QoS Modeling and Evaluation in Cloud Environments

www.modaclouds.eu

Danilo Ardagna
danilo.ardagna@polimi.it
Dipartimento di Elettronica, Informazione e Bioingegneria
Politecnico di Milano

MODAClouds Challenges & Objectives

- MODAClouds: MOdel-Driven Approach for design and execution of applications on multiple Clouds
- Focus on needs of Cloud-based Application Developers and Operators
- Challenges:
  - Avoid vendor lock-in
  - Support risk analysis and management
  - Guarantee quality assurance
- Objective: *to provide methods, a decision support system, an IDE and a runtime environment to support*
  - *High-level design*
  - *Early prototyping*
  - *Semi-automatic code generation*
  - *Automatic (re)deployment of applications on multi-Clouds with guaranteed QoS*

MODAClouds

MODAClouds (www.modaclouds.eu)

- Integrated Project n. 318484
- October 1st 2012 – September 30th 2015

Institute eAustria Timisoara

Imperial College London

SIEMENS

Atos

ca technologies

MODAClouds

European Commission

Our Concept of Multi-Cloud

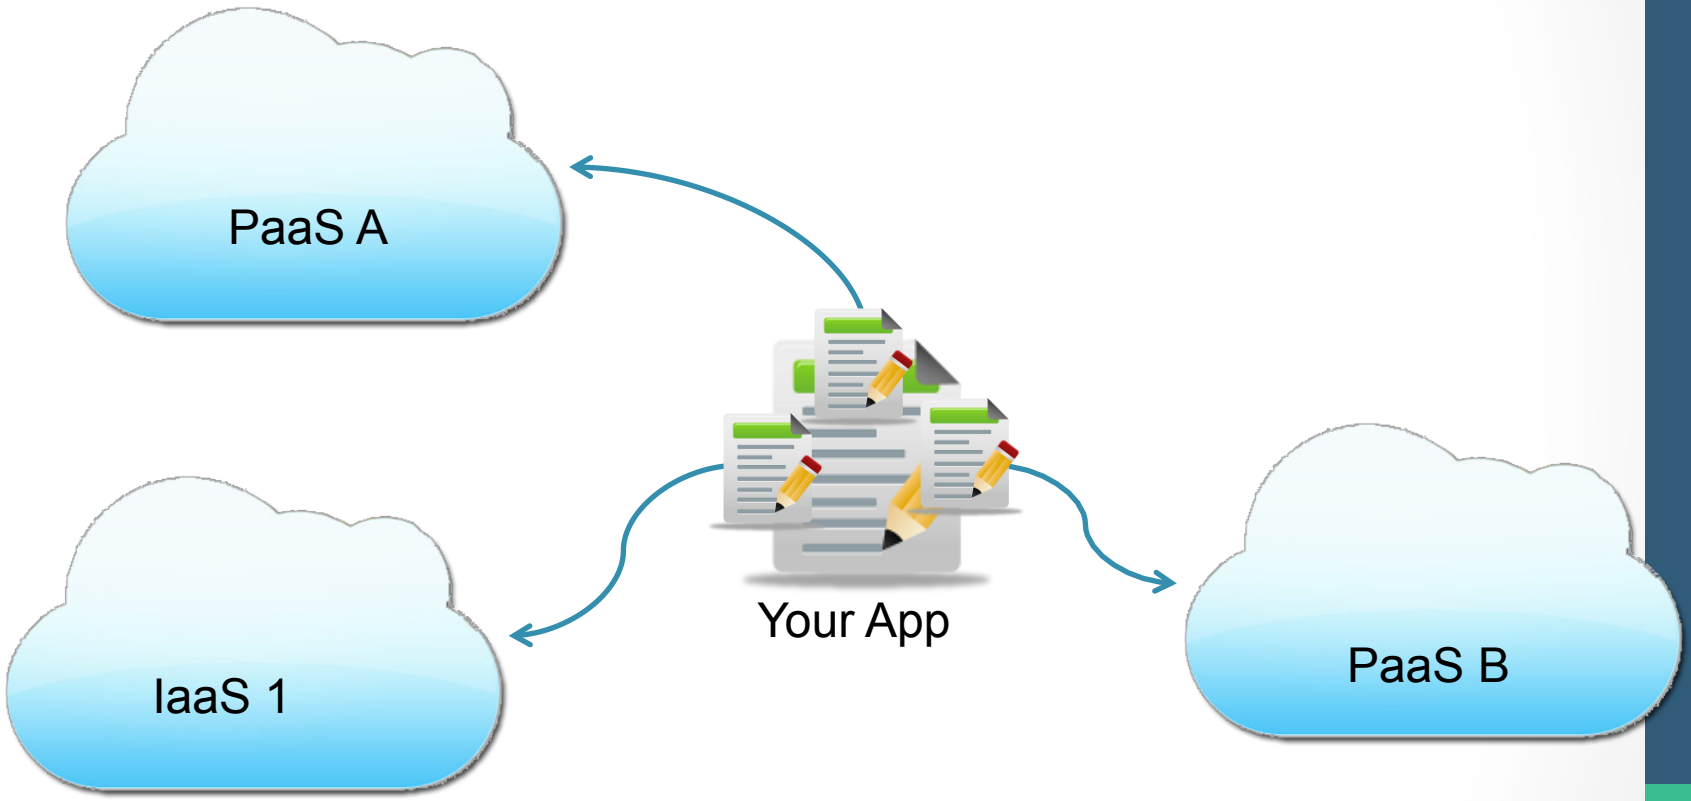

MODAClouds

MODAClouds Vision

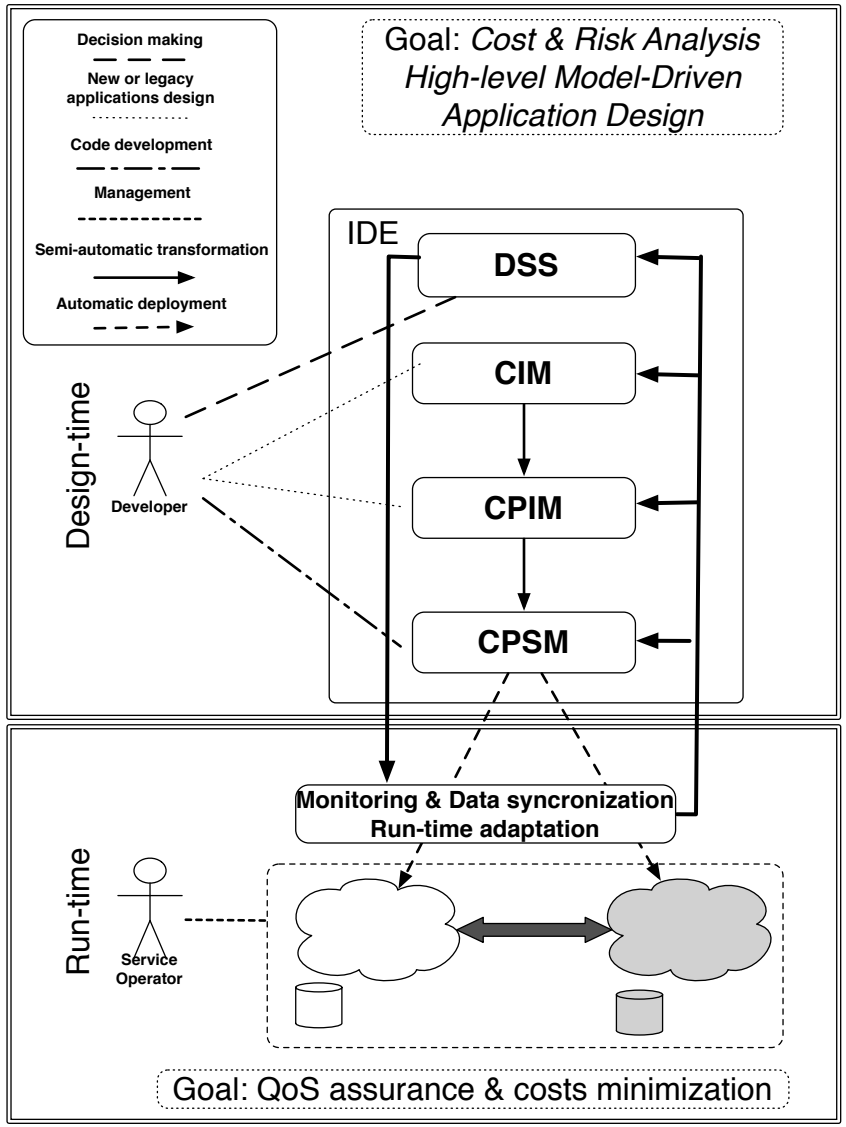

D. Ardagna et al. MODACLOUDS: A Model-Driven Approach for the Design and Execution of Applications on Multiple Clouds. MiSE 2012 Workshop Proceedings.

