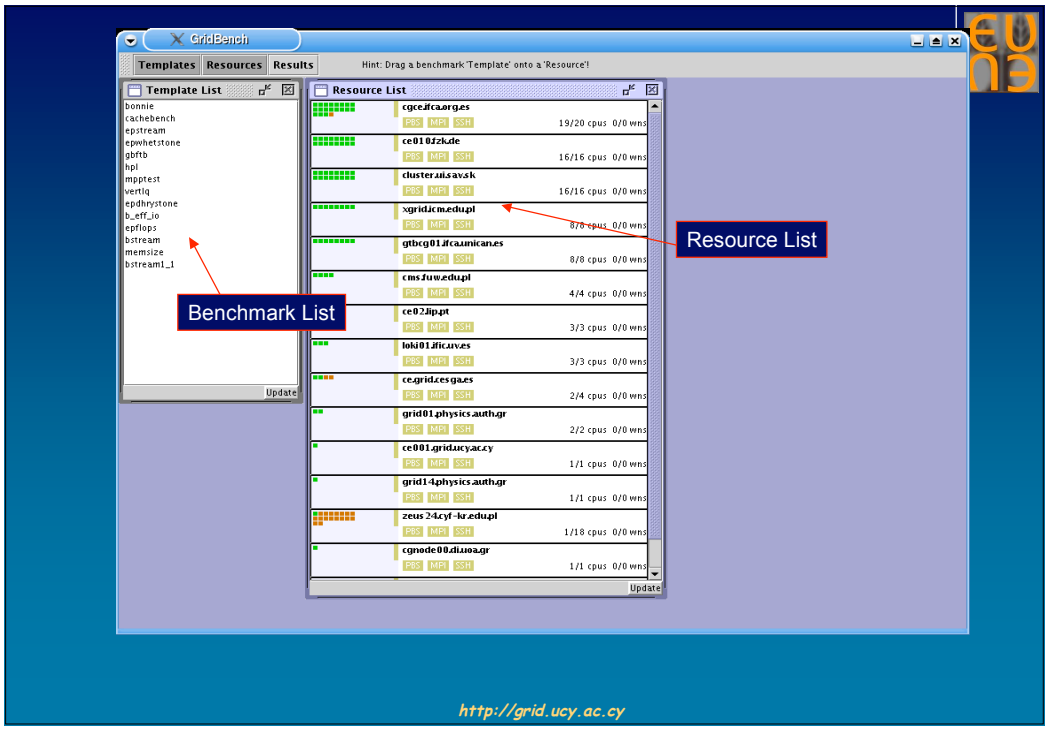
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
bonnie
cachebench
epstream
epwhetstone
gbftb
hpl
mpptest
vertlq
epdhrystone
b_eff_io
epflops
bstream
memsize
bstream1_1
Update
Resource List
cgce.ifca.org.es
ce010.fzk.de
cluster.ui.sav.sk
xgrid.icm.edu.pl
gtbcg01.ifca.unican.es
cms.fuw.edu.pl
ce02.lip.pt
loki01.ific.uv.es
ce.grid.cesga.es
grid01.physics.auth.gr
ce001.grid.ucy.ac.cy
grid14.physics.auth.gr
zeus 24.cyf-kr.edu.pl
cgnode00.di.uoa.gr
Benchmark List
Resource List
http://grid.ucy.ac.cy

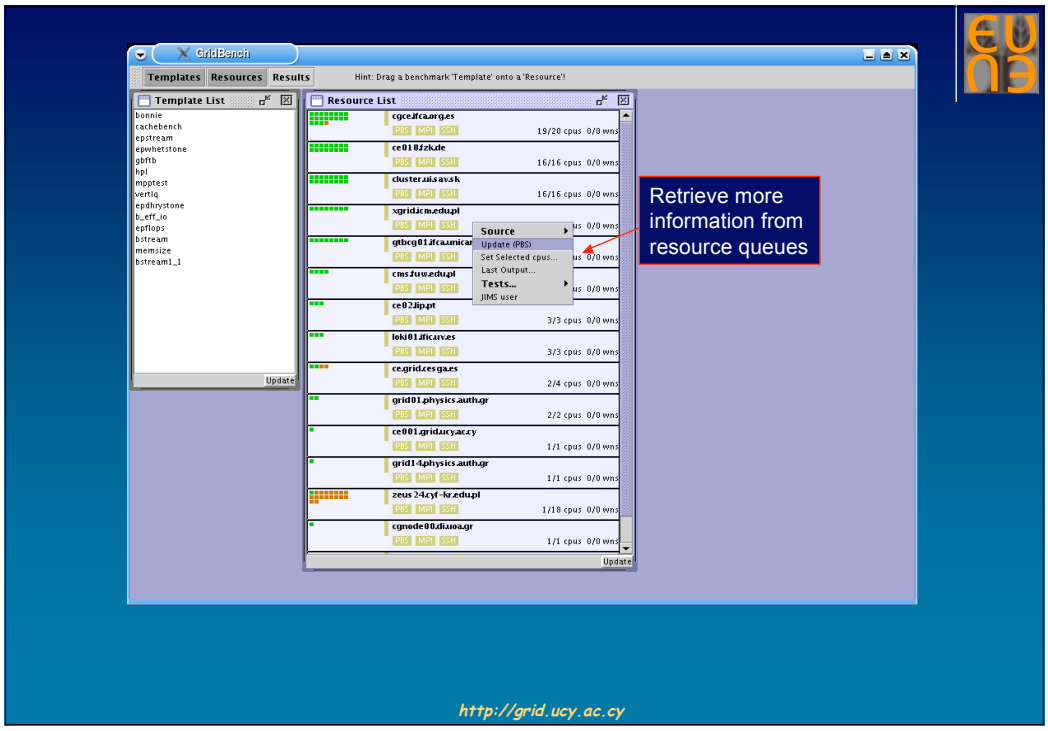
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
Source
Update (PBS)
Set Selected cpus...
Last Output...
Tests...
JIMS user
Retrieve more information from resource queues
http://grid.ucy.ac.cy

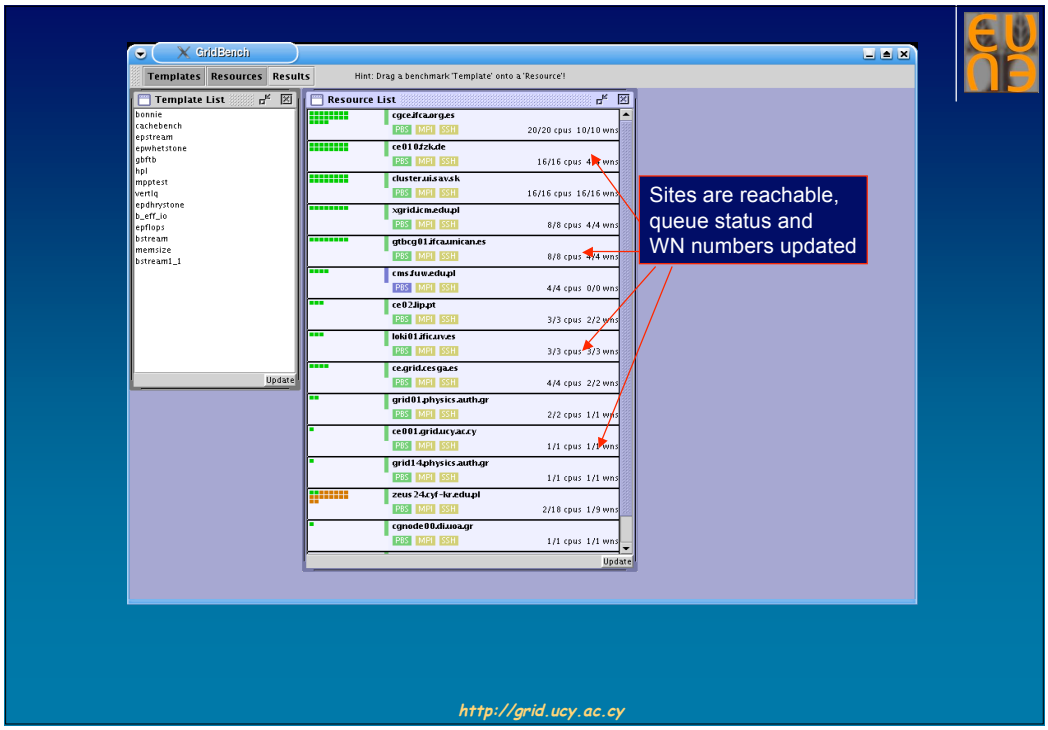
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
bonnie
cachebench
epstream
epwhetstone
gbftb
hpl
mpptest
vertlq
epdhrystone
b_eff_io
epflops
bstream
memsize
bstream1_1
Update
Resource List
cgce.ifca.org.es
ce01.fzk.de
cluster.ui.sav.sk
xgrid.icm.edu.pl
gtbcg01.ifca.unican.es
cms.fuw.edu.pl
ce02.lip.pt
loki01.ific.uv.es
ce.grid.cesga.es
grid01.physics.auth.gr
ce001.grid.ucy.ac.cy
grid14.physics.auth.gr
zeus.24.cyf-kr.edu.pl
cgnode00.di.uoa.gr
Update
Sites are reachable, queue status and WN numbers updated
http://grid.ucy.ac.cy

