



GAMS/ETSAP Liaison Officer R&D Activities

ETSAP Runtime GAMS License and TIMES-Cloud Platform

Gary Goldstein / Michael Bussieck & Alex Meeraus / Amit Kanudia
DecisionWare, Inc. / GAMS Development Corp / KanORS-EMR

IEA-ETSAP Annex XII 4th Workshop
Lisbon, Portugal
December 10-11, 2012

ETSAP Runtime GAMS License (1/2)

- **Motivation** – to provide a reduced cost license (50% off commercial pricing) to all non-academic GAMS users of MARKAL/TIMES
- **Drawback** – those with an ETSAP Runtime license will
 - only be able to run the various variant of MARKAL/TIMES
 - not be able to add new Variables/Equations to the model generator
 - Outreach pricing goes away, though current Outreach licenses may continue to be used/maintained
- **Implementation**
 - a comprehensive list of Variables/Equation names kept and used to “seed” an “empty” restart file *{Already prepared - Excel file with full list processed by a generic GMS routine to create the restart file}*
 - the existing TIMES code will need to have the conditional code that checks for the declaration of certain Variables/Equations adjusted *{~2-days including testing}*
 - all model runs would start from a restart file (either “empty” or with previous run results)

ETSAP Runtime GAMS License (2/2)

- **Licensing Options**

- Academic is unchanged
- New users may either employ a full GAMS license or an ETSAP Runtime license
- Existing users may keep their current licenses or move to an ETSAP Runtime license

- **Pricing**

GAMS Module	List	Outreach	ETSAP Runtime	Savings
GAMS/Base	\$3,200	\$2,250	\$1,600	\$ 650
GAMS/CPLEX-XPRESS-GUROBI	\$9,600	\$9,600	\$4,800	\$4,800
GAMS/CONOPT-MINOS	\$3,200	\$2,250	\$1,600	\$ 650

- New cost for typical GAMS/CPLEX platform = \$6,400 (savings of \$5,450)
- Annual maintenance would be based upon the list, (current) Outreach, or ETSAP Runtime pricing as appropriate

IEA-ETSAP Annex XII 4th Workshop – December 10-11, 2012 Lisbon, Portugal

3

TIMES Cloud Platform (1/9)

- **Motivation is to lower entry/ongoing costs and provide more powerful computing platform when needed**

- “Rent” platforms with licensed version of all software instead of confronting the upfront costs, or to gain temporary access to more powerful machines
- Cloud providers offer a range of hardware (e.g. 68GB of core memory) with up-to-date OS, where hardware investment and maintenance could be a multiple of the cost of a cloud machine.
- Trade-off between software investment and a cloud lease, the amount of use determines what's cheaper, where cloud providers charge by the hours a machine is “on” - if idle that's the users problem.
- Cloud has to potential to handle larger models and to solve a number of scenario runs in a distributed manner such that turnaround time is essentially reduced to the number of runs bundled.

IEA-ETSAP Annex XII 4th Workshop – December 10-11, 2012 Lisbon, Portugal

4

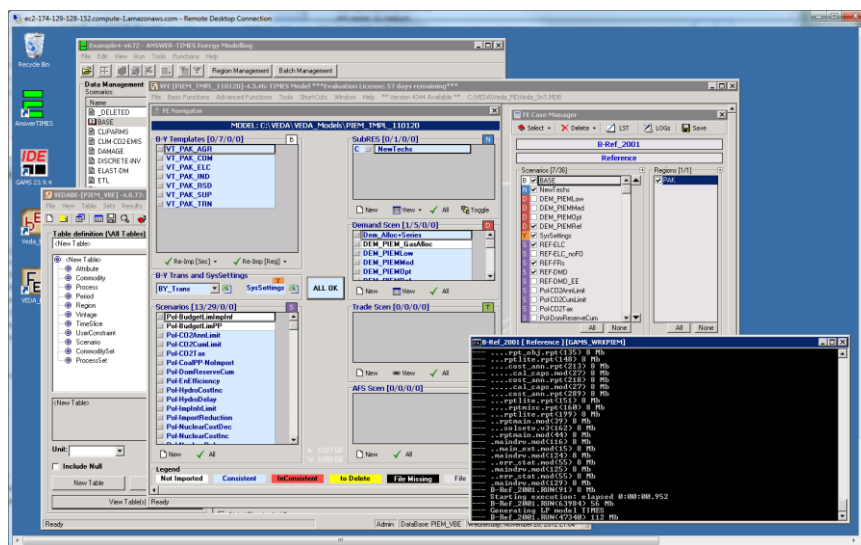
TIMES Cloud Platform (2/9)