An example

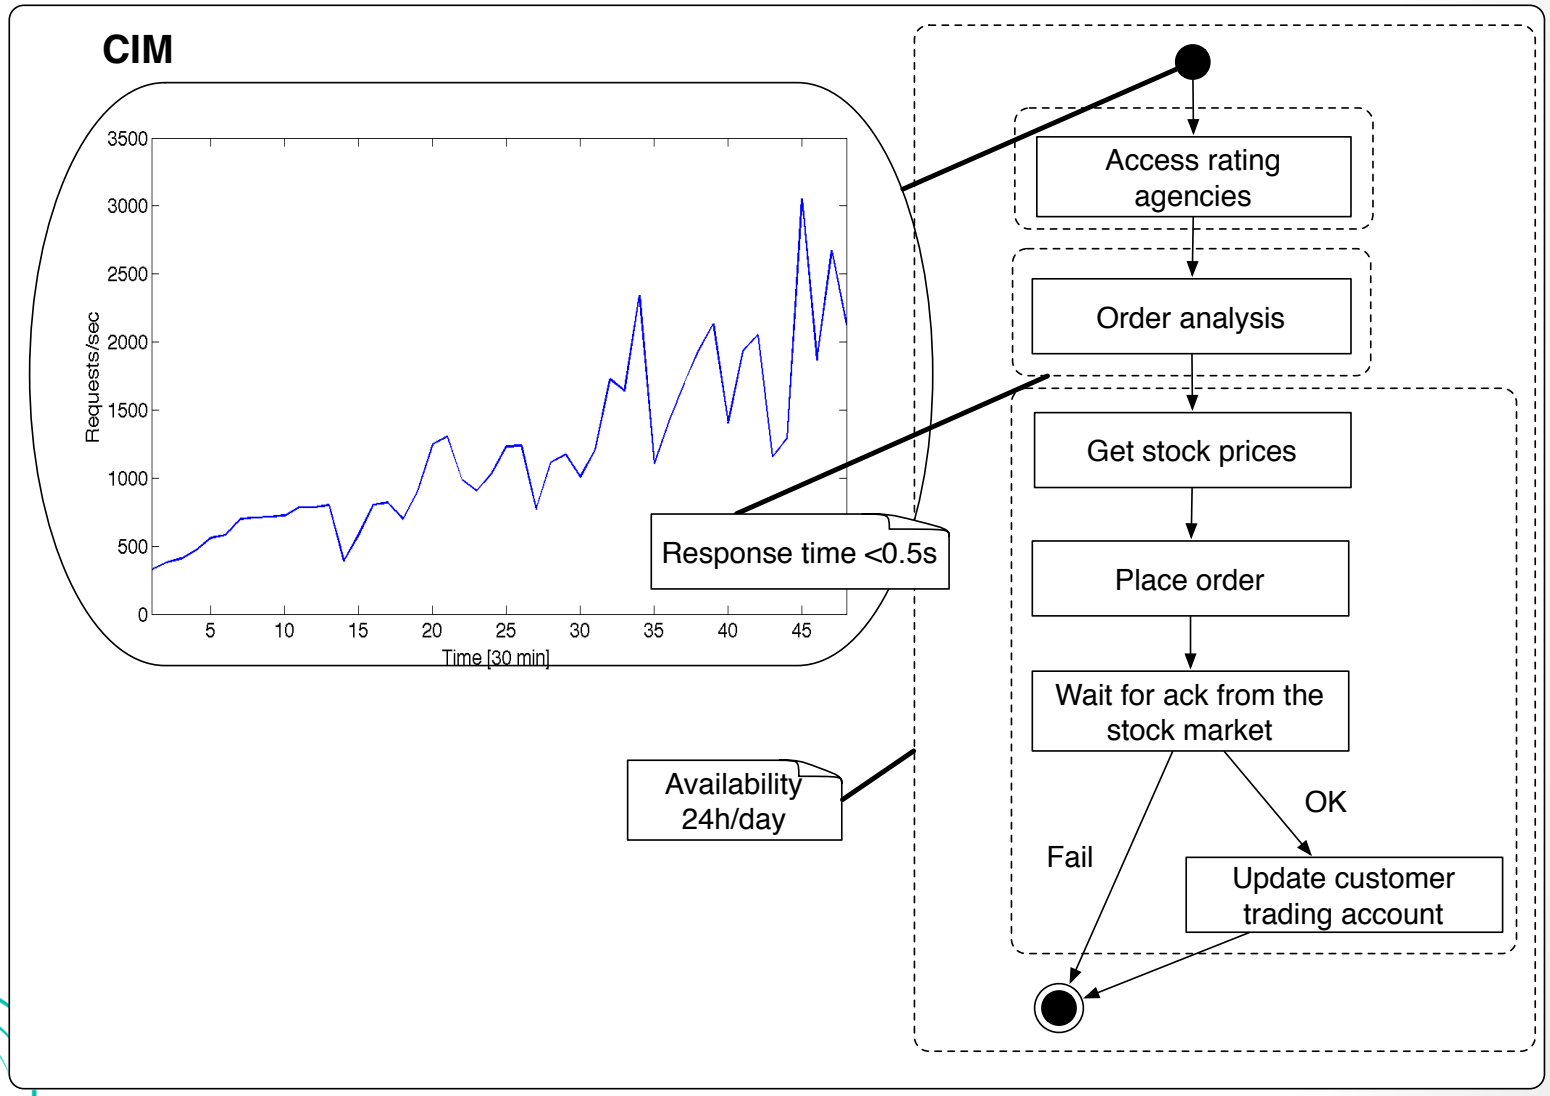

MODAClouds

An example

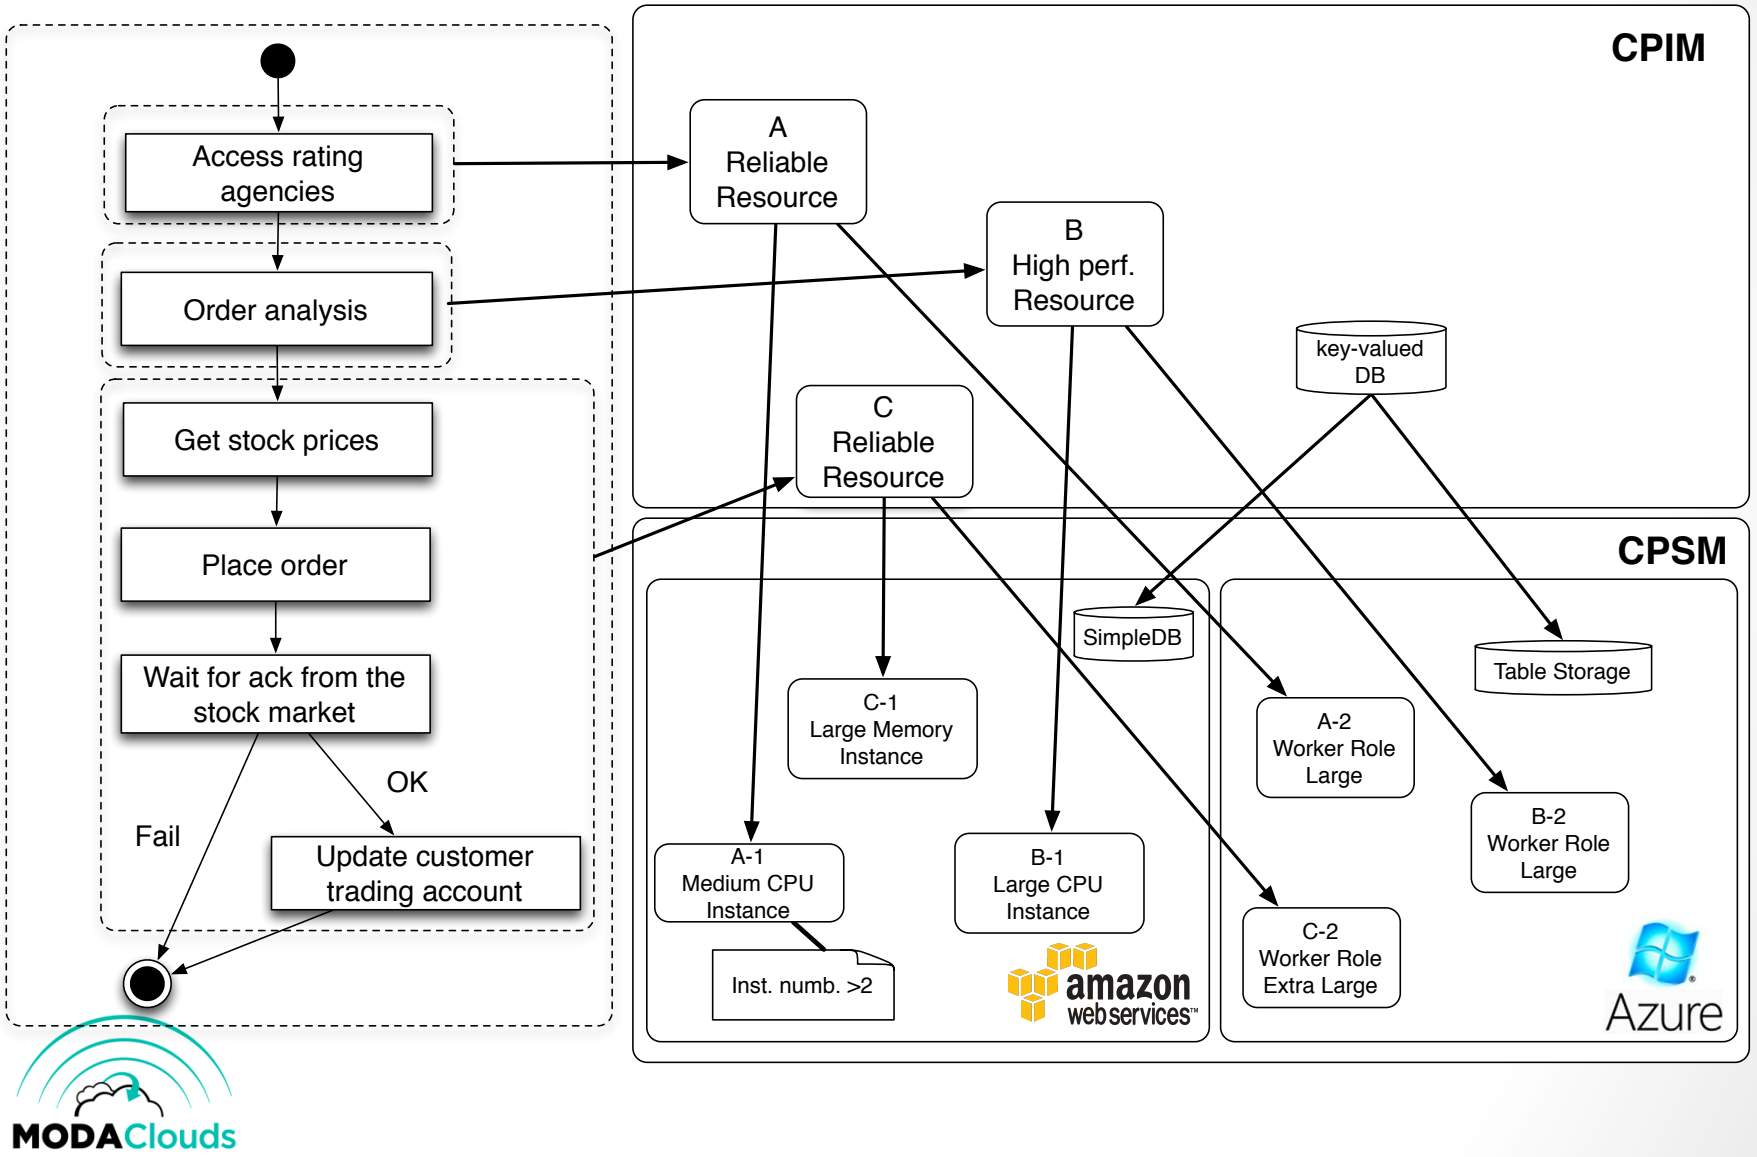

MODAClouds

Cloud Applications Performance Modelling Challenges

- Cloud performance can vary at any point in time
- Elasticity may not ramp at desired speeds
- QoS metrics can be in conflict
- Cost estimate is also difficult:
  - Pricing models vary from a Cloud provider to another
  - Several cost metrics (e.g., $/hour, $/GB-month, $/million I/O, etc…)
  - Costs follow the resource allocation and workload trends, so variable allocations and/or workloads lead to variable costs

MODAClouds

Goals

- Develop a methodology and a software tool for the model driven design, performance, and cost assessment of applications running in the Cloud:
  - Consider generic and specific Cloud
  - Run what-if analysis allowing to compare multiple configurations, Cloud services, Cloud providers, etc…