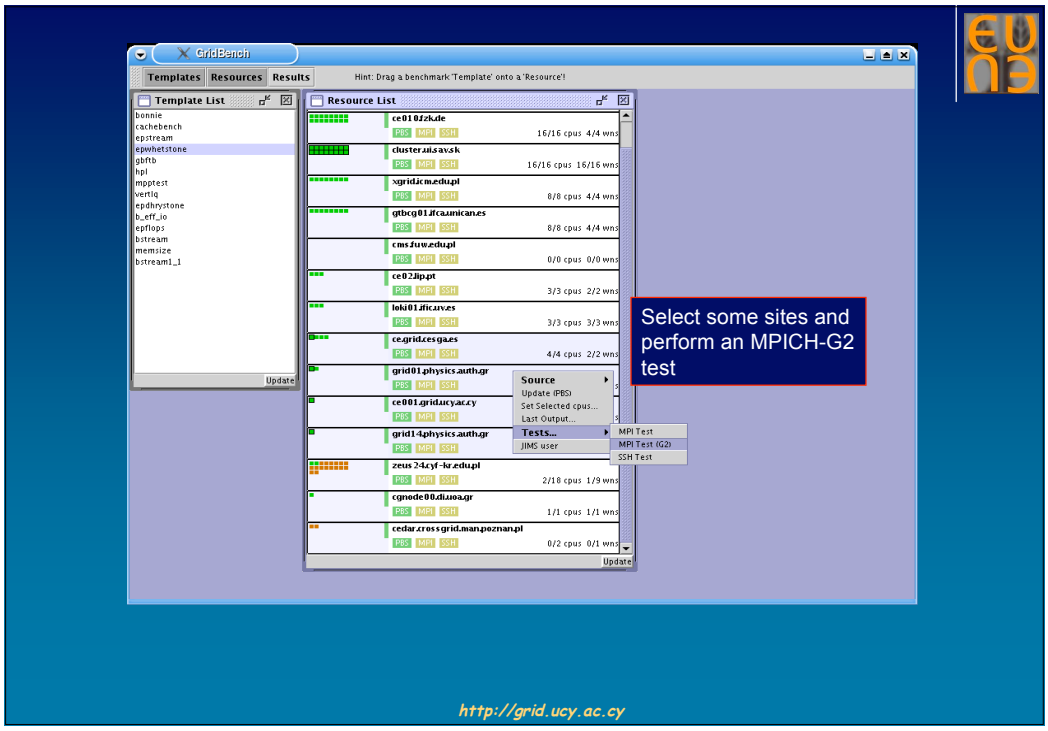
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
bonnie
cachebench
epstream
epwhetstone
gbftb
hpl
mpptest
vertlq
epdhrystone
b_eff_io
epflops
bstream
memsize
bstream1_1
Update
Resource List
ce01.fzk.de
cluster.ui.sav.sk
xgrid.icm.edu.pl
gtbcg01.ifca.unican.es
cms.fuw.edu.pl
ce02.lip.pt
loki01.ific.uv.es
ce.grid.cesga.es
grid01.physics.auth.gr
ce001.grid.ucy.ac.cy
grid14.physics.auth.gr
zeus.24.cyf-kr.edu.pl
cgnode00.di.uoa.gr
cedar.crossgrid.man.poznan.pl
Update
Source
Update (PBS)
Set Selected cpus...
Last Output...
Tests...
JIMS user
MPI Test
MPI Test (G2)
SSH Test
Select some sites and perform an MPICH-G2 test
http://grid.ucy.ac.cy

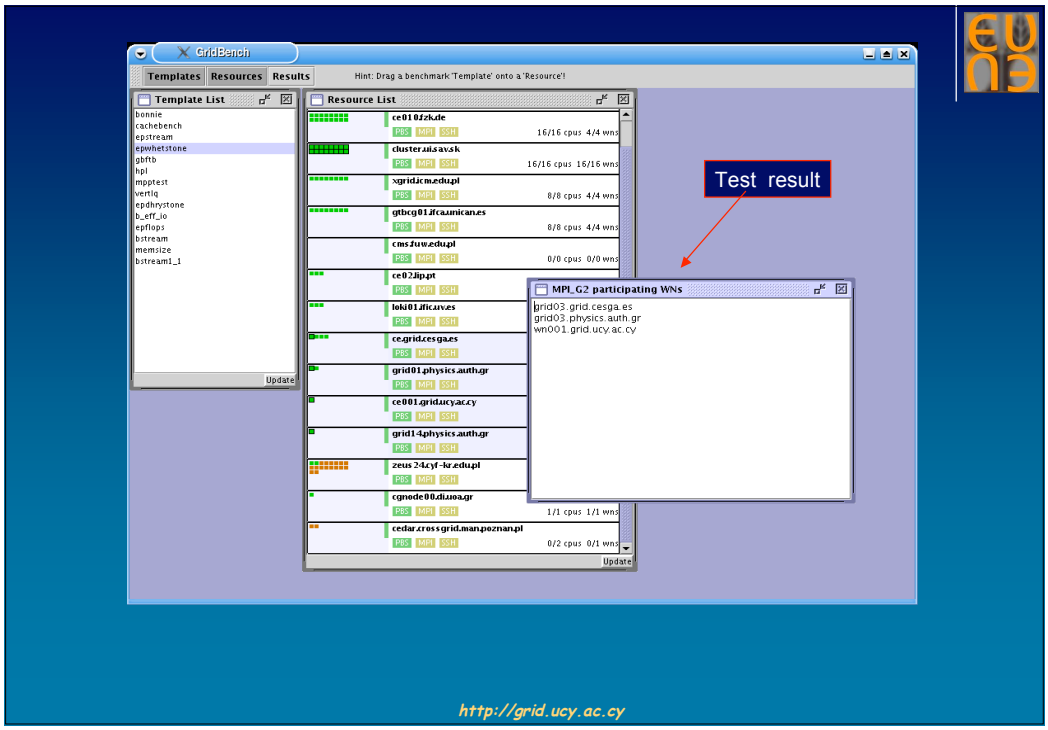
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
bonnie
cachebench
epstream
epwhetstone
gbftb
hpl
mpptest
vertlq
epdhrystone
b_eff_io
epflops
bstream
memsize
bstream1_1
Update
Resource List
ce010.fzk.de
cluster.uis.av.sk
xgrid.icm.edu.pl
gtbcg01.ifca.unican.es
cms.fuw.edu.pl
ce02.lip.pt
loki01.ific.uv.es
ce.grid.cesga.es
grid01.physics.auth.gr
ce001.grid.ucy.ac.cy
grid14.physics.auth.gr
zeus24.cyf-kr.edu.pl
cgnode00.di.uoa.gr
cedar.cross.grid.man.poznan.pl
Update
MPI_G2 participating WNs
grid03.grid.cesga.es
grid03.physics.auth.gr
wn001.grid.ucy.ac.cy
Test result
http://grid.ucy.ac.cy

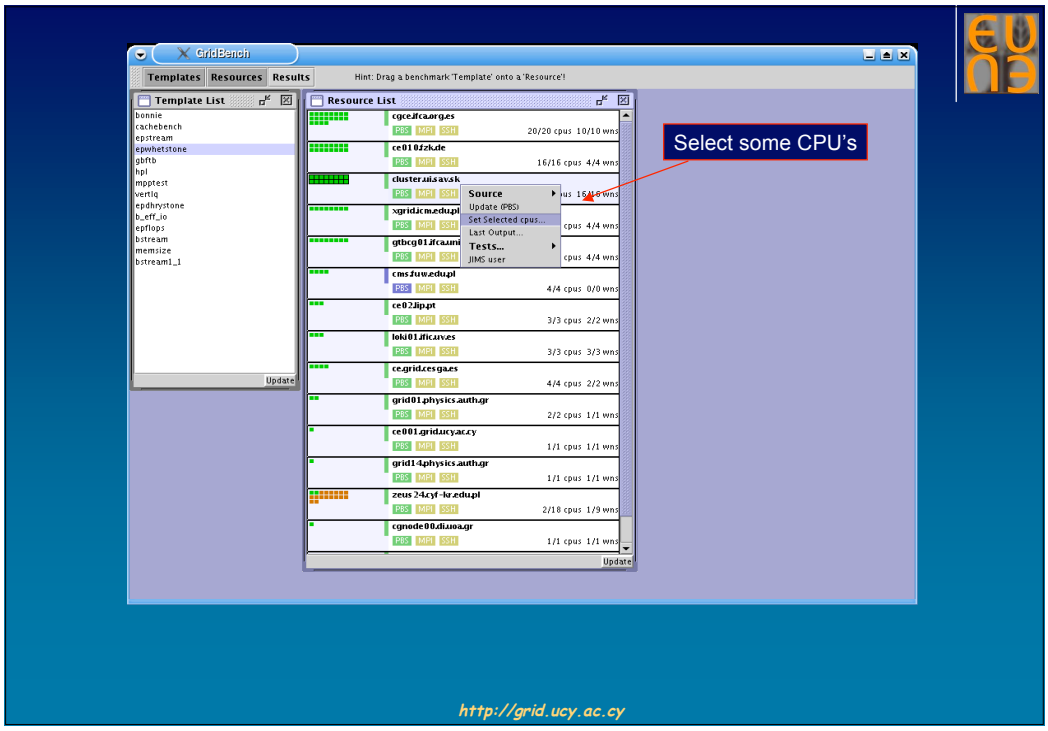
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
bonnie
cachebench
epstream
epwhetstone
gbftb
hpl
mpptest
vertlq
epdhrystone
b_eff_io
epflops
bstream
memsize
bstream1_1
Update
Resource List
cgce.ifca.org.es
ce010.fzk.de
cluster.uis.av.sk
xgrid.icm.edu.pl
gtbcg01.ifca.unican.es
cms.fuw.edu.pl
ce02.lip.pt
loki01.ific.uv.es
ce.grid.cesga.es
grid01.physics.auth.gr
ce001.grid.ucy.ac.cy
grid14.physics.auth.gr
zeus24.cyf-kr.edu.pl
cgnode00.di.uoa.gr
Source
Update (PBS)
Set Selected cpus...
Last Output...
Tests...
JIMS user
Update
Select some CPU's
http://grid.ucy.ac.cy

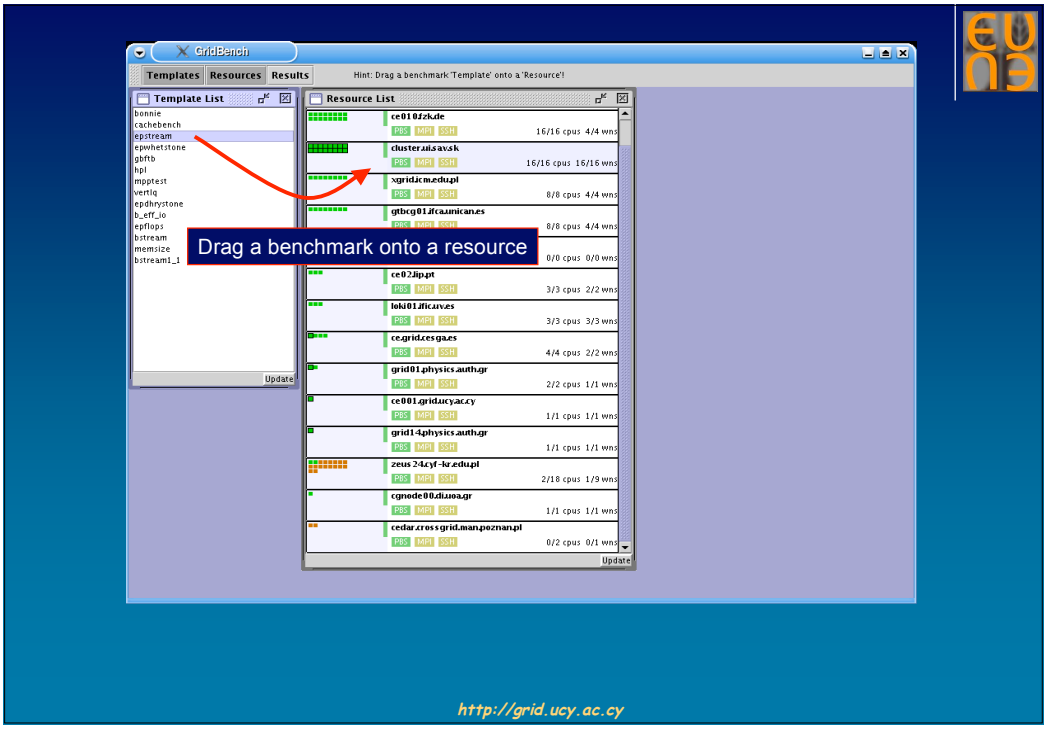
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
bonnie
cachebench
epstream
epwhetstone
gbftb
hpl
mpptest
vertlq
epdhrystone
b_eff_io
epflops
bstream
memsize
bstream1_1
Update
Resource List
Drag a benchmark onto a resource
http://grid.ucy.ac.cy