- **What is the “Cloud”** – a network of computing platforms of varying size and power accessible over the web
- **What is an “Instance”** – a system configured for a particular purpose, with all needed software components (properly licensed) available
- **What is a TIMES-Cloud Instance** – a Windows machine with GAMS/Solver(s), VEDA and/or ANSWER, Excel, and RAR/ZIP available

IEA-ETSAP Annex XII 4th Workshop – December 10-11, 2012 Lisbon, Portugal

5

TIMES Cloud Platform (3/9)



IEA-ETSAP Annex XII 4th Workshop – December 10-11, 2012 Lisbon, Portugal

6

TIMES Cloud Platform (4/9)

- **How does one operate in the Cloud?**
 - Login with a secure password (for your licensed instance) via Start/Remote Desktop
 - Standard Windows platform
 - VEDA/ANSWER operate exactly the same
 - Files that are needed locally or need to get up to the Cloud can either
 - be moved via Copy/Paste
 - the user can “attach” their hard drive to the Cloud
 - large files can be RARed 1st and move pretty quickly
 - Each restart of an instance generates a new IP address, process can be scripted

IEA-ETSAP Annex XII 4th Workshop – December 10-11, 2012 Lisbon, Portugal

7

TIMES Cloud Platform (5/9)

- **What’s been done**
 - Small instance setup with all components
 - Both ANSWER/VEDA (evaluation) up and running TIMES smoothly
 - Large instance cloned
 - Modest and large models run performed
- **Initial Experiments**
 - Small instance
 - PIEM (500k row model) takes 180secs rather than 90secs
 - Large instance
 - Single FACETS (6m rows) +6hr runtime --- bottleneck is Simplex
 - Partial parallel solve of 7-PIEM scenarios, matched to machine horsepower
 - cut run time ~50% by splitting into 2 groups
 - can be done on local multi-CPU machines
 - needs to be handled by VEDA/ANSWER, otherwise manual intervention necessary
- **Don’t expect miracles**

```
Instance ID : i-7b6aa407
IP Address  : 174.129.128.192
Availability Zone : us-east-1d
Instance Size : m1.large
Architecture : AMD64
Total Memory : 7.5 GB
Processing Power : 4 ECUs
I/O Performance : High
```

```
Instance ID : i-7b6af404
IP Address  : 174.129.99.91
Availability Zone : us-east-1e
Instance Size : m2.xlarge
Architecture : AMD64
Total Memory : 68.4 GB
Processing Power : 26 ECUs
I/O Performance : High
```

IEA-ETSAP Annex XII 4th Workshop – December 10-11, 2012 Lisbon, Portugal

8

TIMES Cloud Platform (6/9)

- **Distributed instances for solving in parallel** *{pending}*
 - Cloud machines can be started and stopped with scripts in an automated fashion.
 - Having n scenarios with running times larger than 10 minutes (otherwise inter machine communication will cancel out the gain in parallel processing), one can calculate the number i of worker machines one needs to fire up (and pay for) to get the n scenarios solved in a fixed time frame of k hours [= #scenarios/machines * (solve time + overhead)].
 - A cheap machine (local or on the cloud) can act as a file server to collect the result of the workers, so one can terminate the (expensive) workers after the jobs are done and retrieve the results after more than k hours.

IEA-ETSAP Annex XII 4th Workshop – December 10-11, 2012 Lisbon, Portugal

9

TIMES Cloud Platform (7/9)

- **Amazon boxes sample** <http://aws.amazon.com/ec2/instance-types/>

Standard Instances

First Generation

First generation (M1) Standard instances provide classic platform that is well suited for a wide variety of applications.

M1 Small Instance – default++

1.7 GiB memory
1 EC2 Compute Unit (1 virtual core with 1 EC2 Compute Unit)
160 GB instance storage
32-bit or 64-bit platform
I/O Performance: Moderate
EBS-Optimized Available: No
API name: m1.small

M1 Medium Instance

3.75 GiB memory
2 EC2 Compute Units (1 virtual core with 2 EC2 Compute Units)
315 GB instance storage
32-bit or 64-bit platform
I/O Performance: Moderate
EBS-Optimized Available: No
API name: m1.medium

M1 Large Instance

7.5 GiB memory
4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units)
850 GB instance storage
64-bit platform
I/O Performance: High
EBS-Optimized Available: 500 Mbps
API name: m1.large

M1 Extra Large Instance

15 GiB memory
8 EC2 Compute Units (4 virtual cores with 2 EC2 Compute Units)
1,690 GB instance storage
64-bit platform
I/O Performance: High
EBS-Optimized Available: 1000 Mbps
API name: m1.xlarge

Second Generation

Second generation (M3) Standard instances provide custody of processing performance compared to First Generation instances; applications that require higher absolute CPU and memory from the performance of Second Generation Standard instances, and memcached.