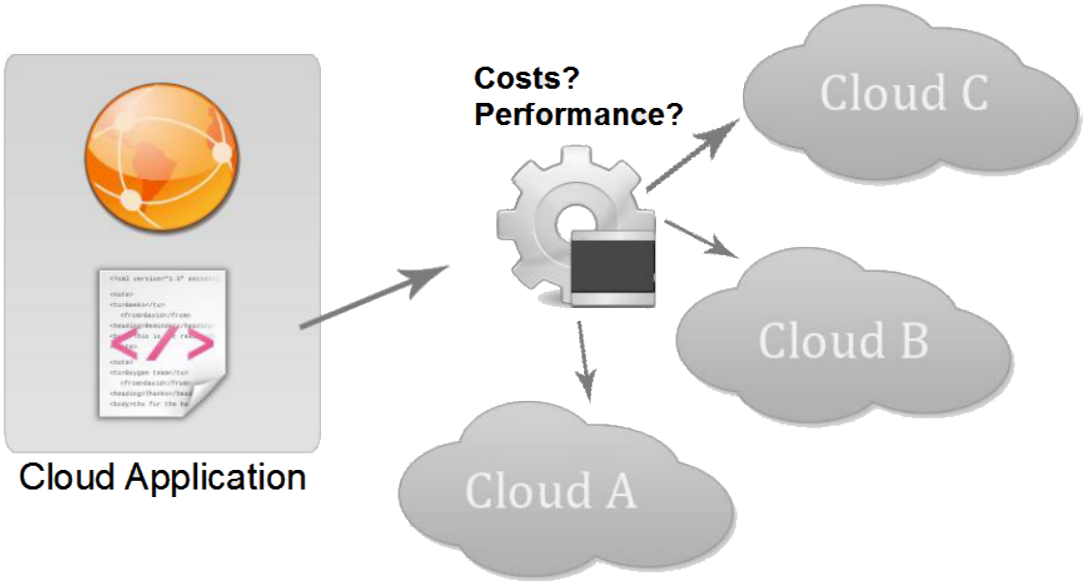

MODAClouds

Running SpecWeb on Amazon EC2

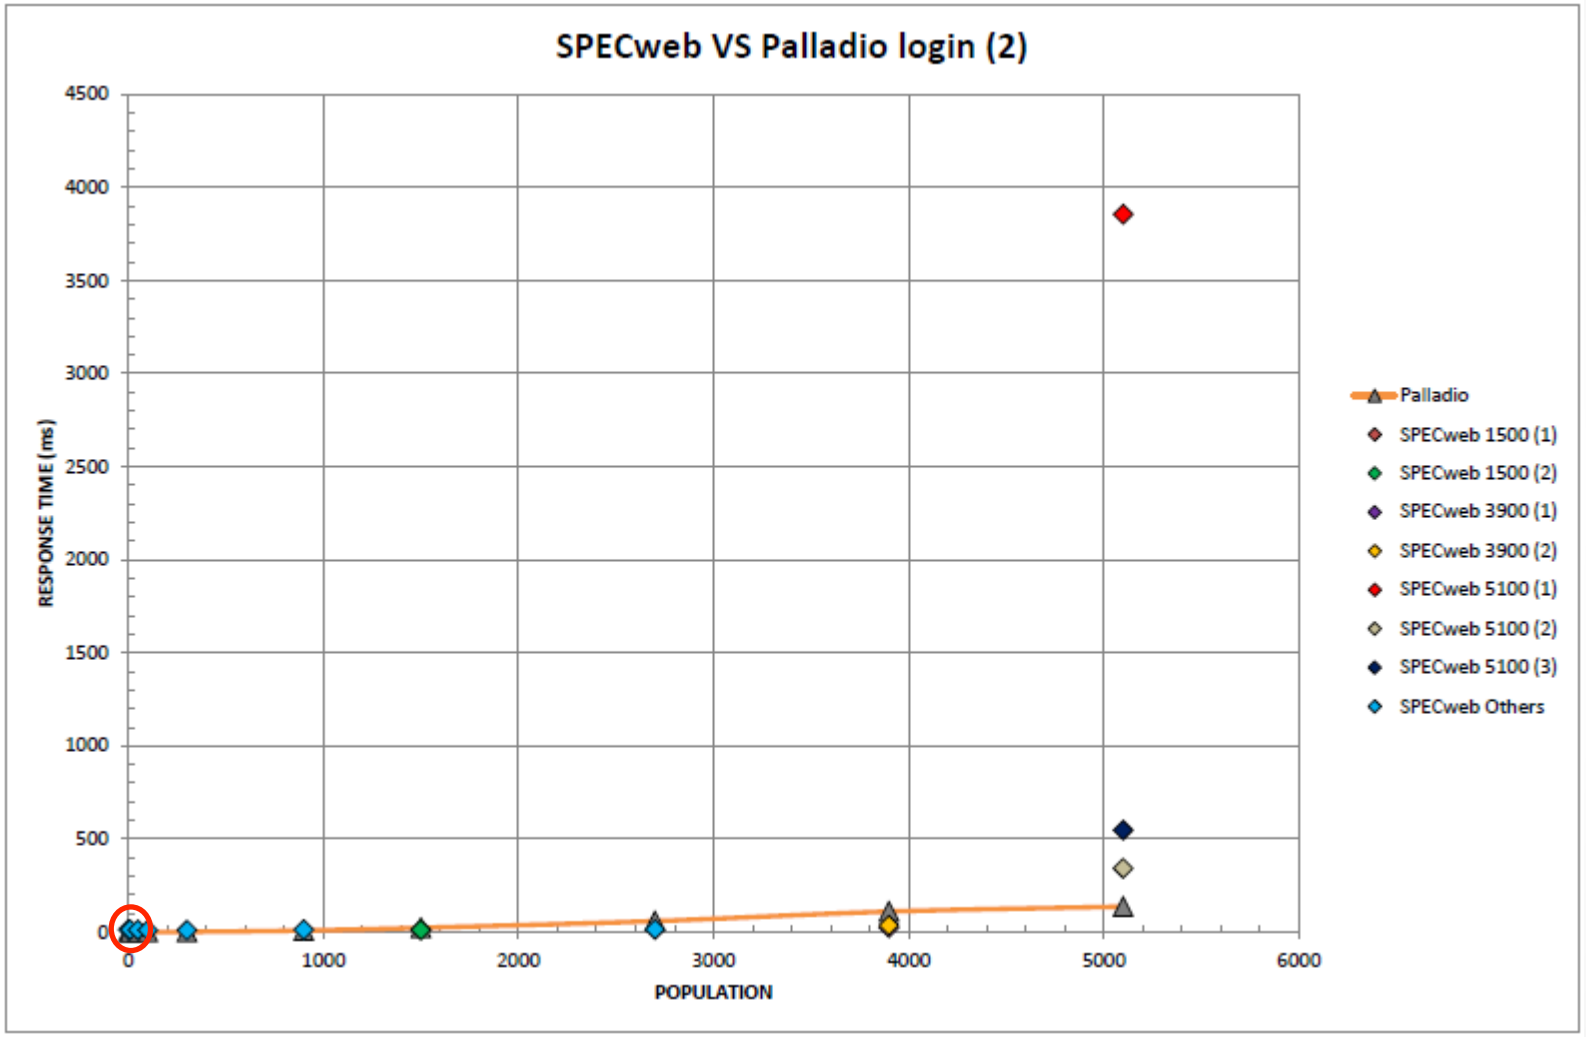

MODAClouds

Running SpecWeb on Amazon EC2

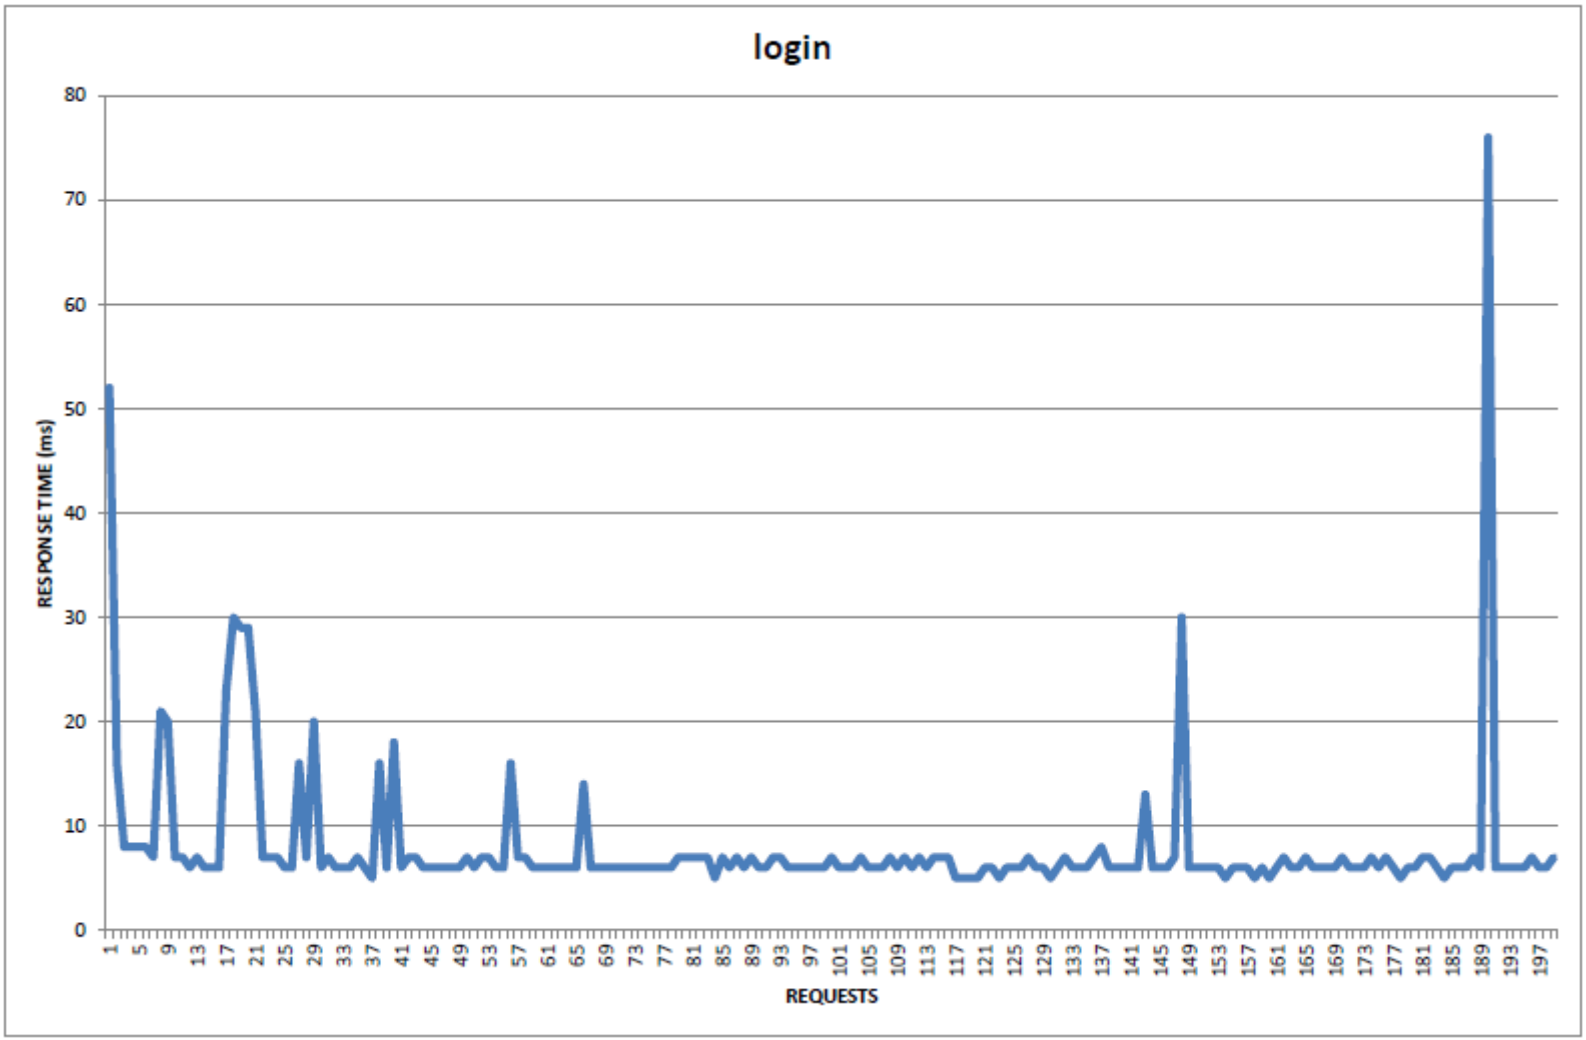

MODAClouds

Running SpecWeb on Amazon EC2

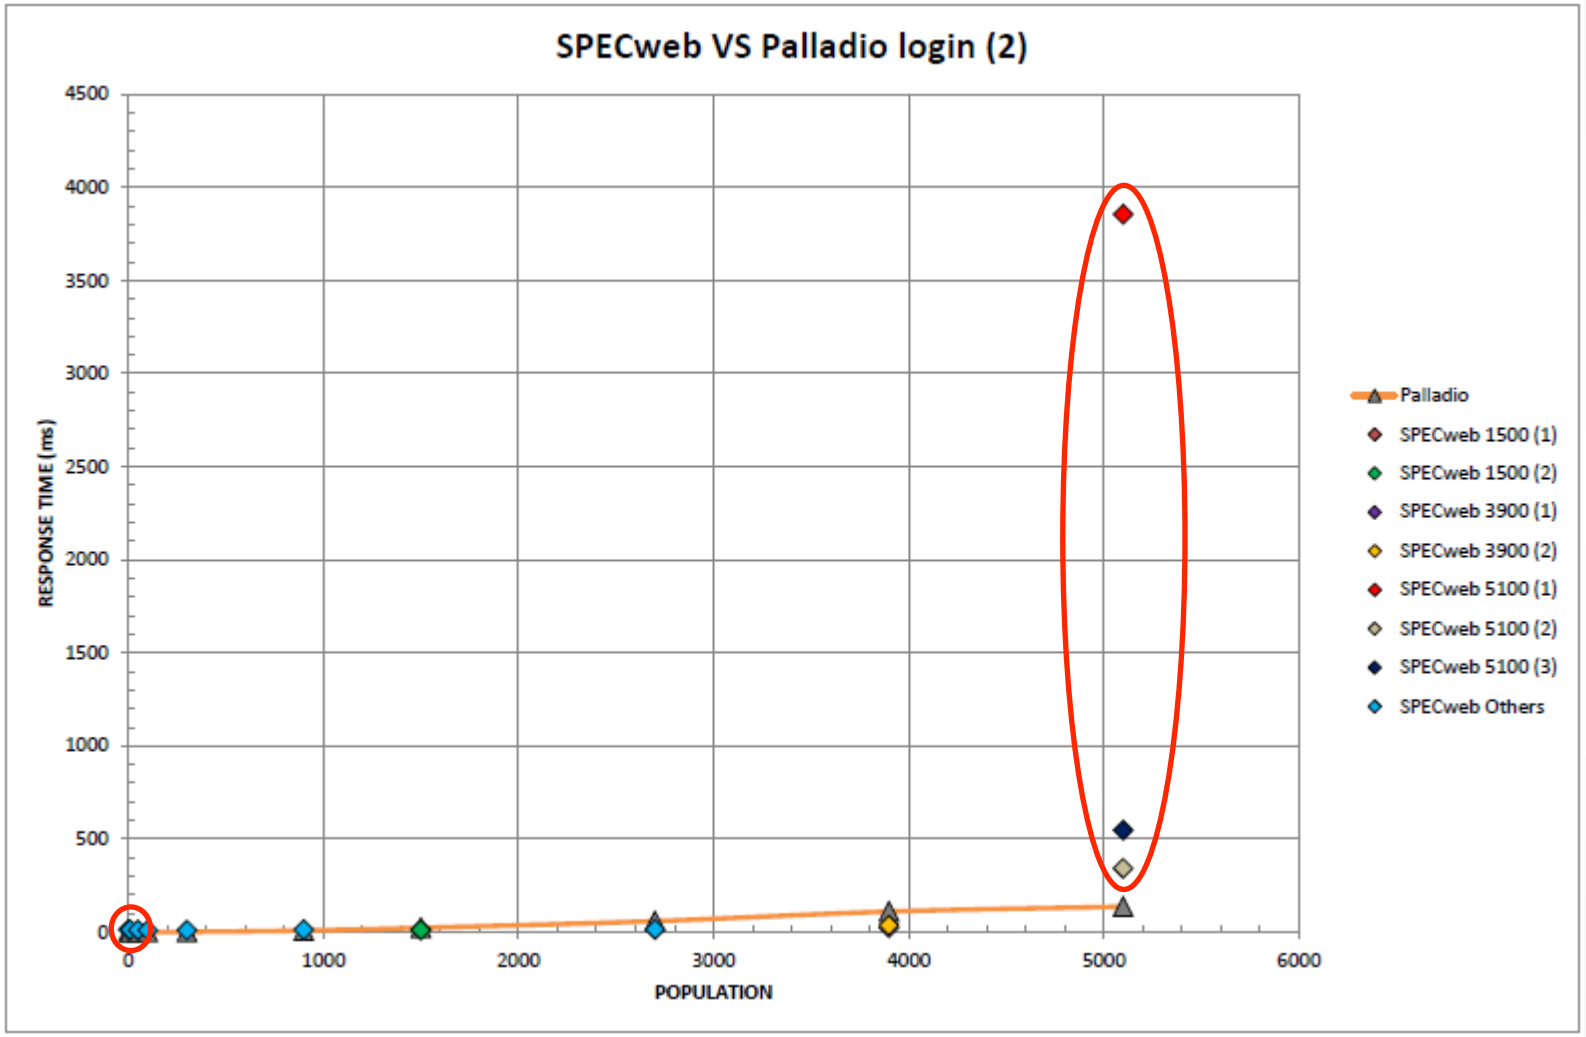

MODAClouds

Running SpecWeb on Amazon EC2

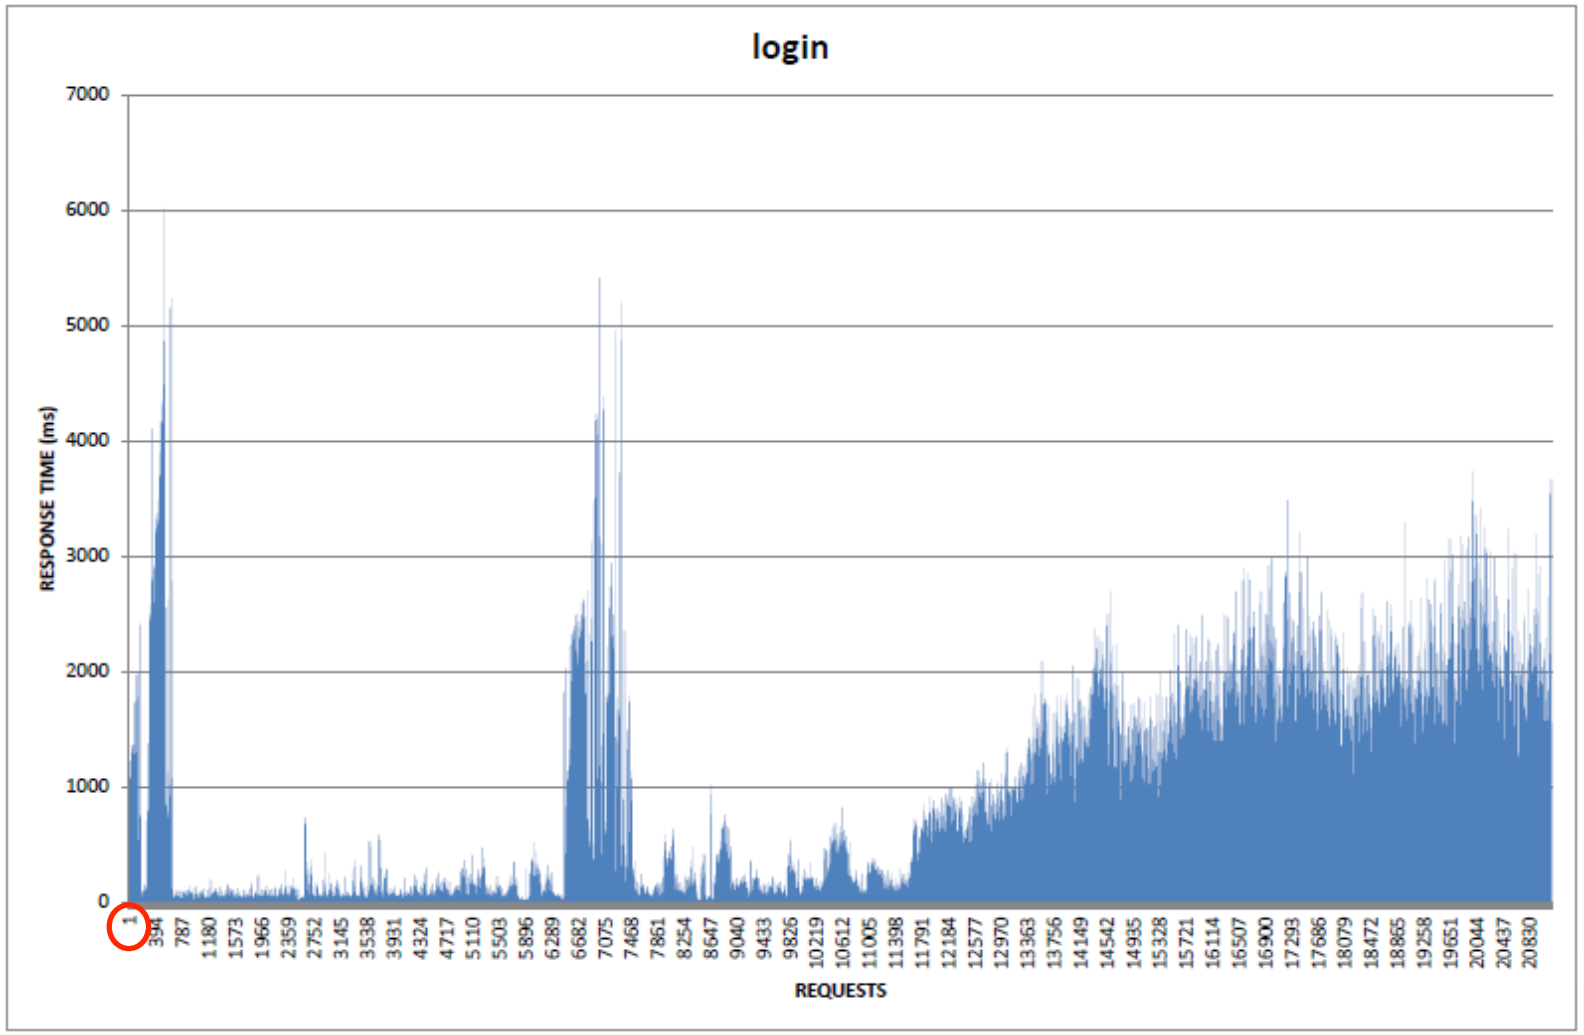

MODAClouds

WAN Effects

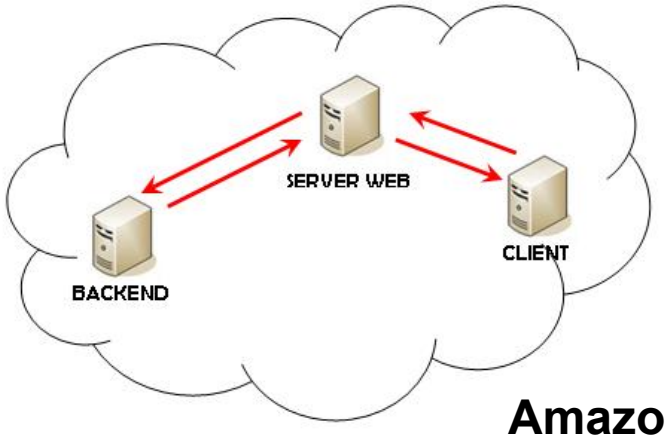

Amazon Network

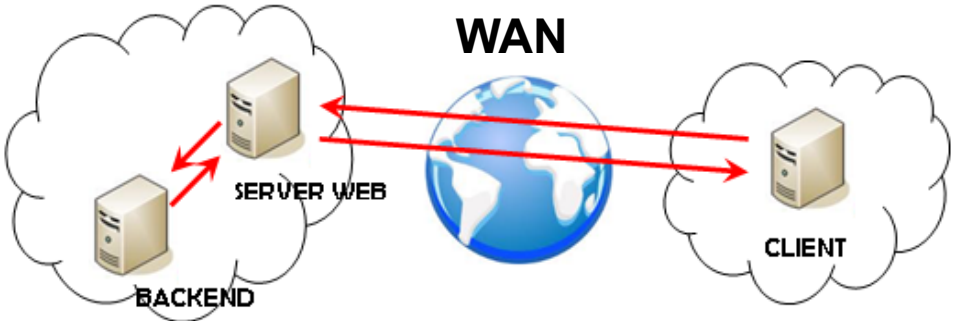

Amazon Network

Amazon Network

MODAClouds

US-East 18.00 – Single site configuration

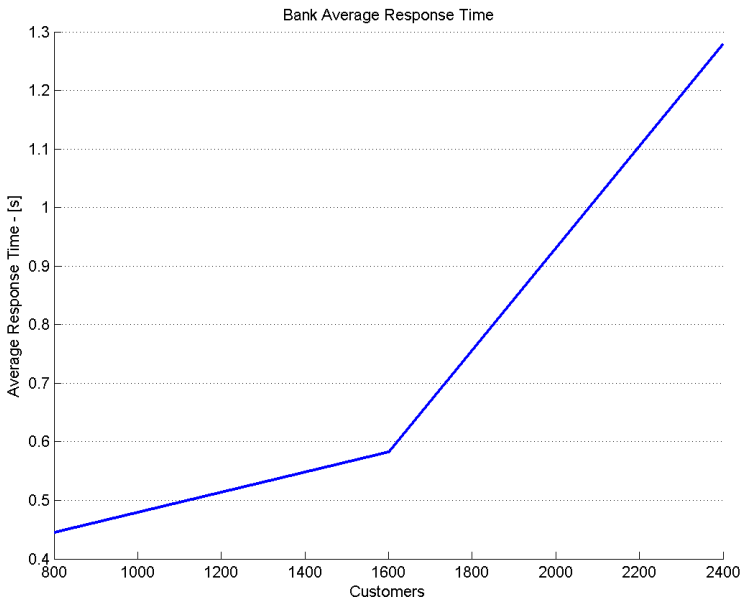

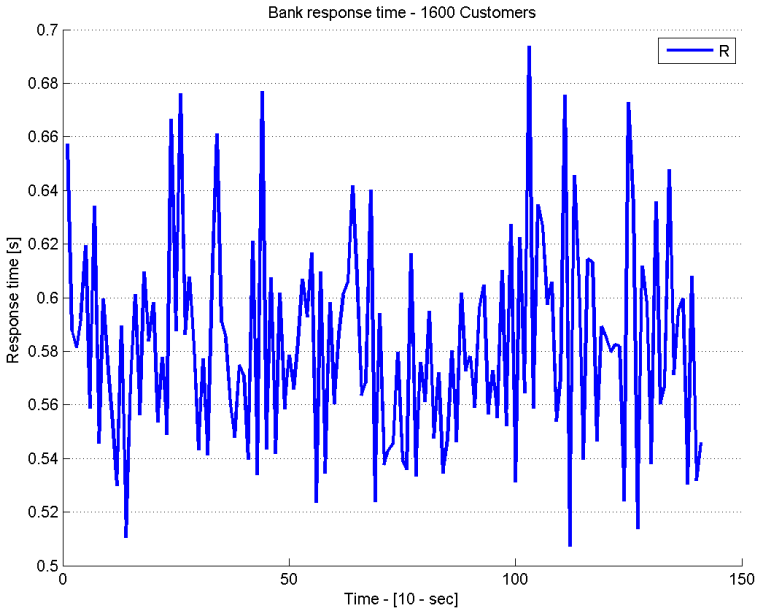

	800 Users	1600 Users	2400 Users
Min	0.4183	0.5073	0.9298
Max	0.4749	0.6939	1.6259
Average	0.4448	0.5823	1.2786
Std. Dev.	0.0100	0.0387	0.1640

MODAClouds

WAN Effects

18.00 US-East

Users	US-East	US-East-West	Latency (Δ)
800	0.4448	1.0342	0.5894
1600	0.5823	1.0946	0.5123
2400	1.2786	1.6444	0.3658

10.00 US-East

Users	US-East	US-East-West	Latency (Δ)
800	0.4639	1.0787	0.6148
1600	0.5653	2.0803	1.5150
2400	NA	NA	NA

MODAClouds