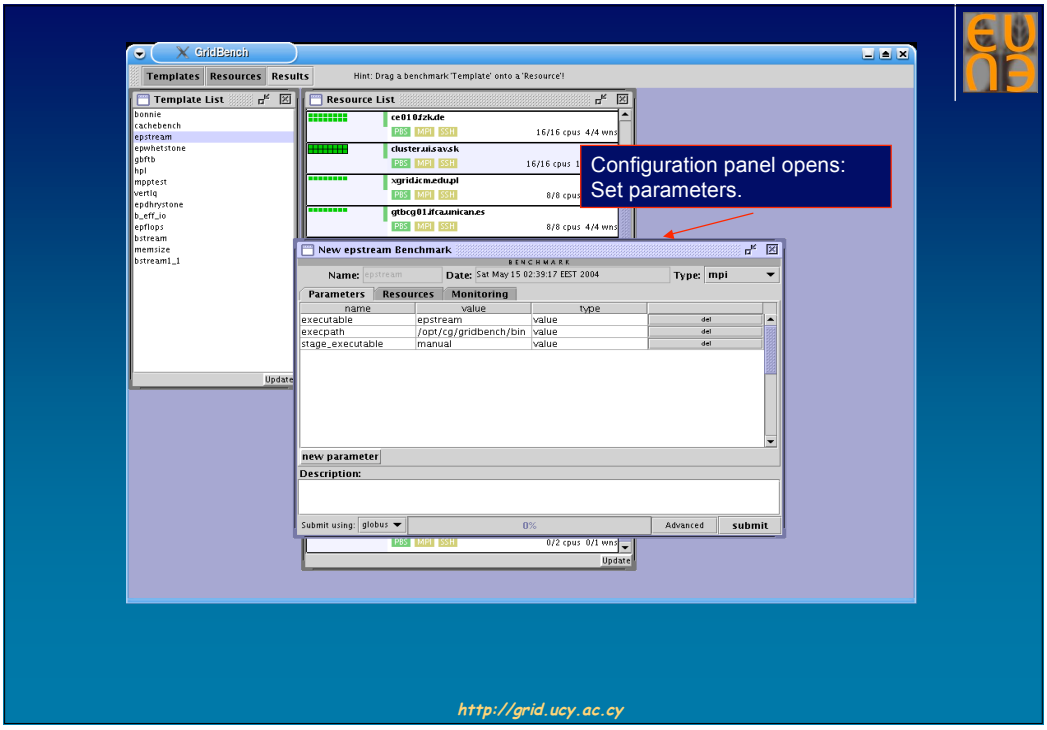
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
Configuration panel opens:
Set parameters.
New epstream Benchmark
BENCHMARK
Name:
Date: Sat May 15 02:39:17 EEST 2004
Type: mpi
Parameters
Resources
Monitoring
name
value
type
executable
epstream
value
execpath
/opt/cg/gridbench/bin
value
stage_executable
manual
value
new parameter
Description:
Submit using:
globus
0%
Advanced
submit
http://grid.ucy.ac.cy

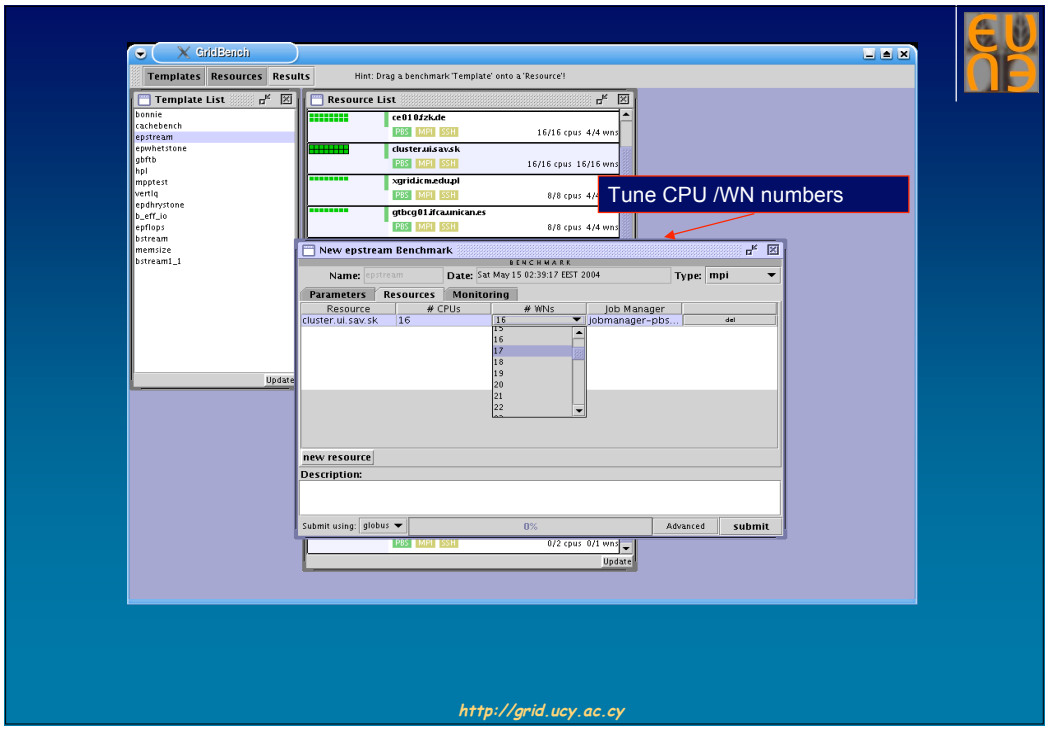
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
Tune CPU /WN numbers
New epstream Benchmark
Name:
Date:
Type:
mpi
Parameters
Resources
Monitoring
Resource
CPUs
WNs
Job Manager
new resource
Description:
Submit using:
globus
Advanced
submit
http://grid.ucy.ac.cy

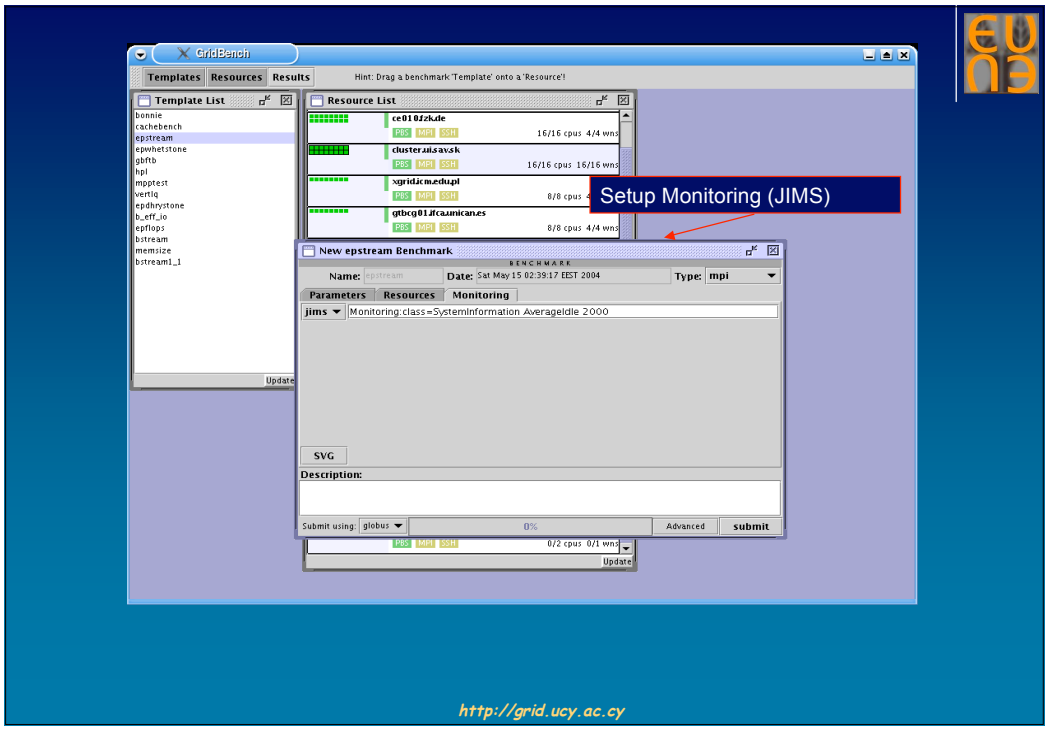
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
Setup Monitoring (JIMS)
New epstream Benchmark
Name:
Date:
Type:
mpi
Parameters
Resources
Monitoring
jims
Monitoring:class=SystemInformation AverageIdle 2000
SVG
Description:
Submit using:
globus
Advanced
submit
http://grid.ucy.ac.cy