M3 Extra Large Instance

15 GiB memory
13 EC2 Compute Units (4 virtual cores with 3.25 EC2 Compute Units)
EBS storage only
64-bit platform
I/O Performance: Moderate
EBS-Optimized Available: No
API name: m3.xlarge

M3 Double Extra Large Instance

30 GiB memory
26 EC2 Compute Units (8 virtual cores with 3.25 EC2 Compute Units)
EBS storage only
64-bit platform
I/O Performance: High
EBS-Optimized Available: No
API name: m3.2xlarge

High-Memory Instances

Instances of this family offer large memory sizes for high throughput and caching applications.

High-Memory Extra Large Instance

17.1 GiB of memory
6.5 EC2 Compute Units (2 virtual cores with 3.25 EC2 Compute Units)
420 GB of instance storage
64-bit platform
I/O Performance: Moderate
EBS-Optimized Available: No
API name: m2.xlarge

High-Memory Double Extra Large Instance

34.2 GiB of memory
13 EC2 Compute Units (4 virtual cores with 3.25 EC2 Compute Units)
850 GB of instance storage
64-bit platform
I/O Performance: High
EBS-Optimized Available: No
API name: m2.2xlarge

High-Memory Quadruple Extra Large Instance

68.4 GiB of memory
26 EC2 Compute Units (8 virtual cores with 3.25 EC2 Compute Units)
1690 GB of instance storage
64-bit platform
I/O Performance: High
EBS-Optimized Available: 1000 Mbps
API name: m2.4xlarge

IEA-ETSAP Annex XII 4th Workshop – December 10-11, 2012 Lisbon, Portugal

10

TIMES Cloud Platform (8/9)

- **Amazon pricing example** <http://aws.amazon.com/ec2/pricing/>
 - Upfront fee
 - Hourly connect charge

Linux		Windows			
Light Utilization Reserved Instances					
Region: US East (N. Virginia)					
		1 yr Term		3 yr Term	
		Upfront	Hourly	Upfront	Hourly
Standard Reserved Instances					
Small (Default)	\$69	\$0.059 per Hour	\$106.30	\$0.051 per Hour	
Medium	\$138	\$0.118 per Hour	\$212.30	\$0.103 per Hour	
Large	\$276	\$0.235 per Hour	\$425.20	\$0.204 per Hour	
Extra Large	\$552	\$0.47 per Hour	\$850.40	\$0.408 per Hour	
Second Generation Standard Reserved Instances					
Extra Large	\$607	\$0.504 per Hour	\$935	\$0.432 per Hour	
Double Extra Large	\$1214	\$1.008 per Hour	\$1870	\$0.864 per Hour	
Micro Reserved Instances					
Micro	\$23	\$0.014 per Hour	\$35	\$0.012 per Hour	
High-Memory Reserved Instances					
Extra Large	\$353	\$0.29 per Hour	\$548	\$0.245 per Hour	
Double Extra Large	\$706	\$0.58 per Hour	\$1096	\$0.49 per Hour	
Quadruple Extra Large	\$1412	\$1.16 per Hour	\$2192	\$0.98 per Hour	
High-CPU Reserved Instances					
Medium	\$178	\$0.165 per Hour	\$273	\$0.153 per Hour	
Extra Large	\$712	\$0.66 per Hour	\$1092	\$0.612 per Hour	
Cluster Compute Reserved Instances					
Quadruple Extra Large	\$1450	\$0.922 per Hour	\$2235	\$0.922 per Hour	
Eight Extra Large	\$1762	\$1.114 per Hour	\$2710	\$1.114 per Hour	
Cluster GPU Reserved Instances					
Quadruple Extra Large	\$2410	\$1.534 per Hour	\$3700	\$1.534 per Hour	
High-I/O Reserved Instances					
Quadruple Extra Large	\$2576	\$1.957 per Hour	\$3884	\$1.63 per Hour	

IEA-ETSAP Annex XII 4th Workshop – December 10-11, 2012 Lisbon, Portugal

11

TIMES Cloud Platform (9/9)

- **Where from here?**
 - Setup and test distributed solve procedure
 - Poll ETSAP and other MARKAL/TIMES users on whether there's interest - if so why and for what?
 - Look into i/o and related issues
 - Investigate other cloud providers
- **Prepare procedure and full pricing information for various instances and configurations**
- **If proceeding**
 - Prepare necessary scripts for the Cloud environment
 - Coordinate and introduce any desired changes to the TIMES environment
- **Appropriateness as an IEW Paper?**

IEA-ETSAP Annex XII 4th Workshop – December 10-11, 2012 Lisbon, Portugal

12

Thank You!

DecisionWare, Inc

- Gary Goldstein gary.a.goldstein@gmail.com

GAMS Development Corp

- Michael Bussieck mbussieck@gams.com
- Alex Meeraus ameeraus@gams.com