WAN Effects

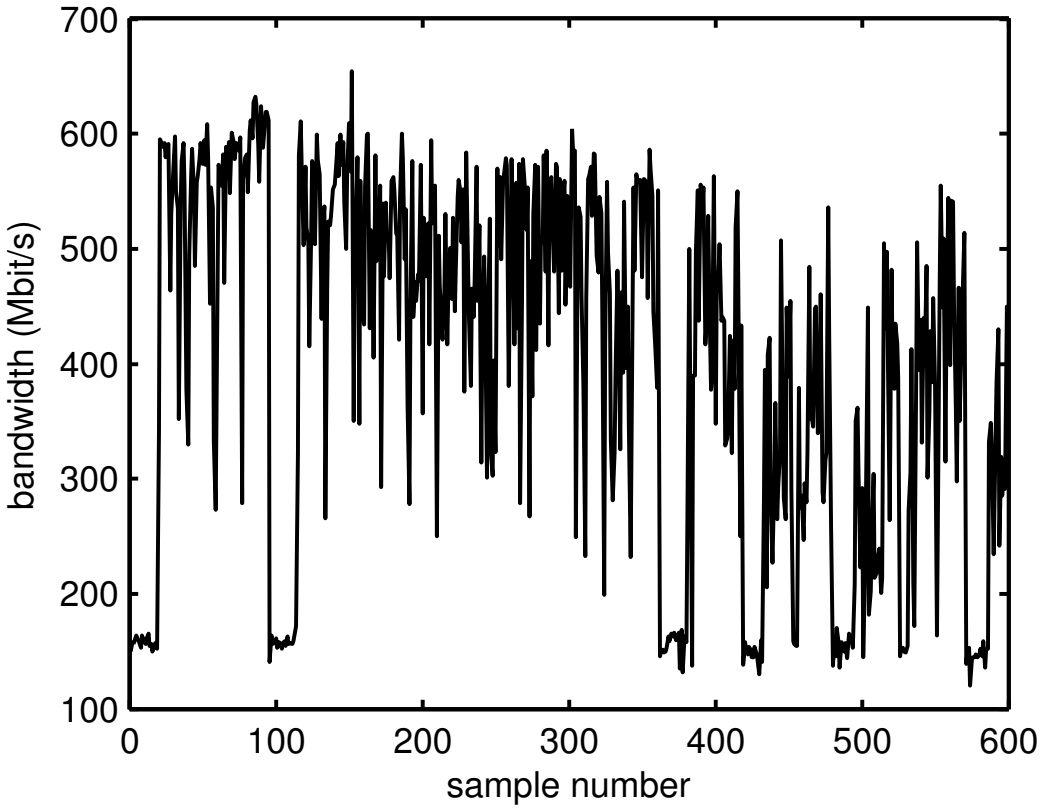

G. Casale, M. Tribastone. Modelling exogenous variability in cloud deployments. SIGMETRICS Performance Evaluation Review 40(4): 73-82 (2013)

MODAClouds

Our Starting Point: The Palladio Framework

- Developed at Uni Oldenburg, Uni Karlsruhe since 2003
- Domain-specific modelling language
- Targeted at:
  - Performance prediction of component-based Software Architectures
  - Multiple models and QoS metrics (CTMC, DTMC, LQN)
  - Support simulation, analytical solutions, design time exploration

MODAClouds

http://sdq.ipd.kit.edu/research/palladio_research_project/

Our Starting Point: The Palladio Framework

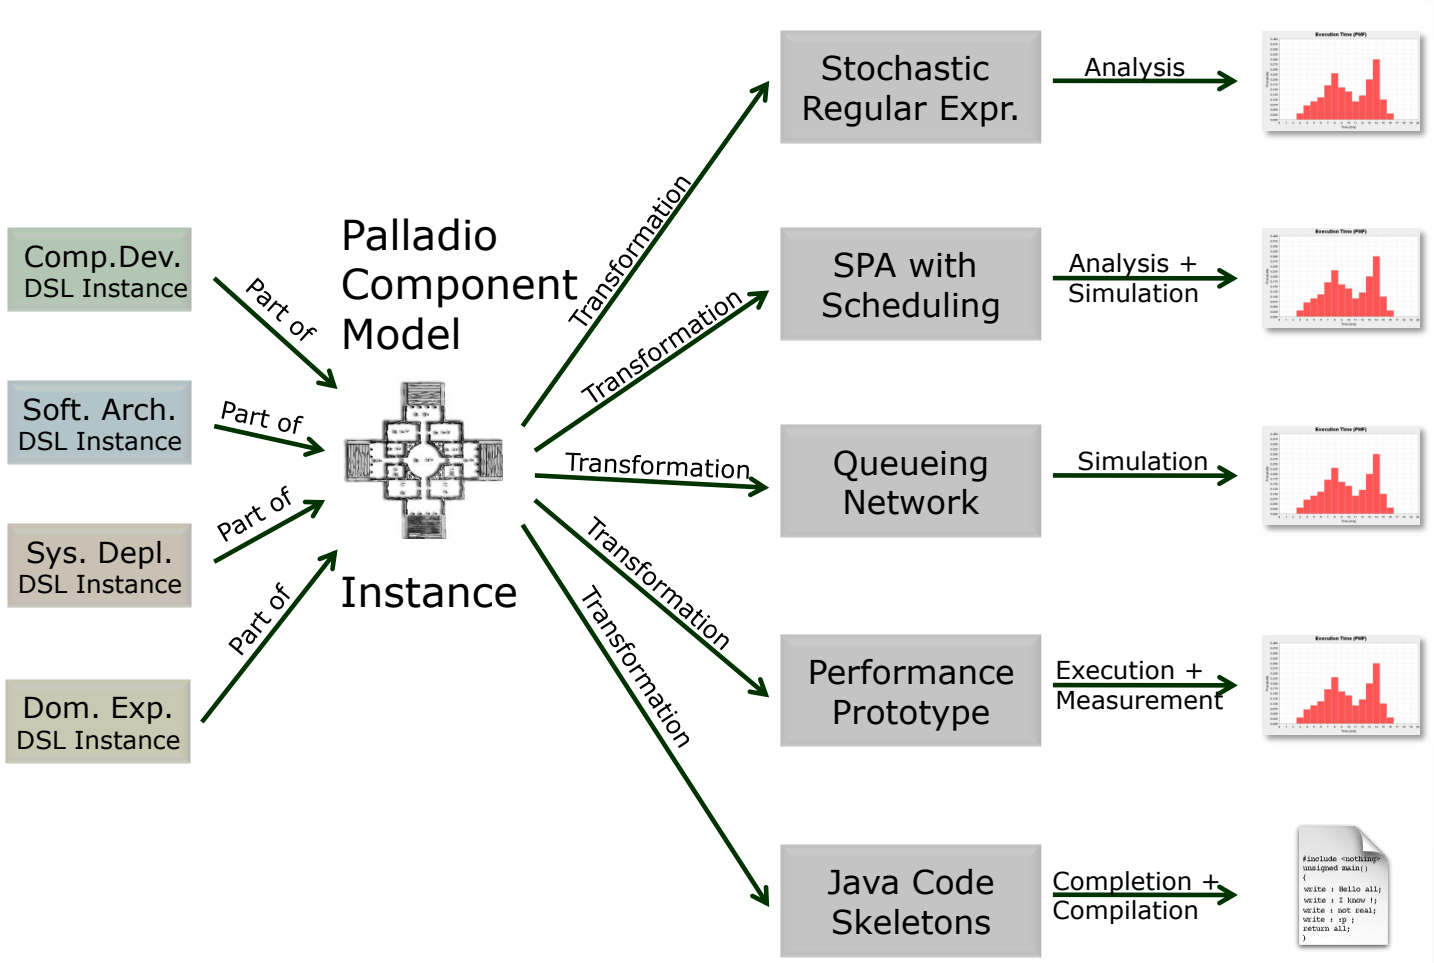

http://sdq.ipd.kit.edu/research/palladio_research_project/

Our Starting Point: The Palladio Framework

- Developed at Uni Oldenburg, Uni Karlsruhe since 2003
- Domain-specific modelling language
- Targeted at:
  - Performance prediction of component-based Software Architectures
  - Multiple models and QoS metrics (CTMC, DTMC, LQN)
  - Support simulation, analytical solutions, design time exploration