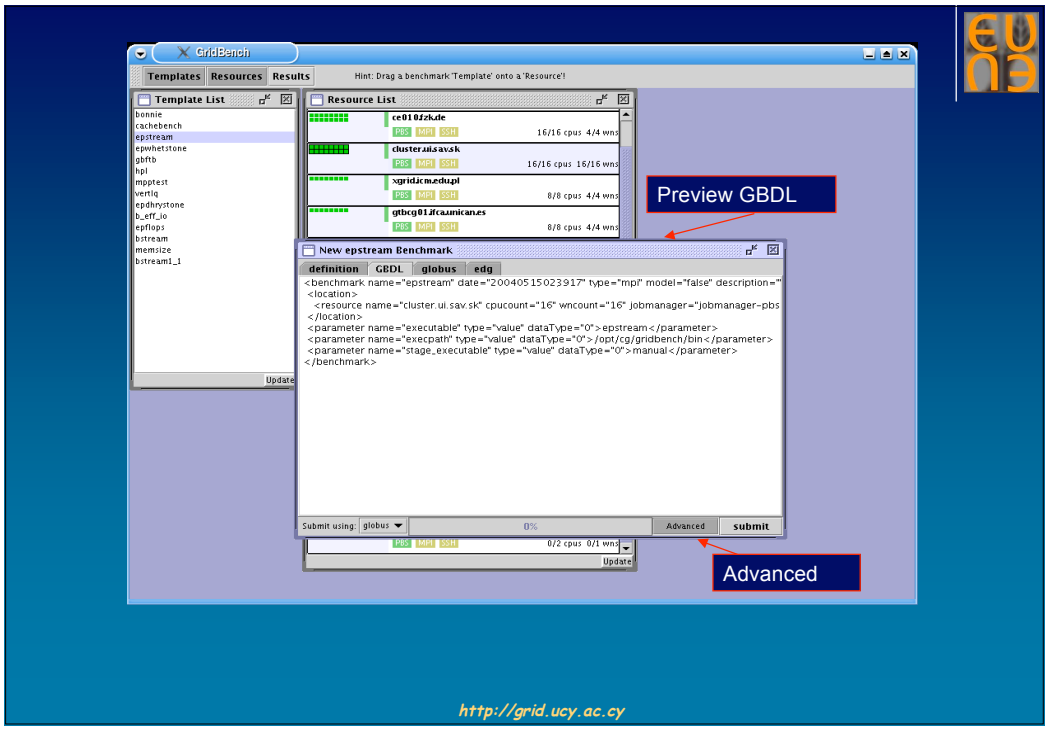
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
bonnie
cachebench
epstream
epwhetstone
gbftb
hpl
mpptest
vertlq
epdhrystone
b_eff_io
epflops
bstream
memsize
bstream1_1
Update
Resource List
ce01.fzk.de
cluster.ui.sav.sk
16/16 cpus 16/16 wns
xgrid.icm.edu.pl
8/8 cpus 4/4 wns
gtbcg01.ifca.unican.es
0/2 cpus 0/1 wns
Preview GBDL
New epstream Benchmark
definition
GBDL
globus
edg
<benchmark name="epstream" date="20040515023917" type="mpi" model="false" description="
<location>
<resource name="cluster.ui.sav.sk" cpucount="16" wncount="16" jobmanager="jobmanager-pbs
</location>
<parameter name="executable" type="value" dataType="0">epstream</parameter>
<parameter name="execpath" type="value" dataType="0">/opt/cg/gridbench/bin</parameter>
<parameter name="stage_executable" type="value" dataType="0">manual</parameter>
</benchmark>
Submit using:
globus
0%
Advanced
submit
Advanced
http://grid.ucy.ac.cy

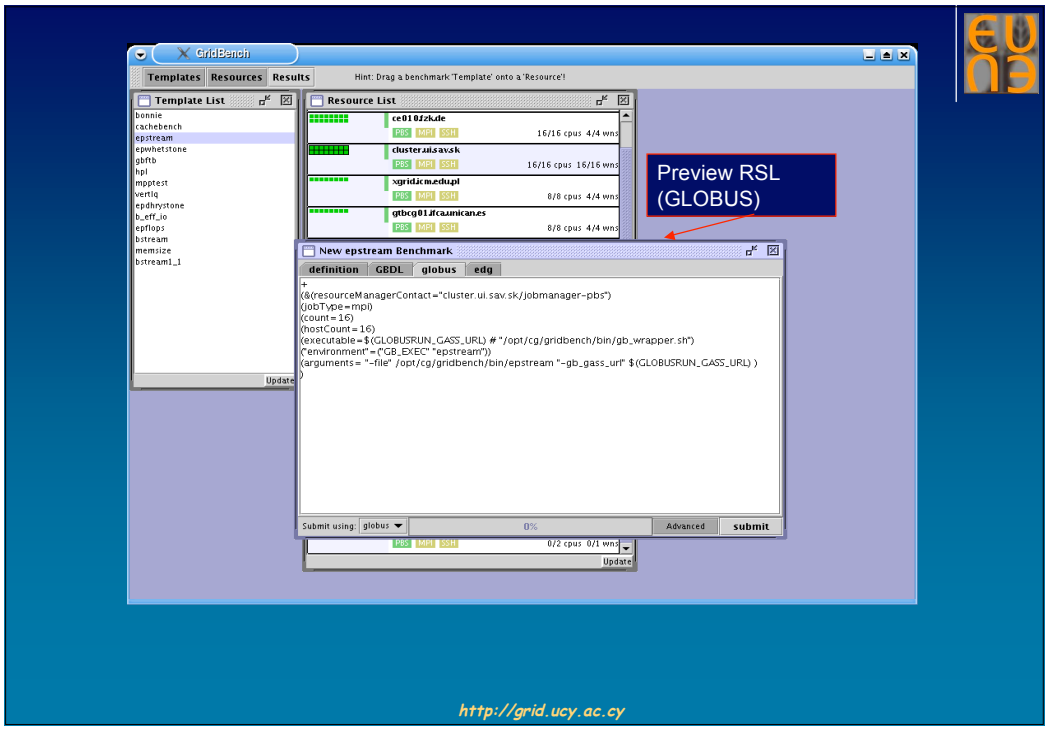
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
bonnie
cachebench
epstream
epwhetstone
gbftb
hpl
mpptest
vertlq
epdhrystone
b_eff_io
epflops
bstream
memsize
bstream1_1
Update
Resource List
ce01.fzk.de
cluster.ui.sav.sk
16/16 cpus 16/16 wns
xgrid.icm.edu.pl
8/8 cpus 4/4 wns
gtbcg01.ifca.unican.es
0/2 cpus 0/1 wns
Preview RSL (GLOBUS)
New epstream Benchmark
definition
GBDL
globus
edg
+
(&(resourceManagerContact="cluster.ui.sav.sk/jobmanager-pbs")
(jobType=mpi)
(count=16)
(hostCount=16)
(executable=$(GLOBUSRUN_GASS_URL) # "/opt/cg/gridbench/bin/gb_wrapper.sh")
("environment"=("GB_EXEC" "epstream"))
(arguments= "-file" /opt/cg/gridbench/bin/epstream "-gb_gass_url" $(GLOBUSRUN_GASS_URL))
)
Submit using:
globus
0%
Advanced
submit
http://grid.ucy.ac.cy

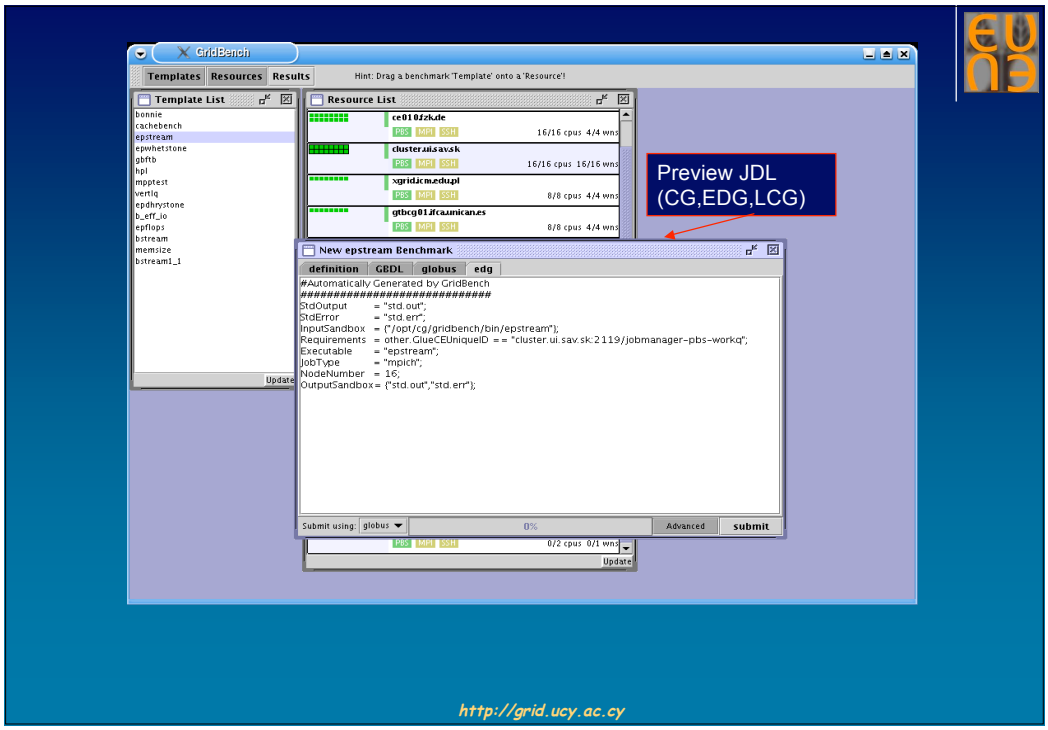
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
Preview JDL
(CG,EDG,LCG)
New epstream Benchmark
definition
GBDL
globus
edg
#Automatically Generated by GridBench
StdOutput = "std.out";
StdError = "std.err";
InputSandbox = {"/opt/cg/gridbench/bin/epstream"};
Requirements = other.GlueCEUniqueID == "cluster.ui.sav.sk:2119/jobmanager-pbs-workq";
Executable = "epstream";
JobType = "mpich";
NodeNumber = 16;
OutputSandbox = {"std.out","std.err"};
Submit using:
globus
Advanced
submit
http://grid.ucy.ac.cy

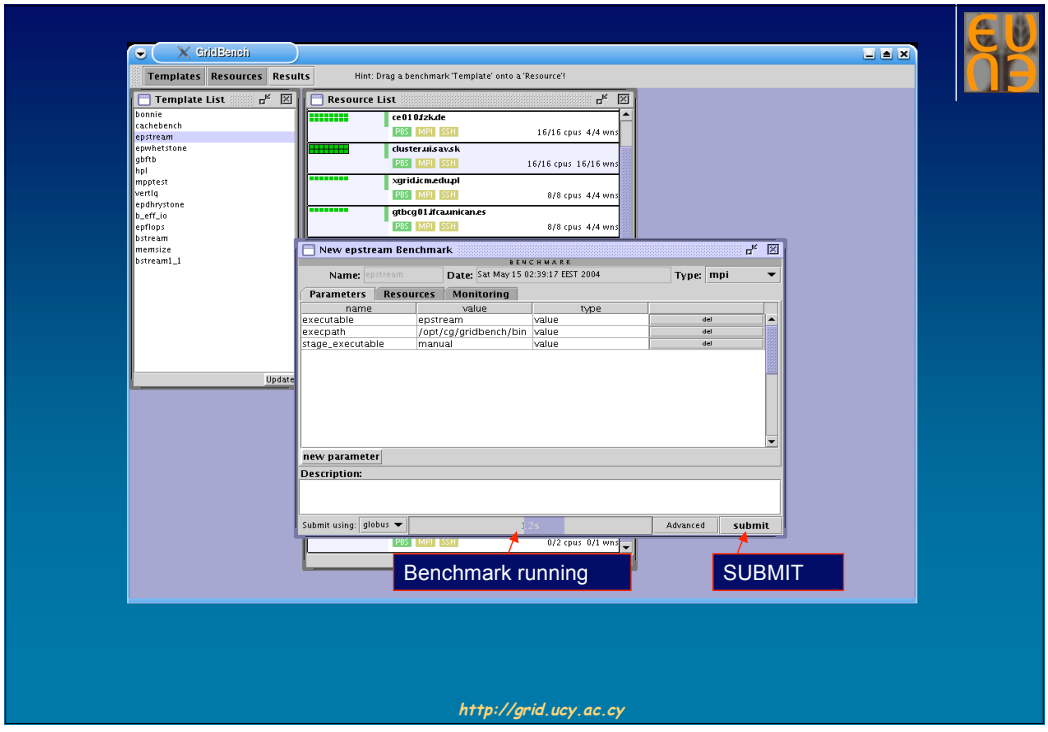
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
New epstream Benchmark
BENCHMARK
Name:
Date: Sat May 15 02:39:17 EEST 2004
Type: mpi
Parameters
Resources
Monitoring
name
value
type
executable
epstream
value
execpath
/opt/cg/gridbench/bin
value
stage_executable
manual
value
new parameter
Description:
Submit using:
globus
Advanced
submit
Benchmark running
SUBMIT
http://grid.ucy.ac.cy