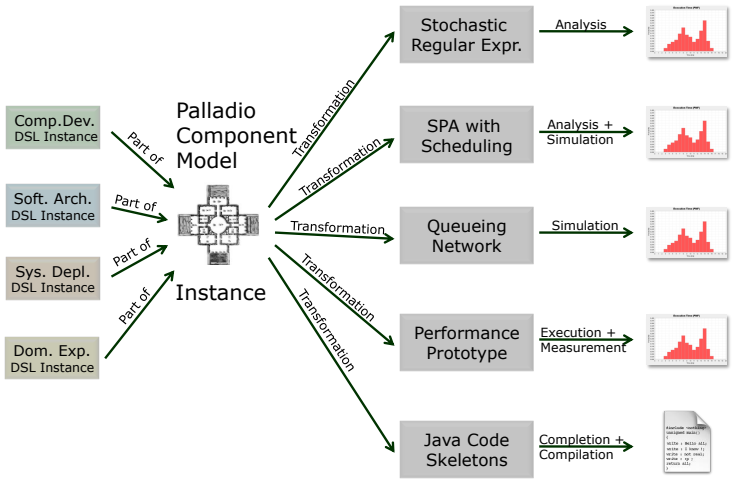

http://sdq.ipd.kit.edu/research/palladio_research_project/

MODAClouds

SPECWeb running example

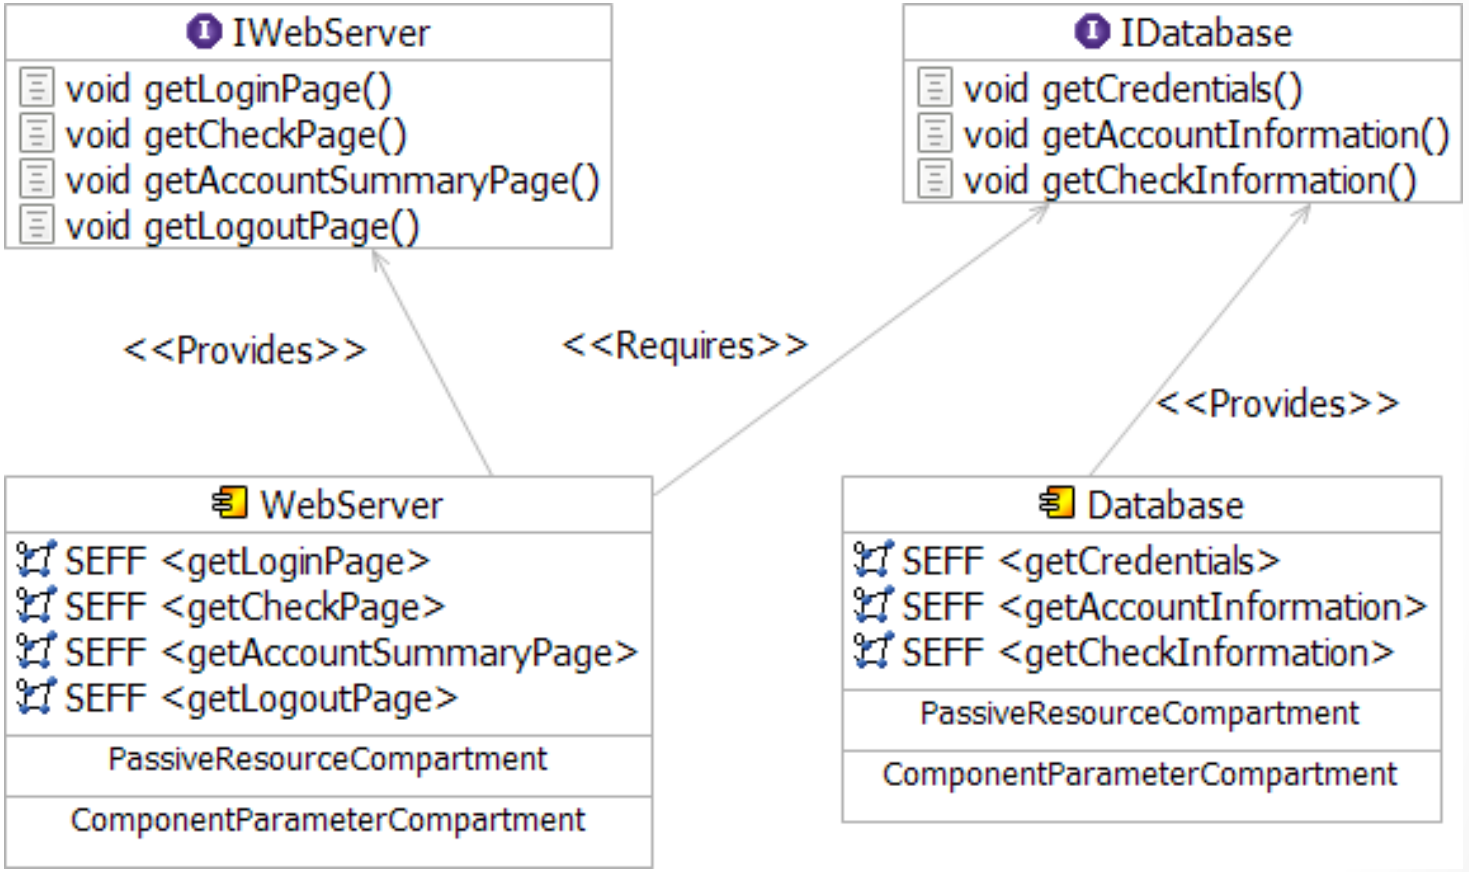

MODAClouds

SPECWeb running example

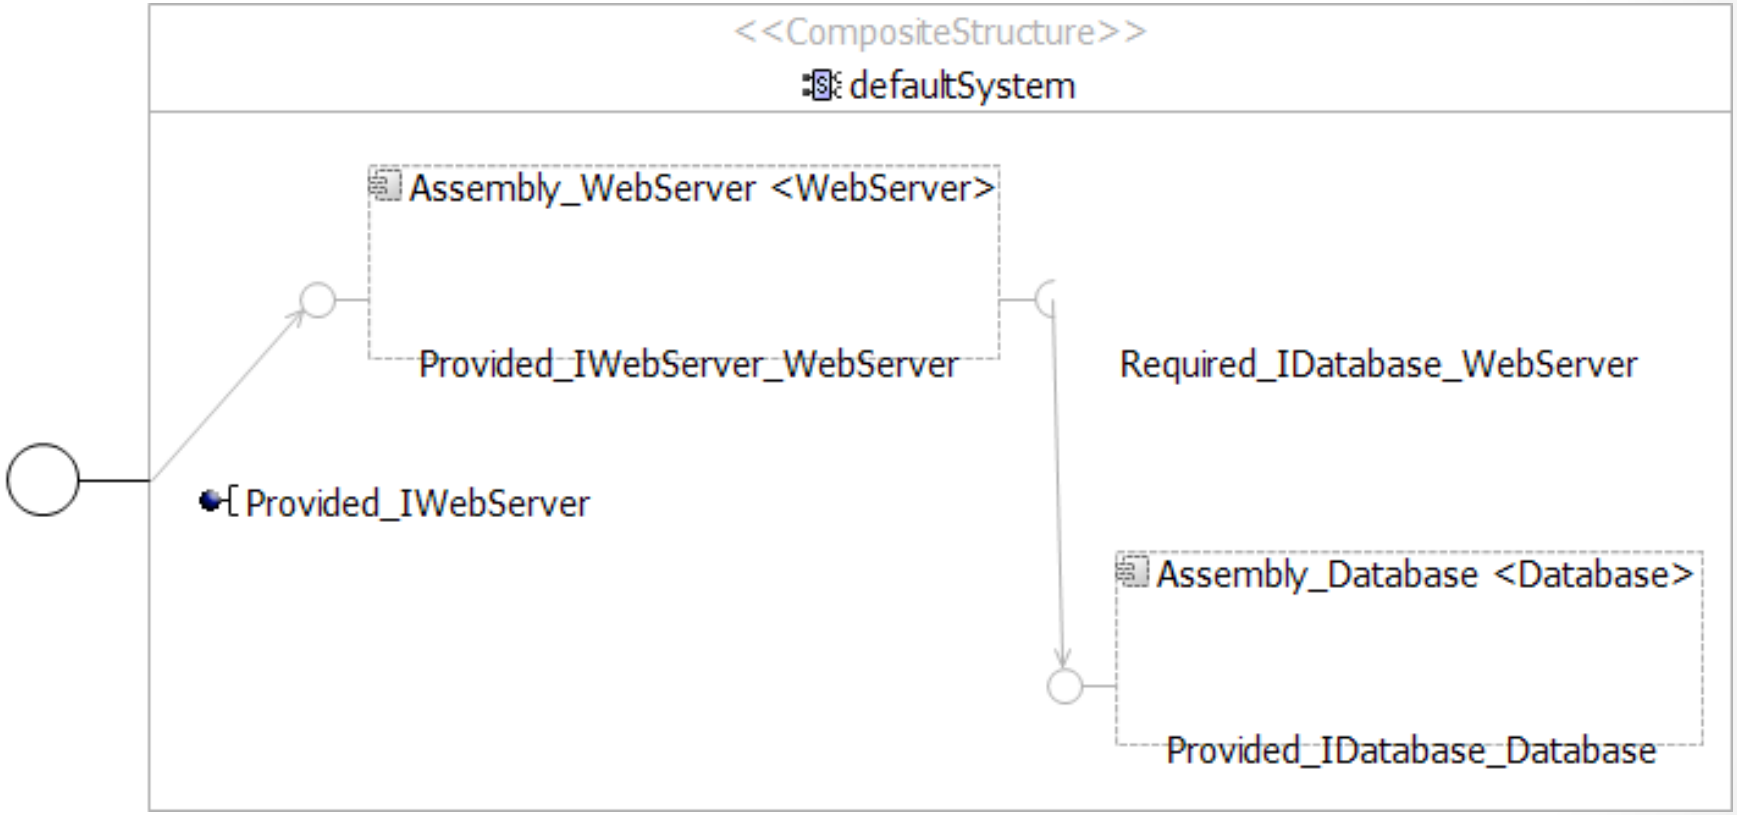

MODAClouds

SPECWeb running example

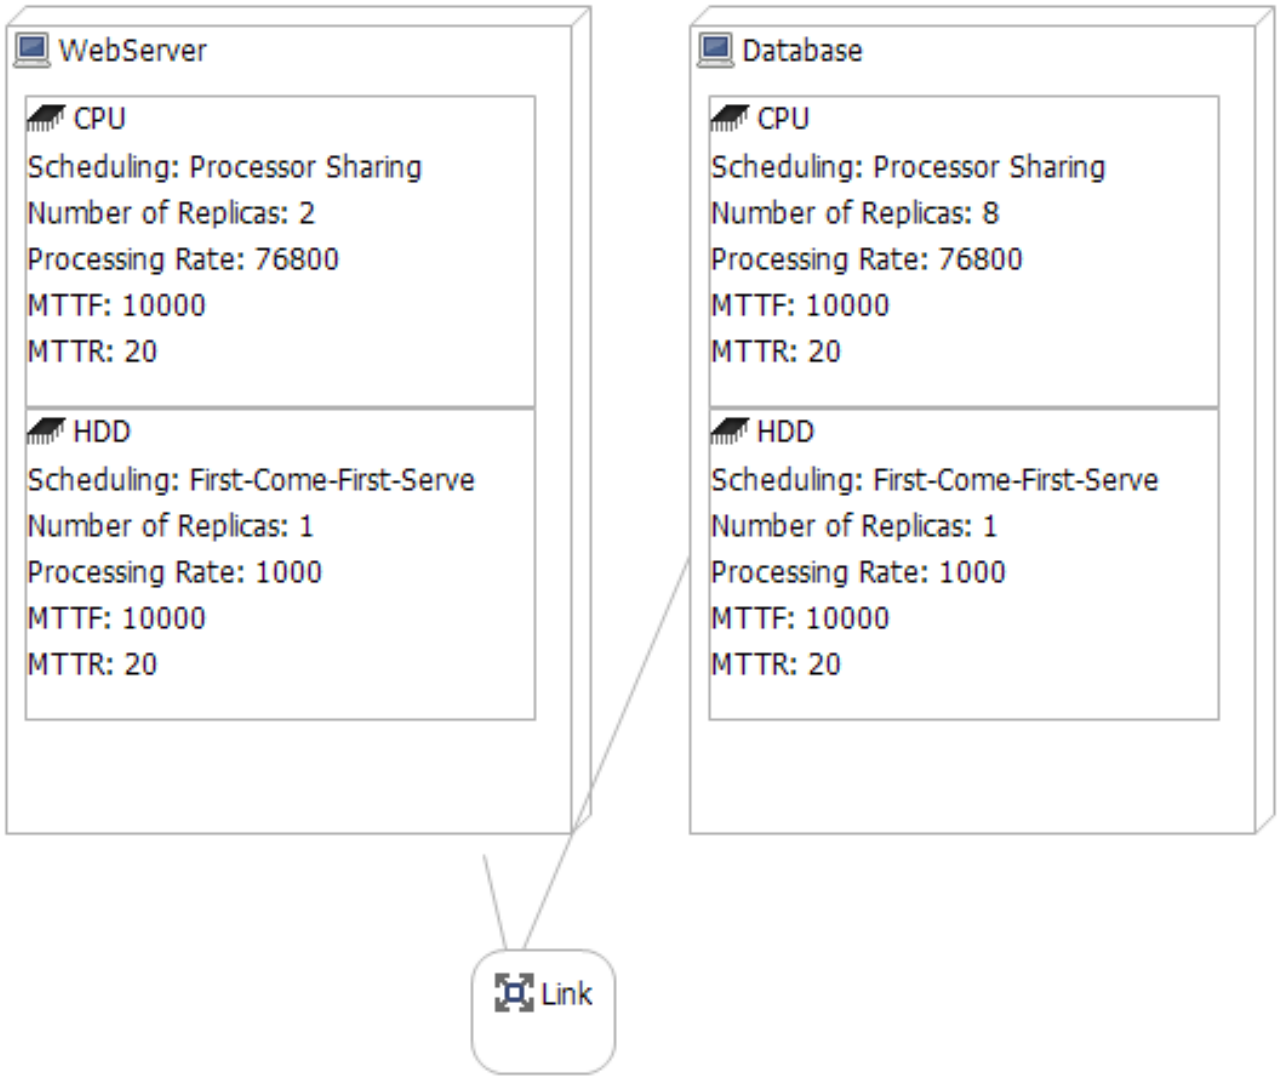

MODAClouds

SPECWeb running example

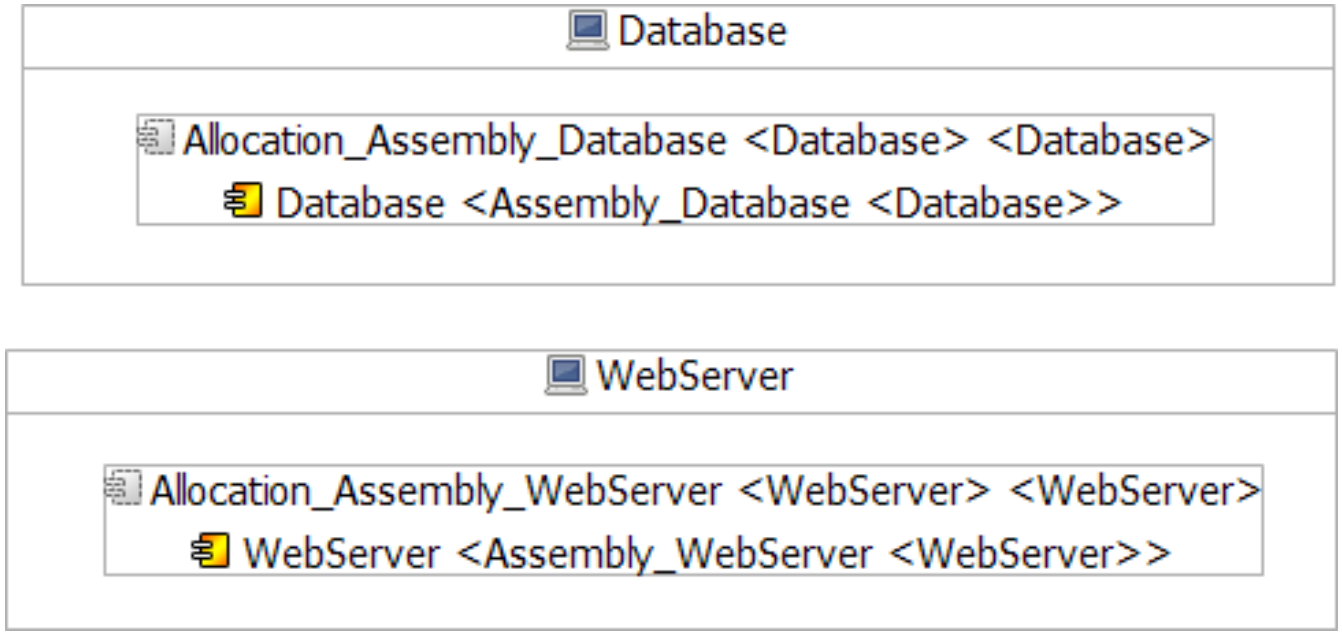

MODAClouds

SPECWeb running example

MODAClouds

Solution Outline

- The Palladio Framework allows to perform performance and cost analysis:
  - The Palladio Component Model (PCM) can be used as a CIM representing the application behavior independently from the hosting system
  - However, the framework does not support neither Cloud systems neither 24 hours analysis

- SPACE4CLOUD (Systems PerformAnce and Cost Evaluation for CLOUD) tool is intended to extend the Palladio Framework:
  - To support performance and cost evaluation of Clouds
  - Address explicitly the peculiarities of Clouds (workload fluctuations, burstiness, performance variability)
  - This is obtained leveraging the CPIM and CPSMs meta-models

MODAClouds

Cloud Provider Analysis

- In order to define the general CPIM and the specific CPSMs, for each considered provider we analysed:
  - Types and features of Cloud services
  - Pricing models
  - Scaling capabilities/features

Instance Type	Reservation Fee		Per Hour Fee	
	1 Year	3 Years	Linux/Unix	Windows
Micro Instance	$ 62	$ 100	$ 0.005	$ 0.011
Small Instance	$ 276.25	$ 425	$ 0.02	$ 0.04
Large Instance	$ 1105	$ 1700	$ 0.08	$ 0.16
Extra Large Instance	$ 2210	$ 3400	$ 0.16	$ 0.32
High-Memory Extra Large Instance	$ 1600	$ 2415	$ 0.114	$ 0.184
High-Memory Double Extra Large Instance	$ 3200	$ 4830	$ 0.227	$ 0.367
High-Memory Quadruple Extra Large Instance	$ 6400	$ 9660	$ 0.454	$ 0.734
High-CPU Medium Instance	$ 553	$ 850	$ 0.04	$ 0.105
High-CPU Extra Large Instance	$ 2210	$ 3400	$ 0.16	$ 0.42
Cluster Compute Quadruple Extra Large	$ 4060	$ 6300	$ 0.297	$ 0.477
Cluster Compute Eight Extra Large	$ 5000	$ 7670	$ 0.361	$ 0.571
Cluster GPU Quadruple Extra Large	$ 6830	$ 10490	$ 0.494	$ 0.794

MODAClouds

M
an

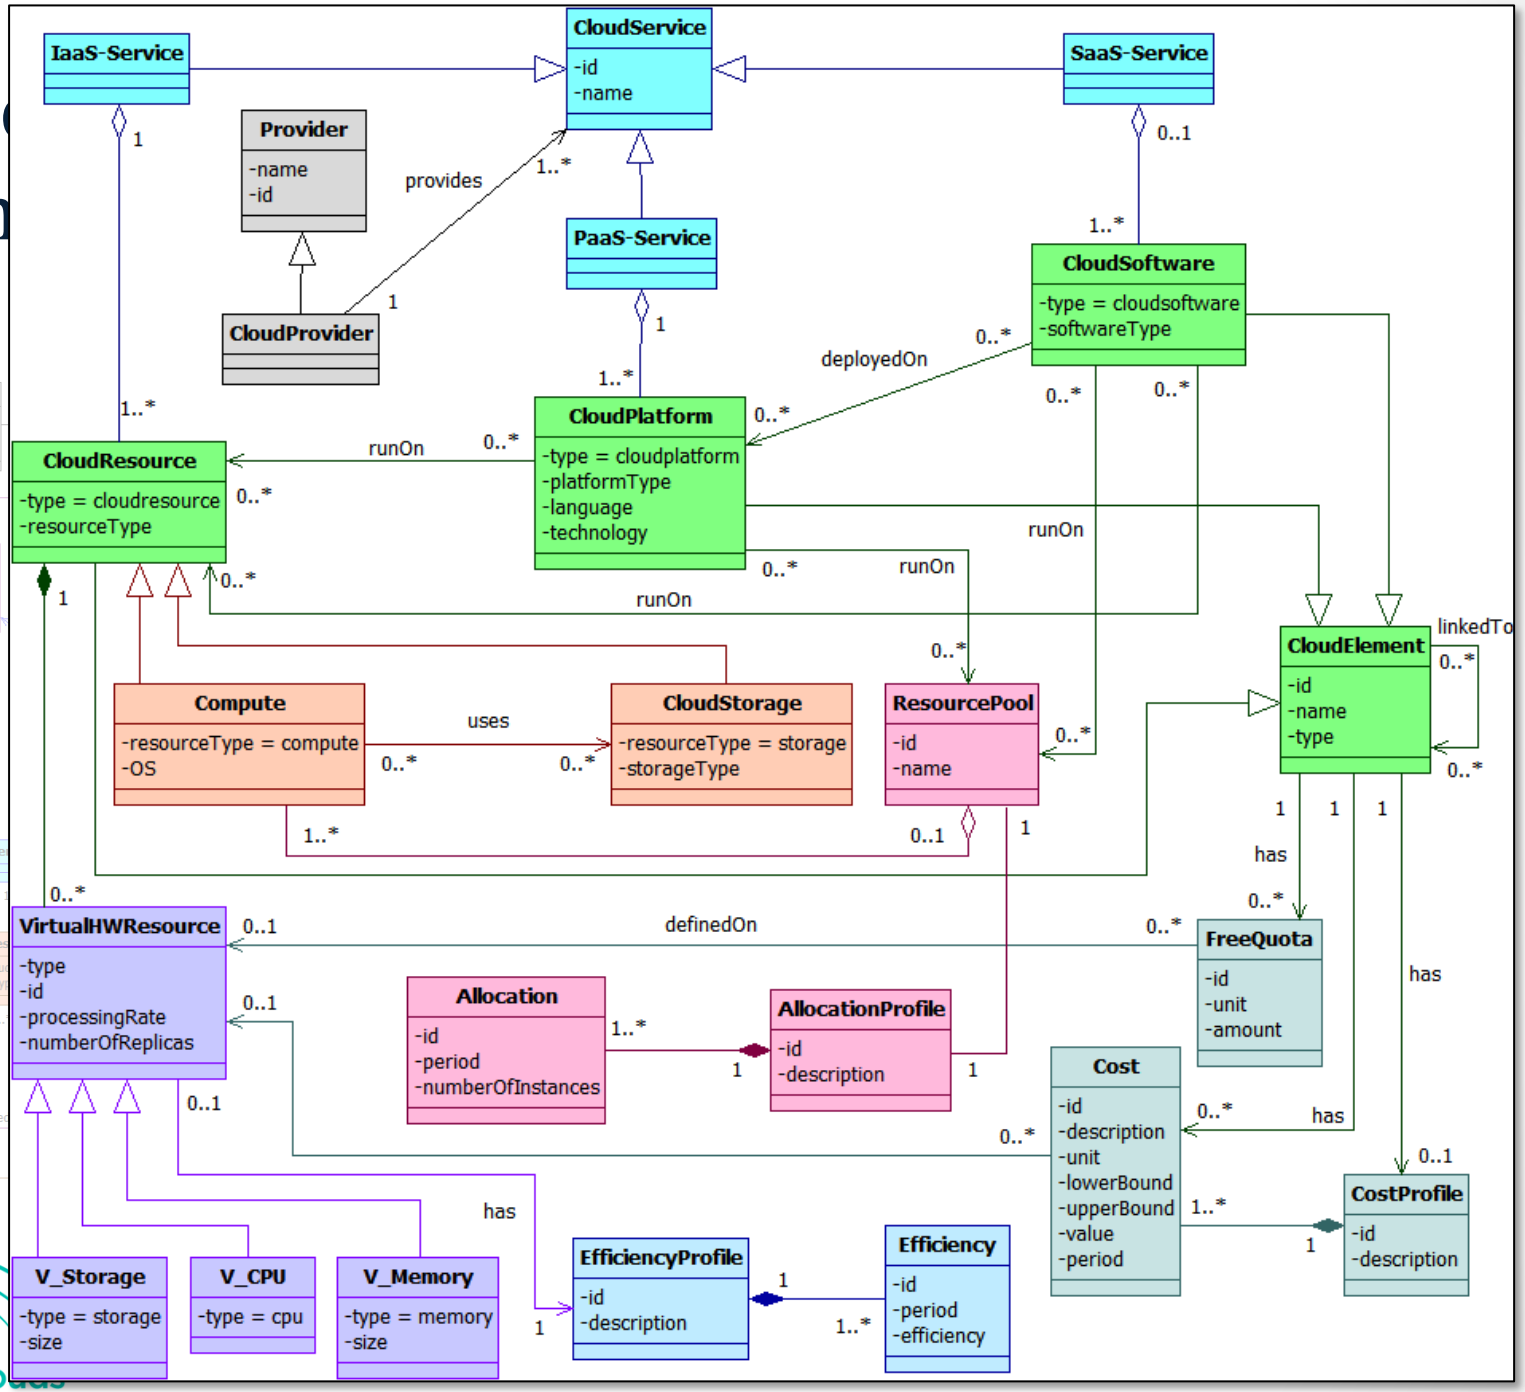

MODACloud

Meta-Models for Cloud Systems Performance and Cost Evaluation – Amazon CPSM (EC2)