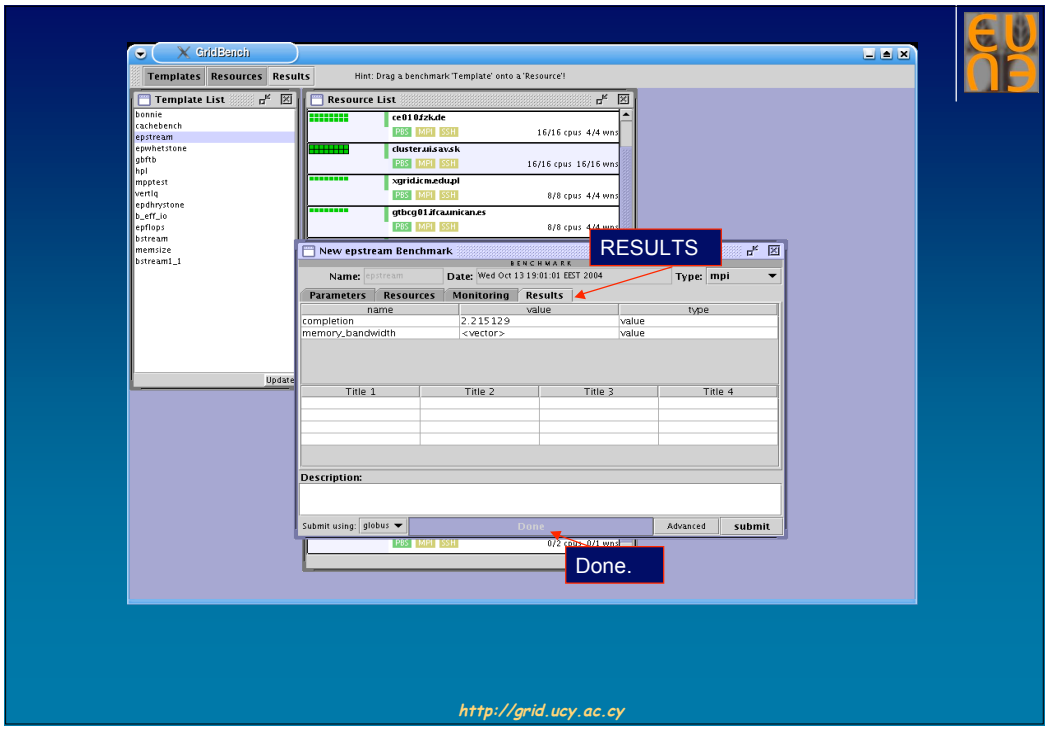
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
New epstream Benchmark
RESULTS
Done.
http://grid.ucy.ac.cy

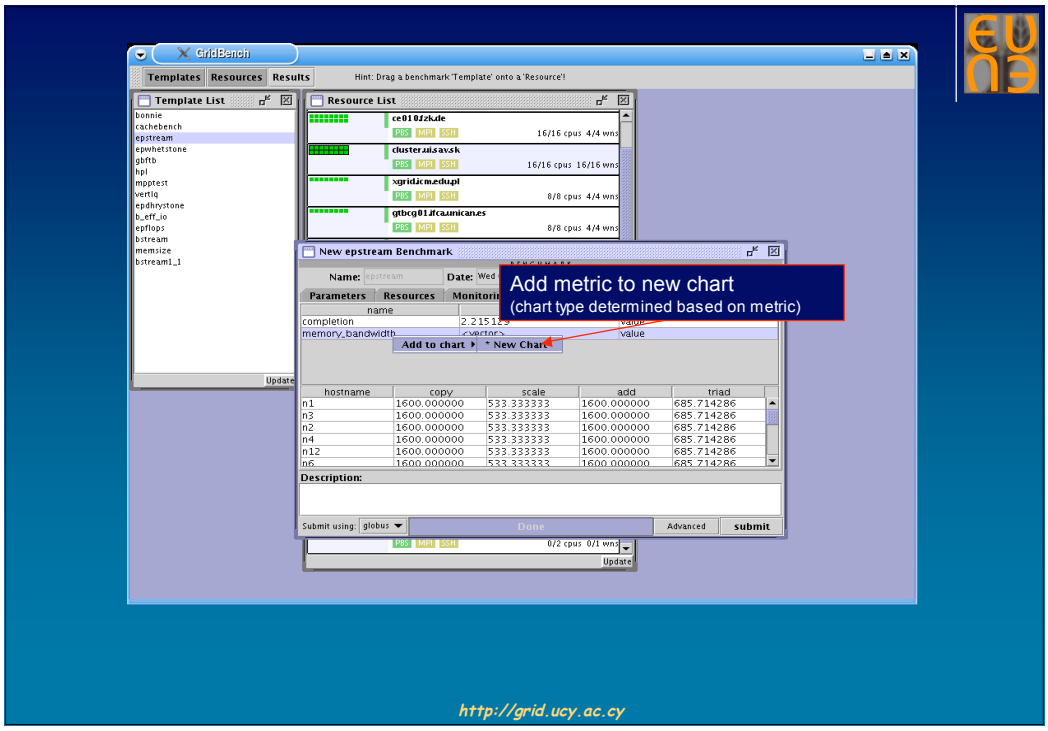
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
New epstream Benchmark
Add metric to new chart
(chart type determined based on metric)
http://grid.ucy.ac.cy

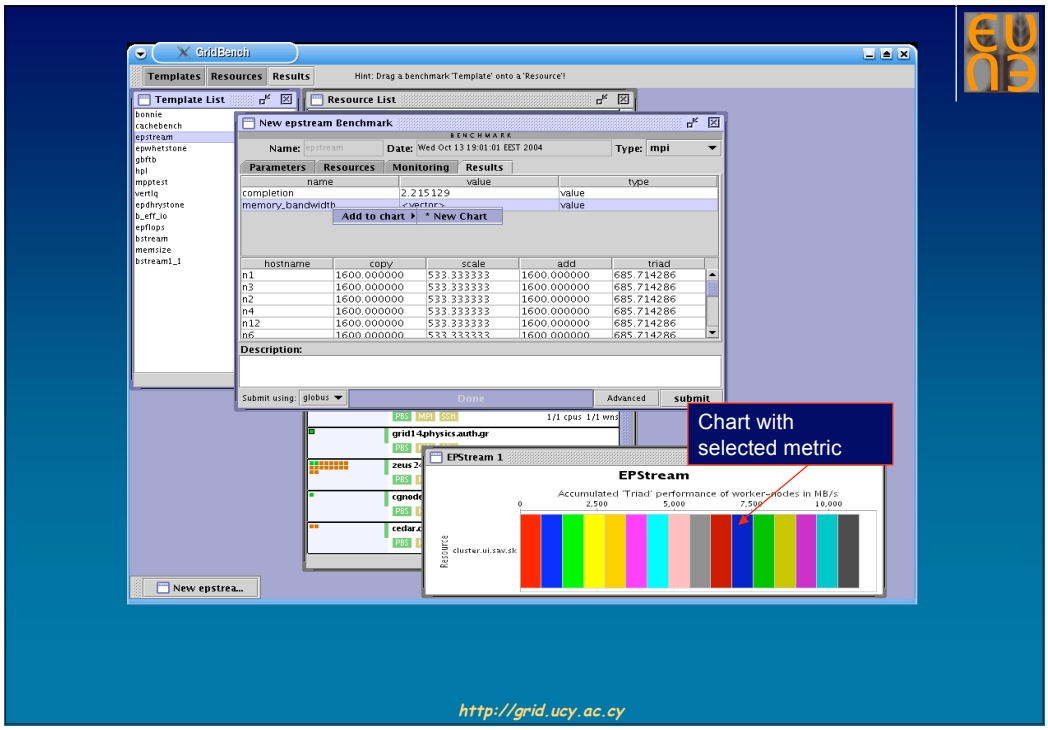
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
New epstream Benchmark
Name:
Date: Wed Oct 13 19:01:01 EEST 2004
Type: mpi
Parameters
Resources
Monitoring
Results
name
value
type
completion
2.215129
memory_bandwidth
Add to chart
* New Chart
hostname
copy
scale
add
triad
Description:
Submit using: globus
Done
Advanced
submit
Chart with selected metric
EPStream 1
EPStream
Accumulated 'Triad' performance of worker-nodes in MB/s
Resource
http://grid.ucy.ac.cy

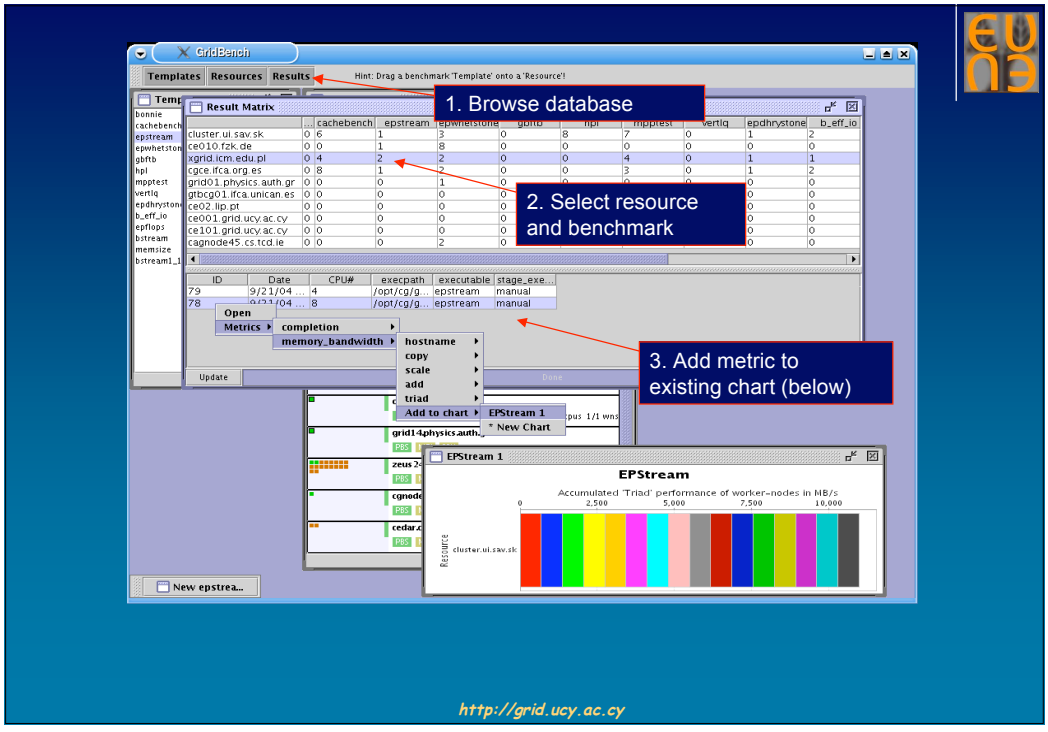
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
1. Browse database
Result Matrix
2. Select resource and benchmark
Open
Metrics
completion
memory_bandwidth
hostname
copy
scale
add
triad
Add to chart
EPStream 1
* New Chart
3. Add metric to existing chart (below)
EPStream 1
EPStream
Accumulated 'Triad' performance of worker-nodes in MB/s
Resource
http://grid.ucy.ac.cy

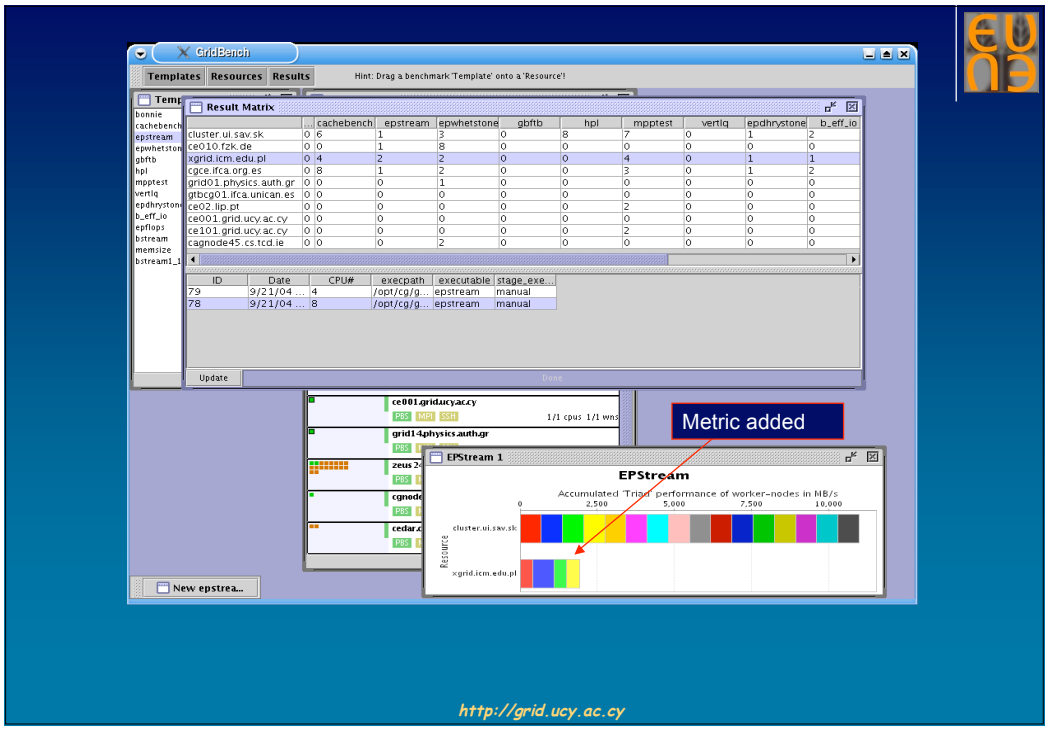
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Result Matrix
Update
Metric added
EPStream 1
EPStream
Accumulated Triad performance of worker-nodes in MB/s
New epstrea...
http://grid.ucy.ac.cy

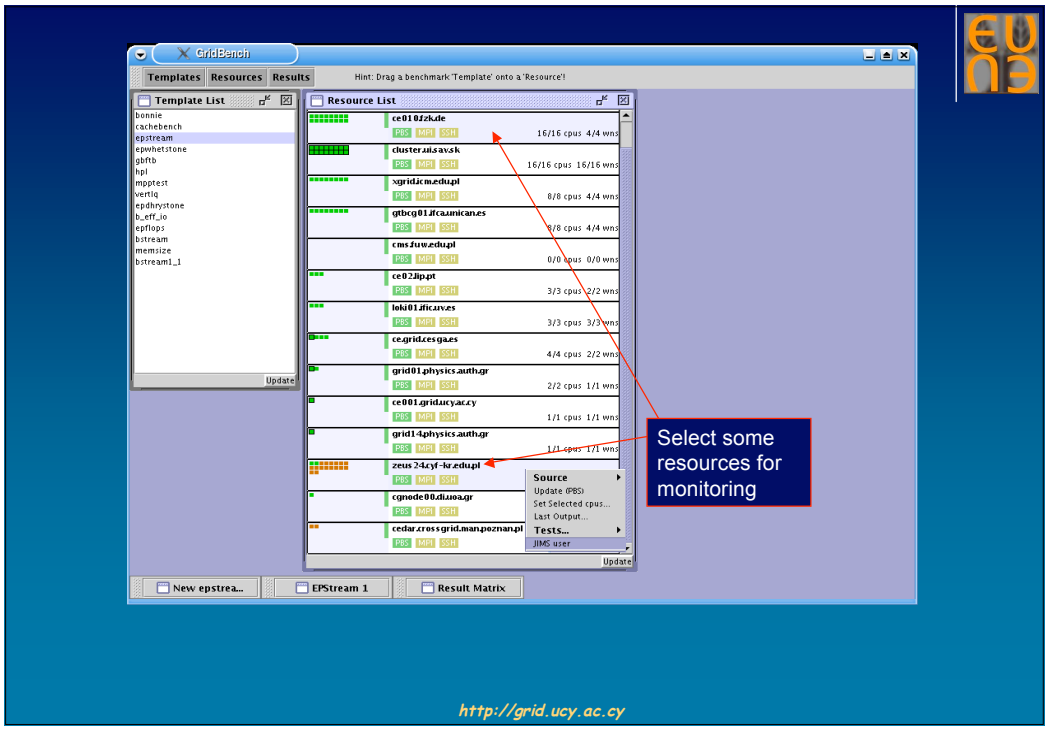
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
Update
Source
Update (PBS)
Set Selected cpus...
Last Output...
Tests...
JIMS user
Select some resources for monitoring
New epstrea...
EPStream 1
Result Matrix
http://grid.ucy.ac.cy