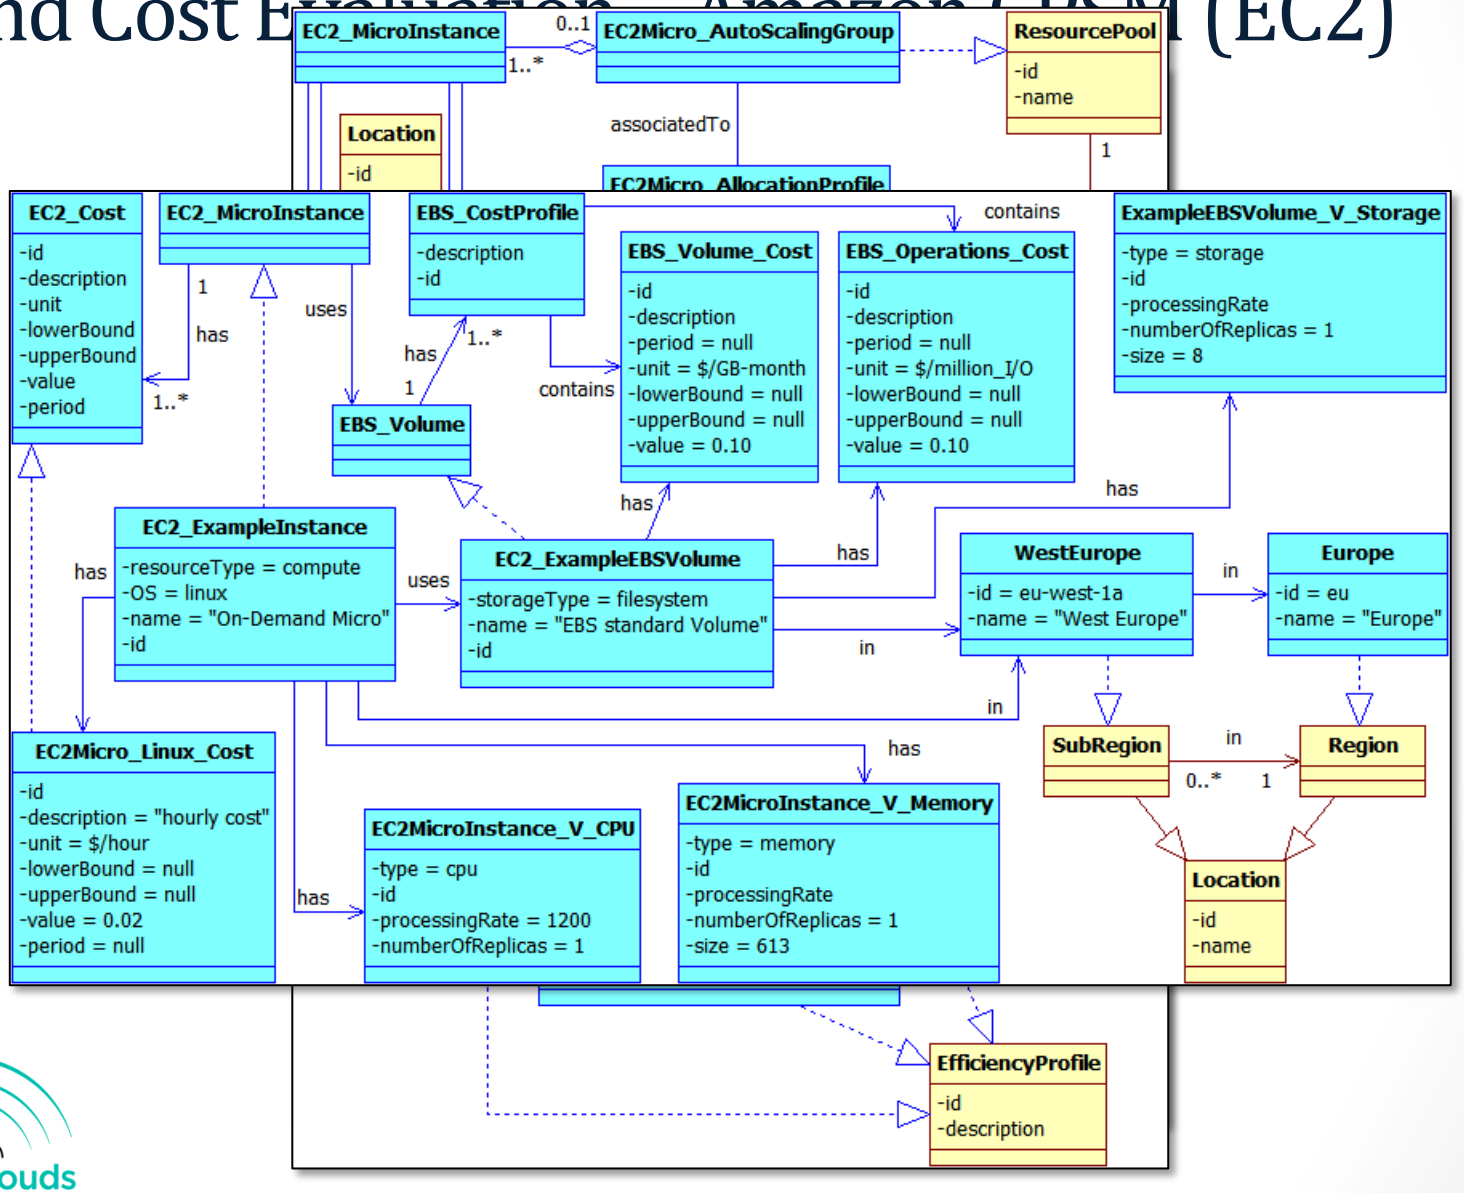

MODAClouds

Mapping the CPIM to the PCM

- In order to extend Palladio to use cloud resources as hosting systems, we needed:
  - A general CPIM definition
  - A specific CPSM definition for each considered provider
  - A mapping between the CPIM/CPSMs and the PCM. In particular, the mapping allows to represent cloud resources as Palladio processing resources

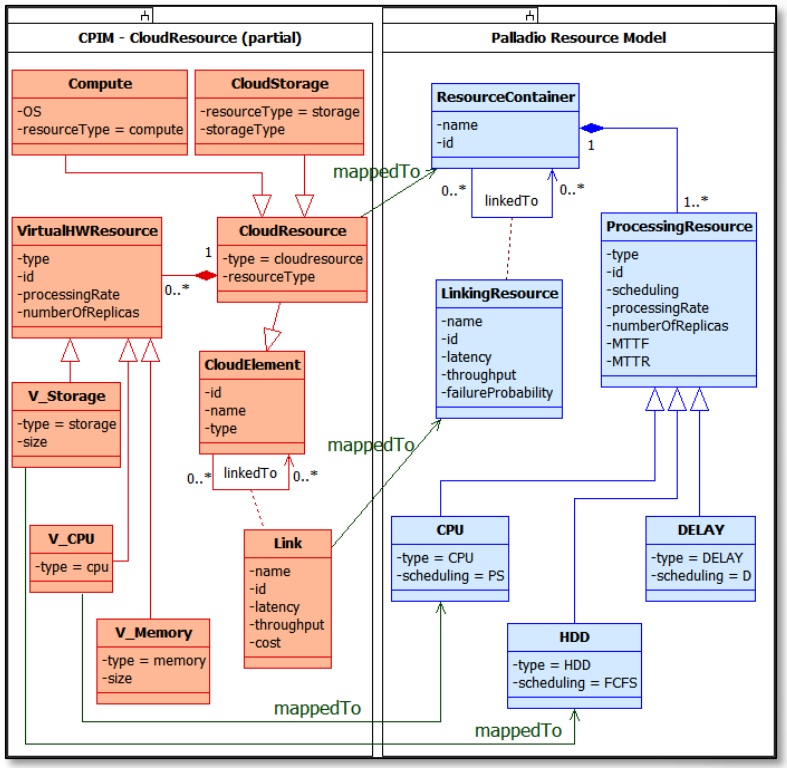

MODAClouds

Advanced Performance Features

- Random Environments:
  - Continuous-time Markov chain
  - Model systems jumping between stages characterizing system working condition (e.g., fast/slow)

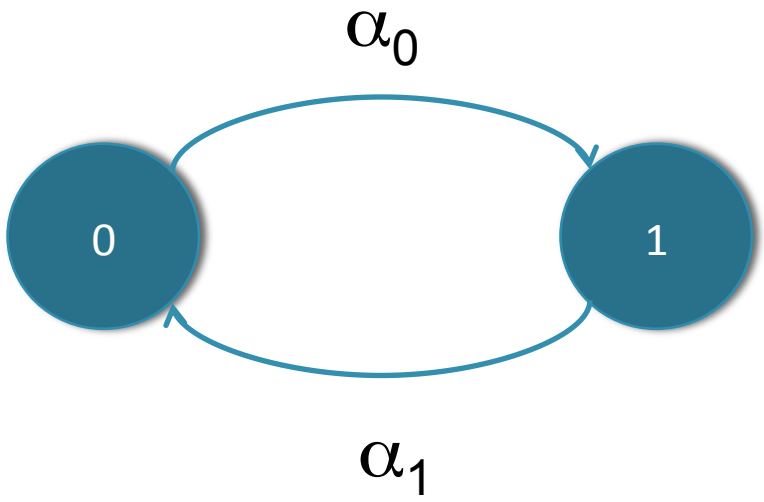

MODAClouds

Advanced Performance Features

- Random Environments:
  - Continuous-time Markov chain
  - Model systems jumping between stages characterizing system working condition (e.g., fast/slow)
- General service demand distribution:
  - Service times follow a Coxian distribution
  - Any distribution function can be approximated arbitrarily closely by a Coxian distribution

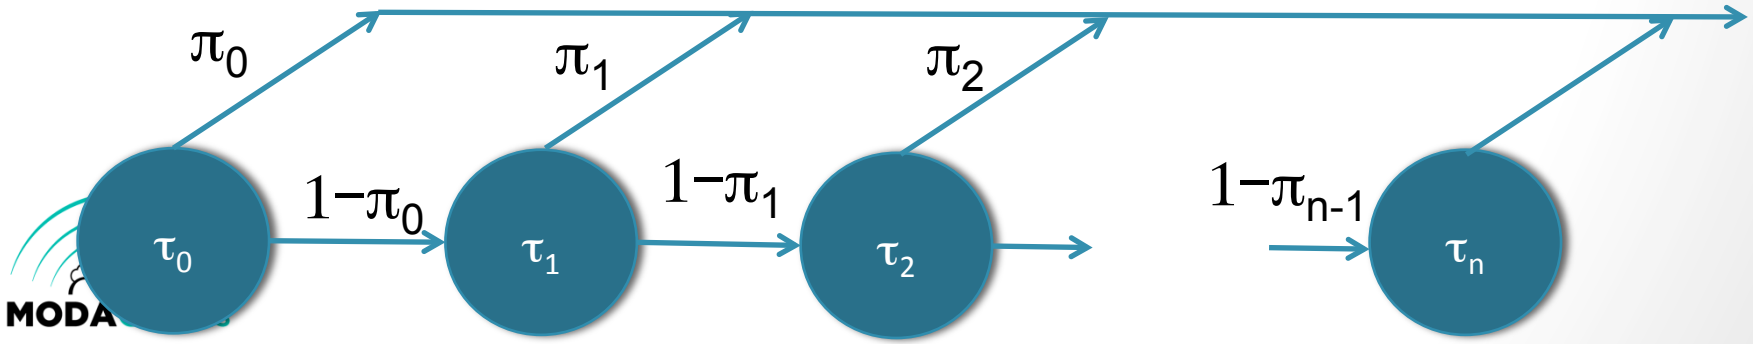

Advanced Performance Features

- Random Environments:
  - Continuous-time Markov chain
  - Model systems jumping between stages characterizing system working condition (e.g., fast/slow)
- General service demand distribution:
  - Service times follow a Coxian distribution
  - Any distribution function can be approximated arbitrarily closely by a Coxian distribution

Solution based on the “Blending Algorithm”

G. Casale, M. Tribastone. Modelling exogenous variability in cloud deployments. SIGMETRICS Performance Evaluation Review 40(4): 73-82 (2013)

MODAClouds

Cloud Providers Comparison

- Leverage SPACE4CLOUD to compare Amazon and Flexiscale by:
  1. Choosing similar machines for the SPECWeb components
  2. Using realistic workload and allocation profiles