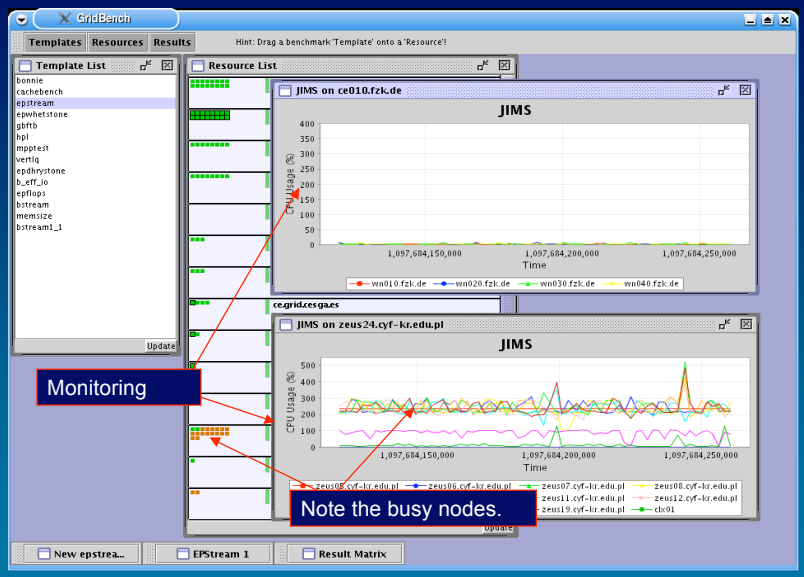

http://grid.ucy.ac.cy

GridBench results: CPU performance

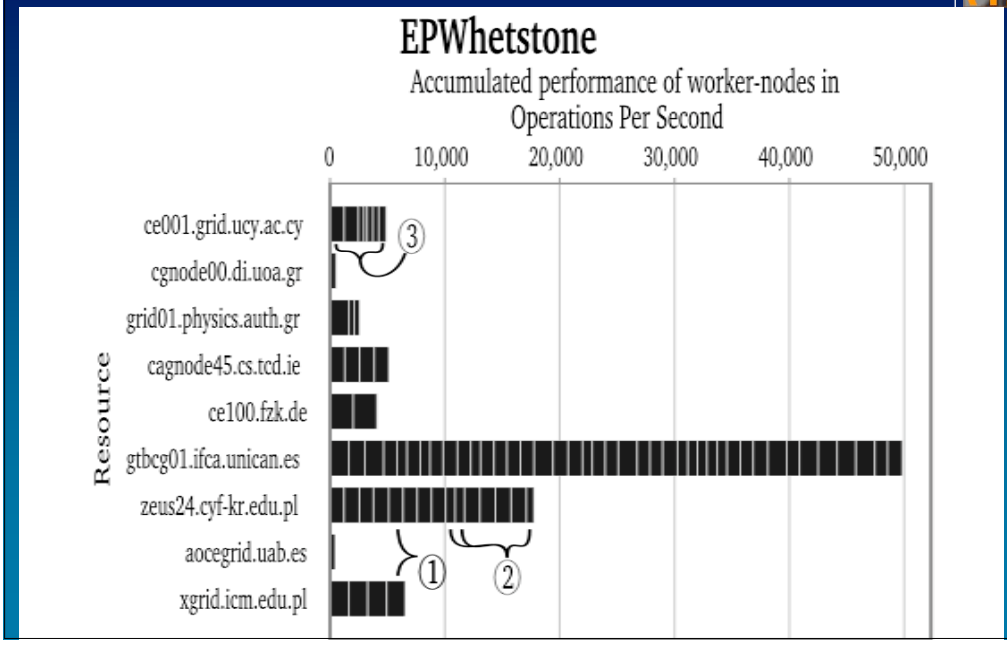

GridBench results: Memory performance

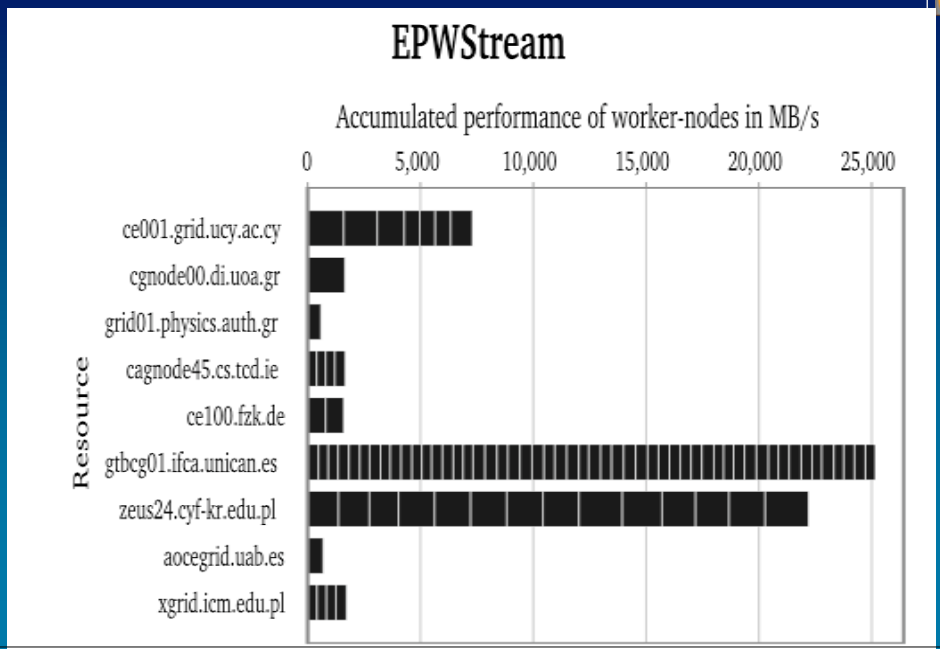

GridBench results: Available Memory

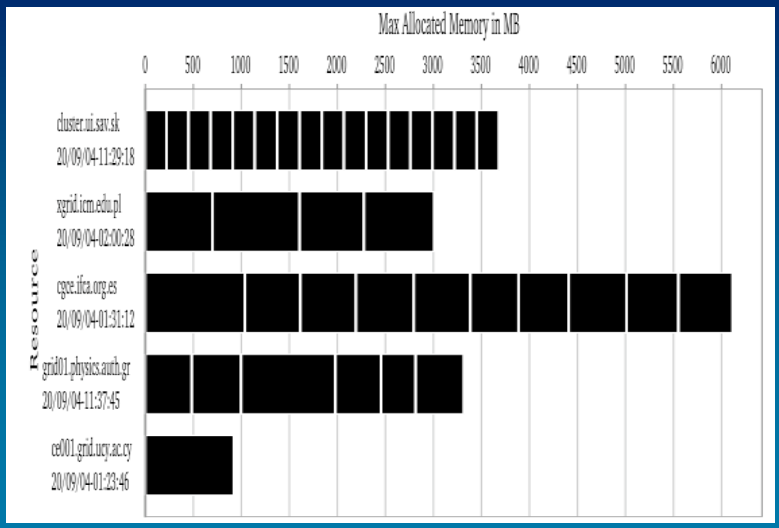

GridBench results: Communication

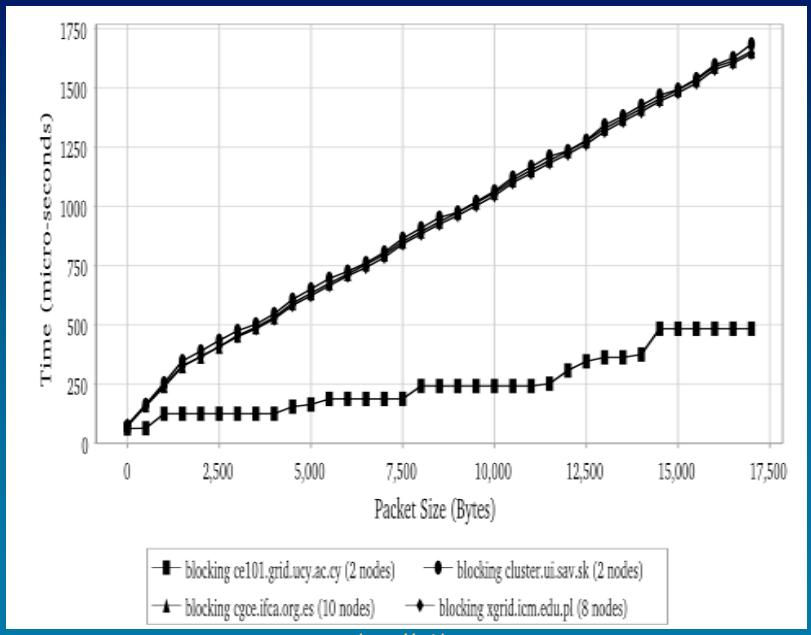

GridBench results: cache

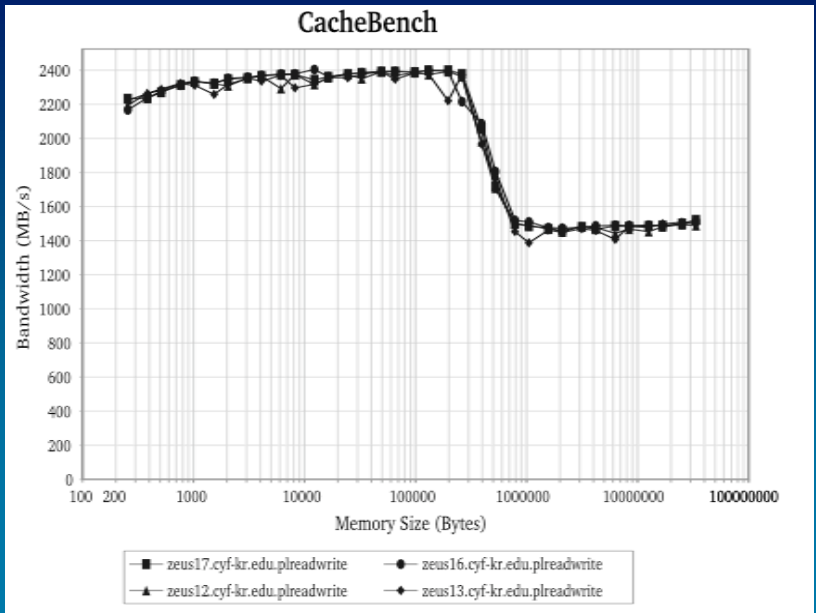

Resource Selection with GridBench

- Where should I run my fluid-flow code for Surgical Planning?

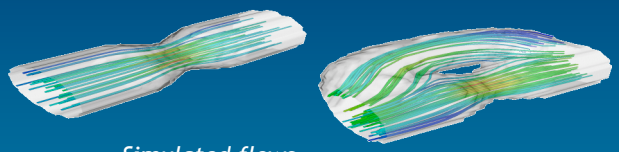

Simulated flows

http://grid.ucy.ac.cy

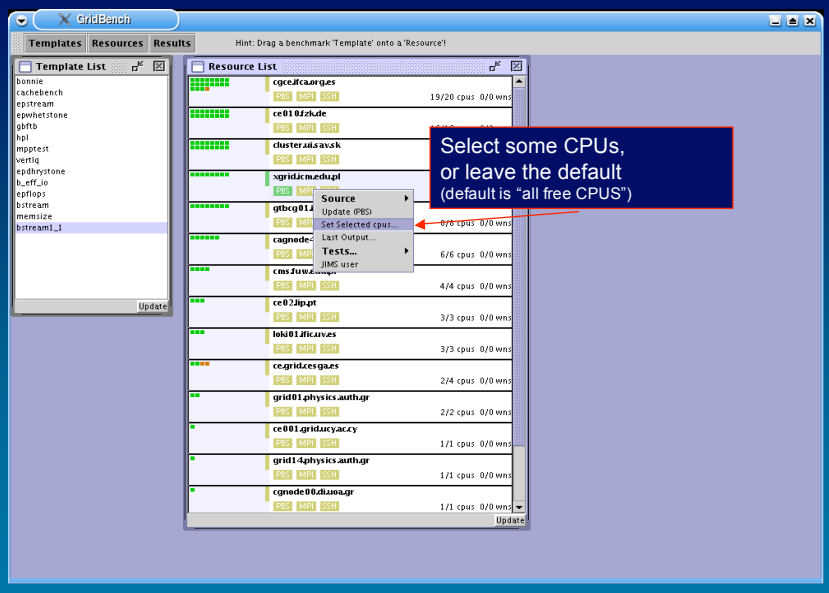

http://grid.ucy.ac.cy

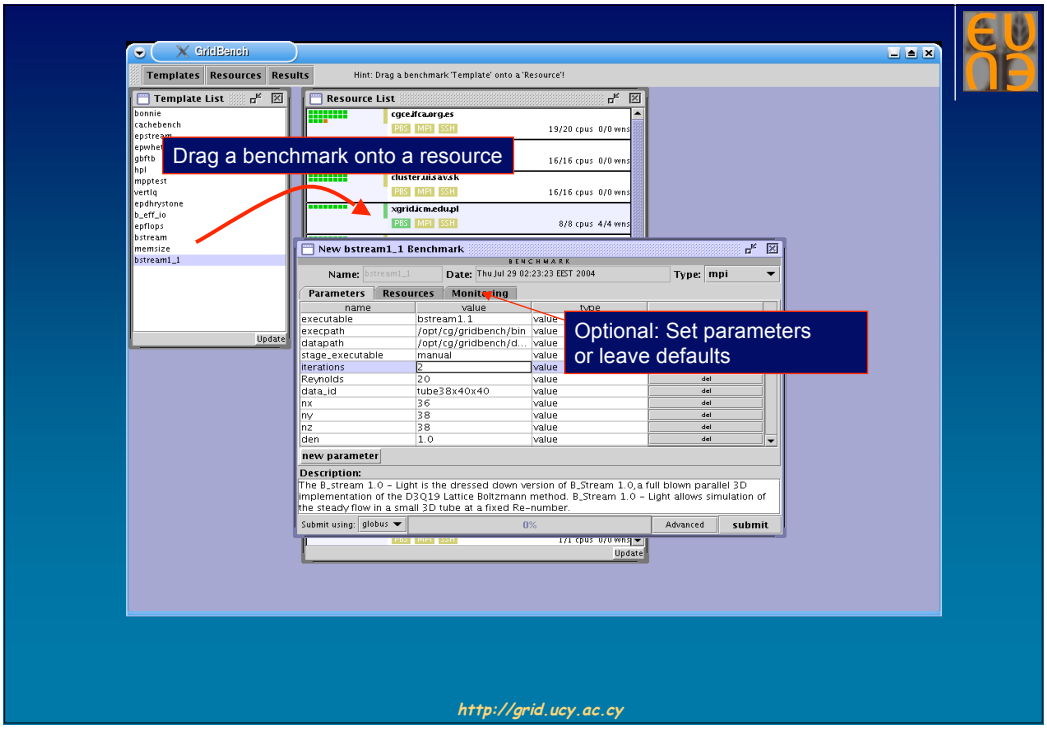
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
Drag a benchmark onto a resource
New bstream1_1 Benchmark
Optional: Set parameters or leave defaults
http://grid.ucy.ac.cy

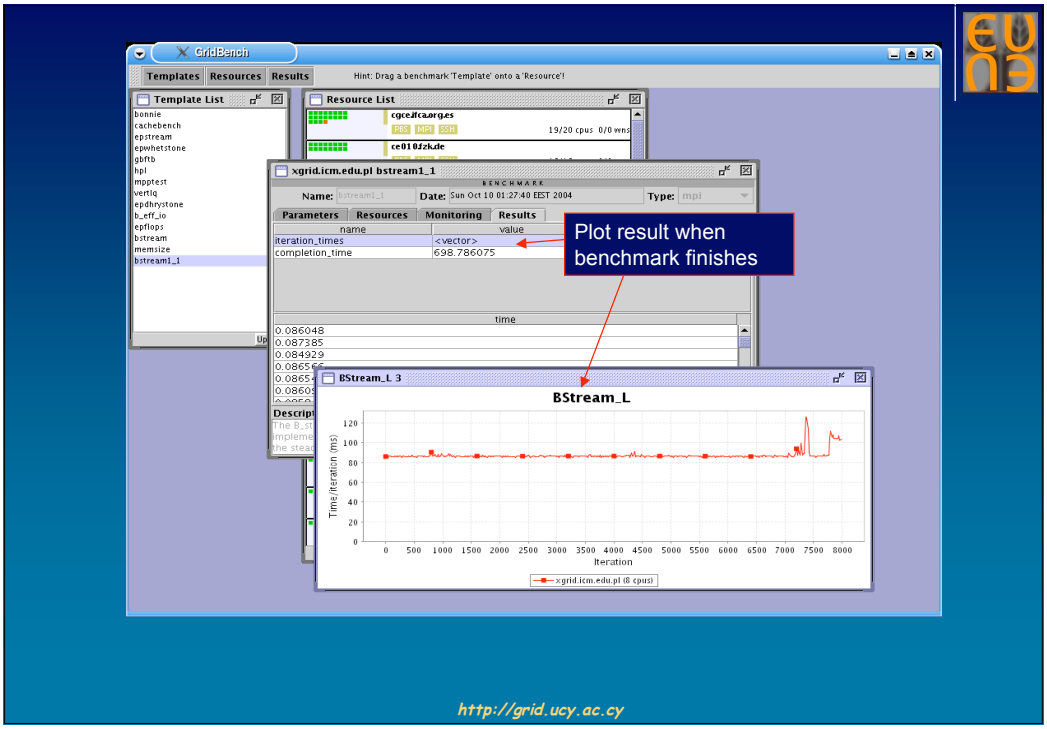
GridBench
Templates
Resources
Results
Hint: Drag a benchmark 'Template' onto a 'Resource'!
Template List
Resource List
xgrid.icm.edu.pl bstream1_1
Plot result when benchmark finishes
BStream_L 3
BStream_L
Time/iteration (ms)
Iteration
http://grid.ucy.ac.cy

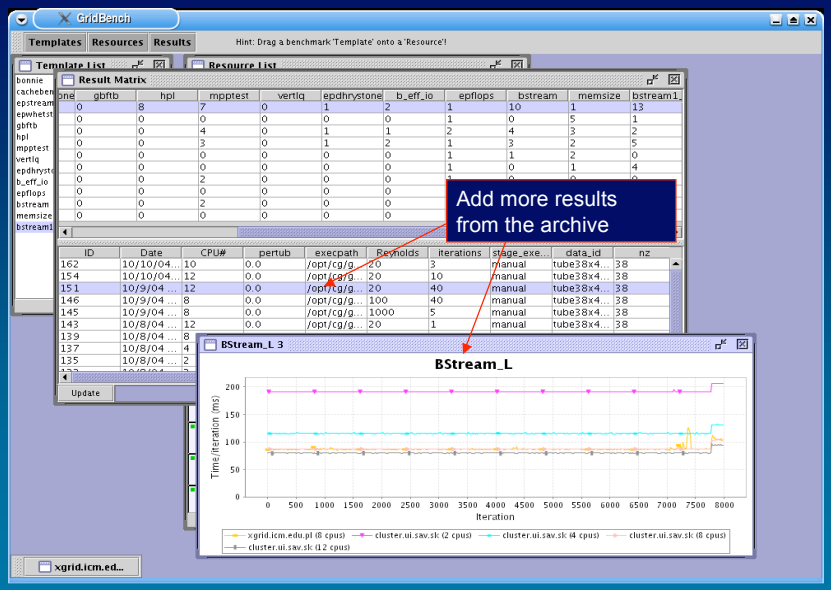

Conclusions

- Virtualization and resource heterogeneity turn Grid Benchmarking into a:
  - Challenging and expensive process.
  - Necessary undertaking for performance-based decisions.
- Isolated metrics are of little use. We need instead, structured sets of metrics:
  - Describing collectively the performance capacity of an abstract representation of a Grid infrastructure.
  - Amenable to statistical distillation, to derive higher-level, qualitative metrics.
  - Whose storage, organization, visualization, and interpretation raises several difficulties.
- Virtualization and the lack of central control, put the accuracy of benchmarking measurements to question. The combination of metrics with monitoring information is required to filter-out invalid measurements.

Conclusions

- The size, the openness, and the complexity of the Grid, make it susceptible to a variety of frequent, partial failures. Thus:
  - Functionality benchmarking is equally important from an end-user's perspective.
  - Benchmarking a mechanism for driving automatic remote healing.
- The dynamic nature of Grid infrastructures necessitate a periodic "refreshment" of performance metrics. Thus:
  - Grid benchmarking can be established as an automated central Grid service running with special privileges.
  - Benchmarks can be used as a quick, "end-to-end" test of a Grid's "health."
  - Benchmarks can be used for the automatic auditing of resource providers by a VO administration: compliance to policies, reliability of information services, etc.

Current and Future Work

- Extending GBDL and the GBDL translator to support:
  - The description of configuration-parameter selection constraints and guidelines.
  - The automatic selection of configuration parameters.
  - Interoperability with UNICORE middleware.
- Expanding the GridBench suite.
- Automating the process of metrics filtering and decision support.
- Using GridBench-metrics for resource selection and brokerage.
- Deriving higher-level metrics to express "quality features" of Grid infrastructures: Homogeneity, Trustworthiness of GIS, Health of the infrastructure, Reliability & Robustness.

ΤΕΛΟΣ

http://grid.ucy.ac.cy

Questions?

GridBench

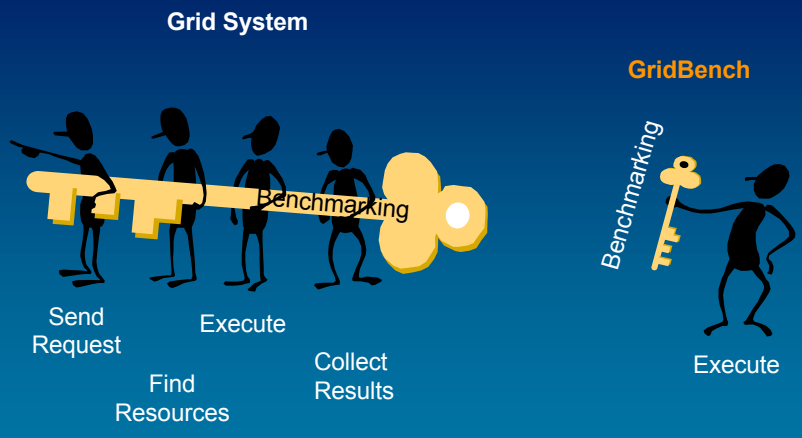

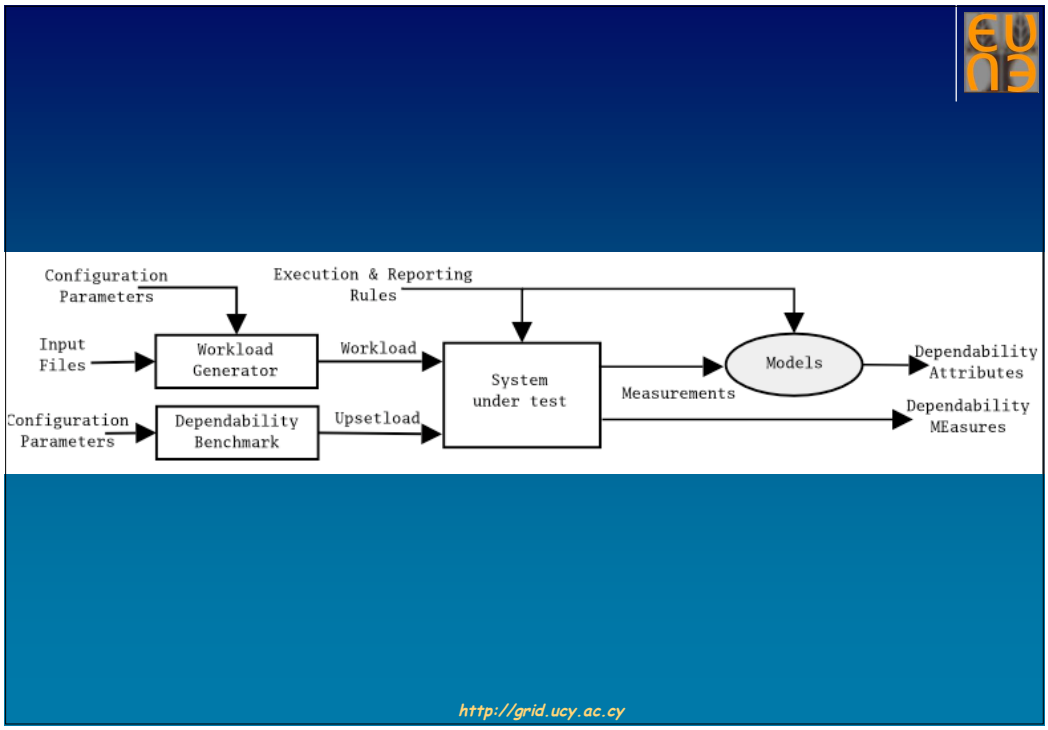
Configuration Parameters
Execution & Reporting Rules
Input Files
Workload Generator
Workload
System under test
Models
Dependability Attributes
Measurements
Configuration Parameters
Dependability Benchmark
Upsetload
Dependability MEasures
http://grid.ucy.ac.cy