MODAClouds

Cloud Providers Comparison

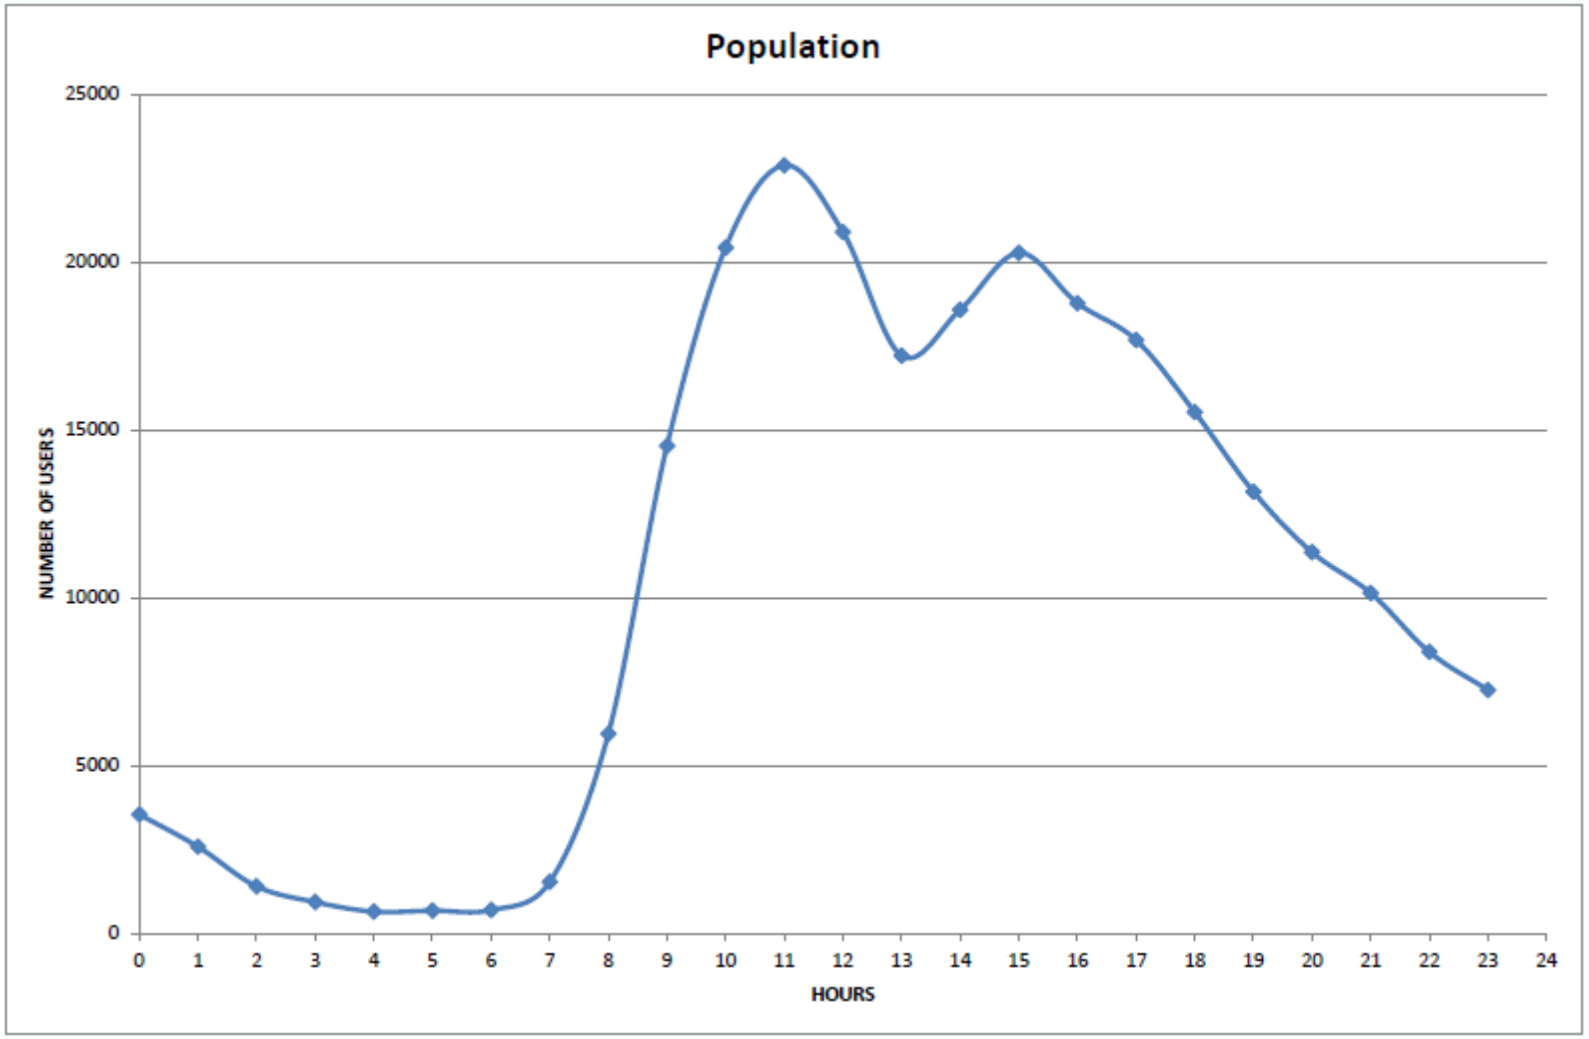

MODAClouds

Cloud Providers Comparison

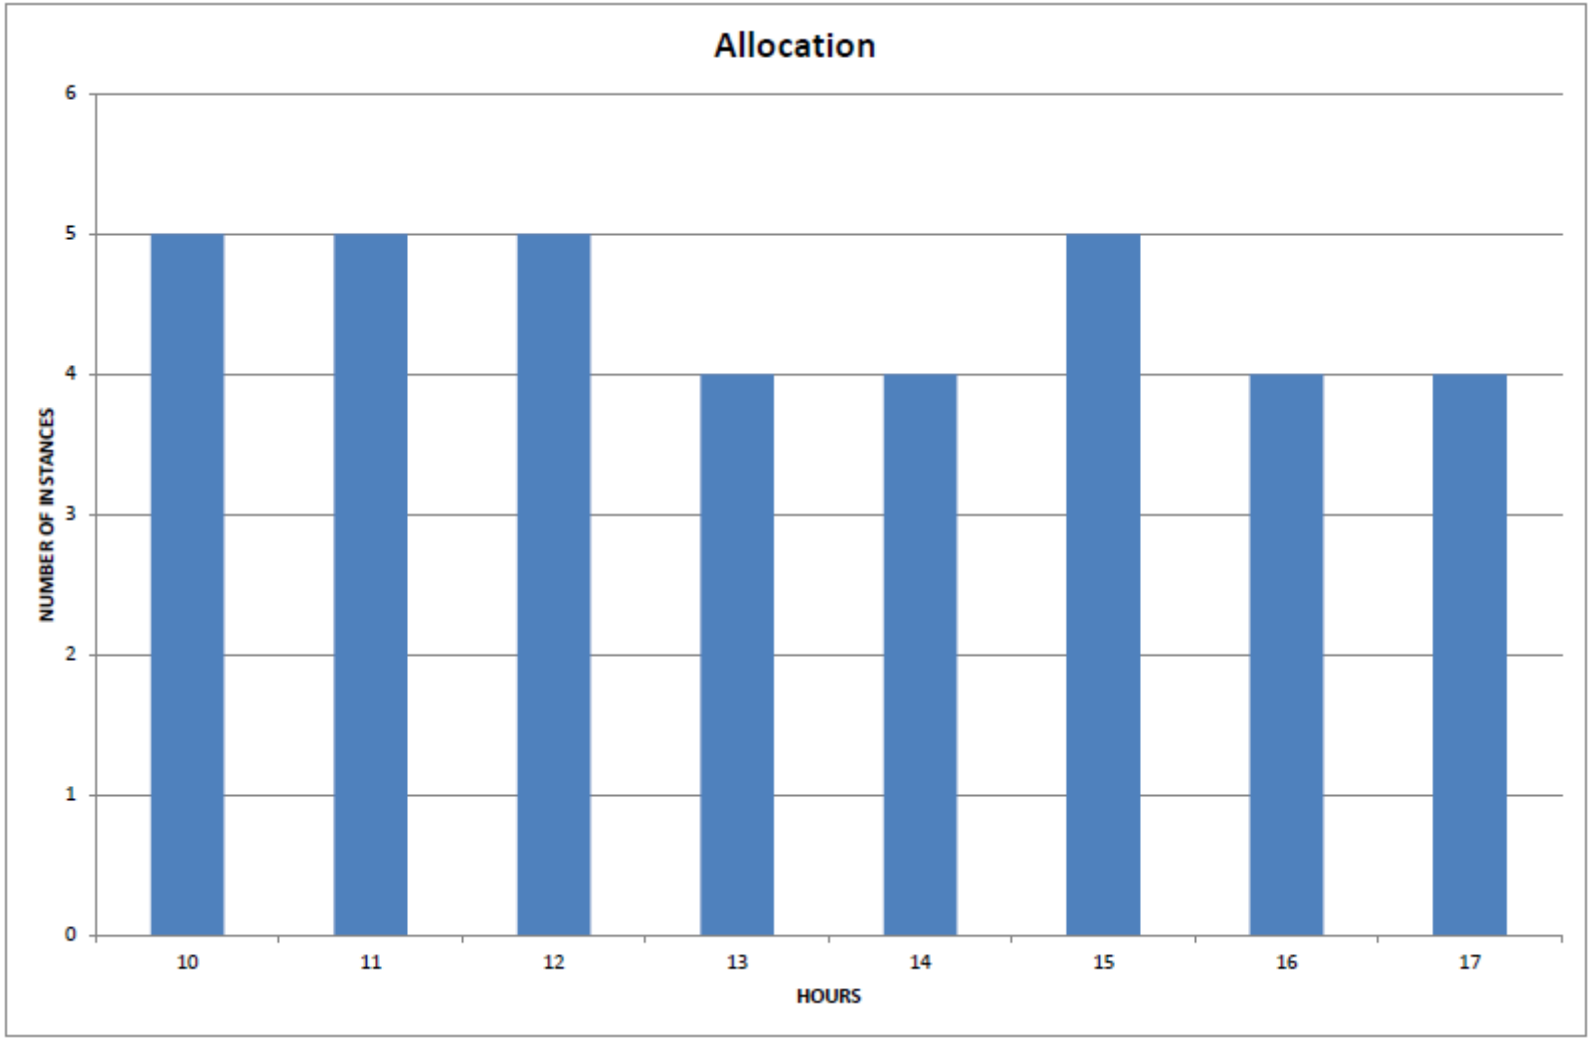

MODAClouds

Cloud Providers Comparison - Results

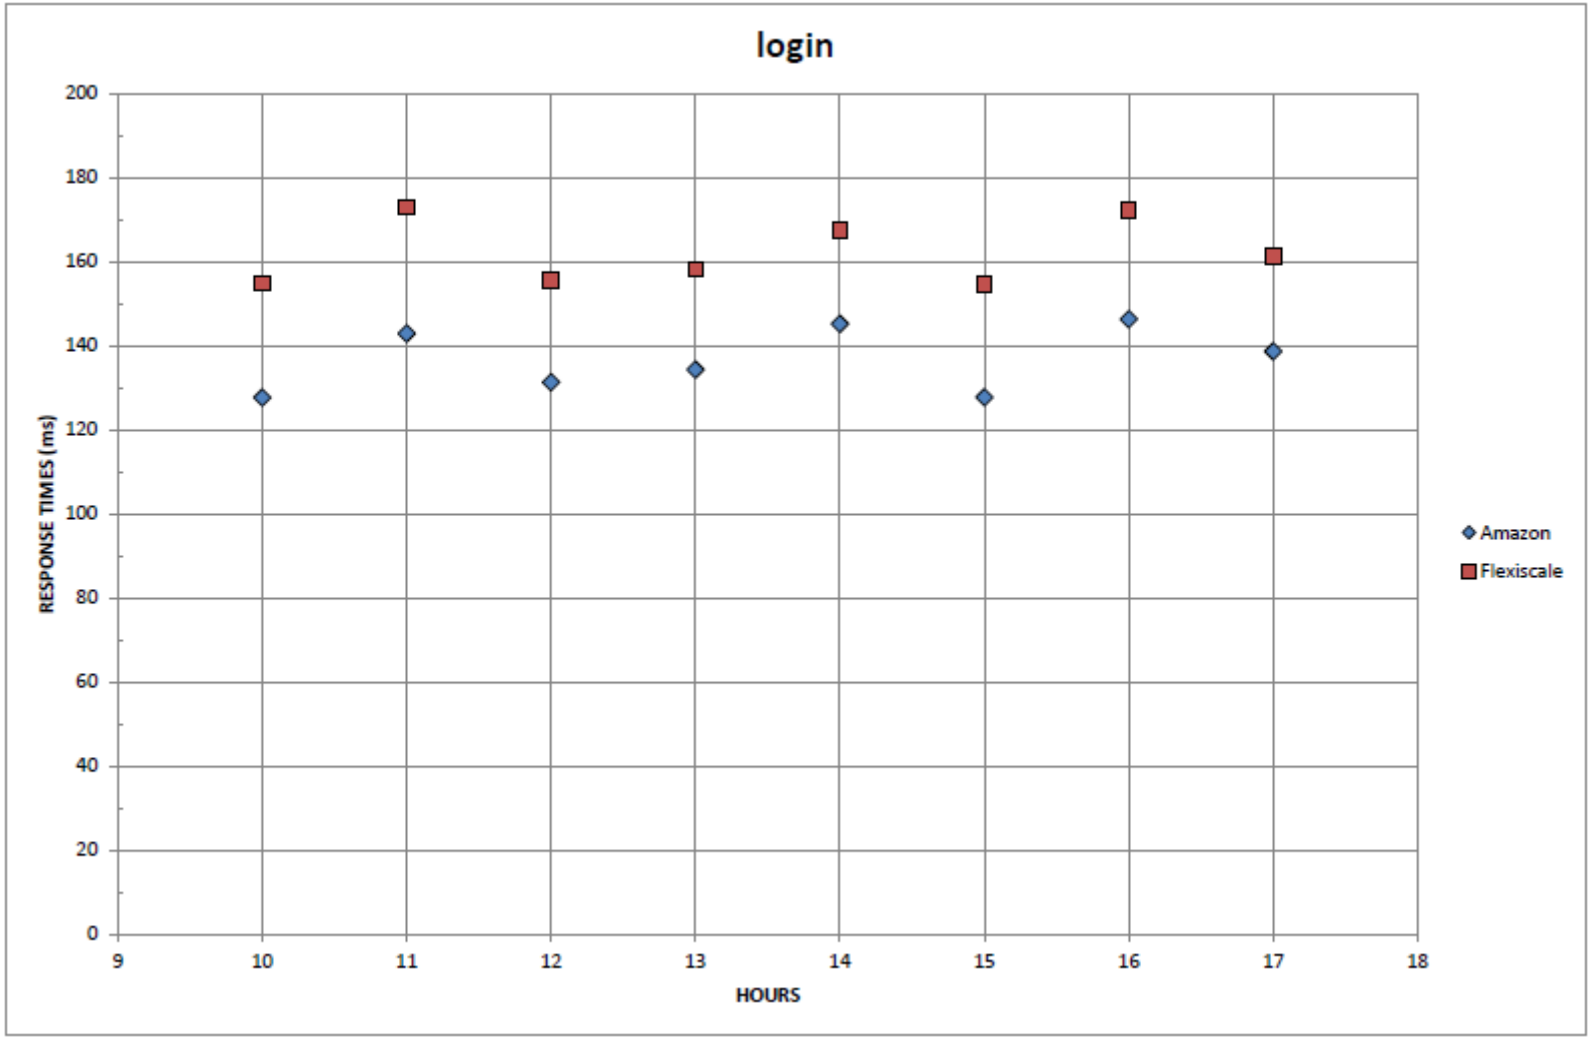

MODAClouds

Cloud Providers Comparison - Results

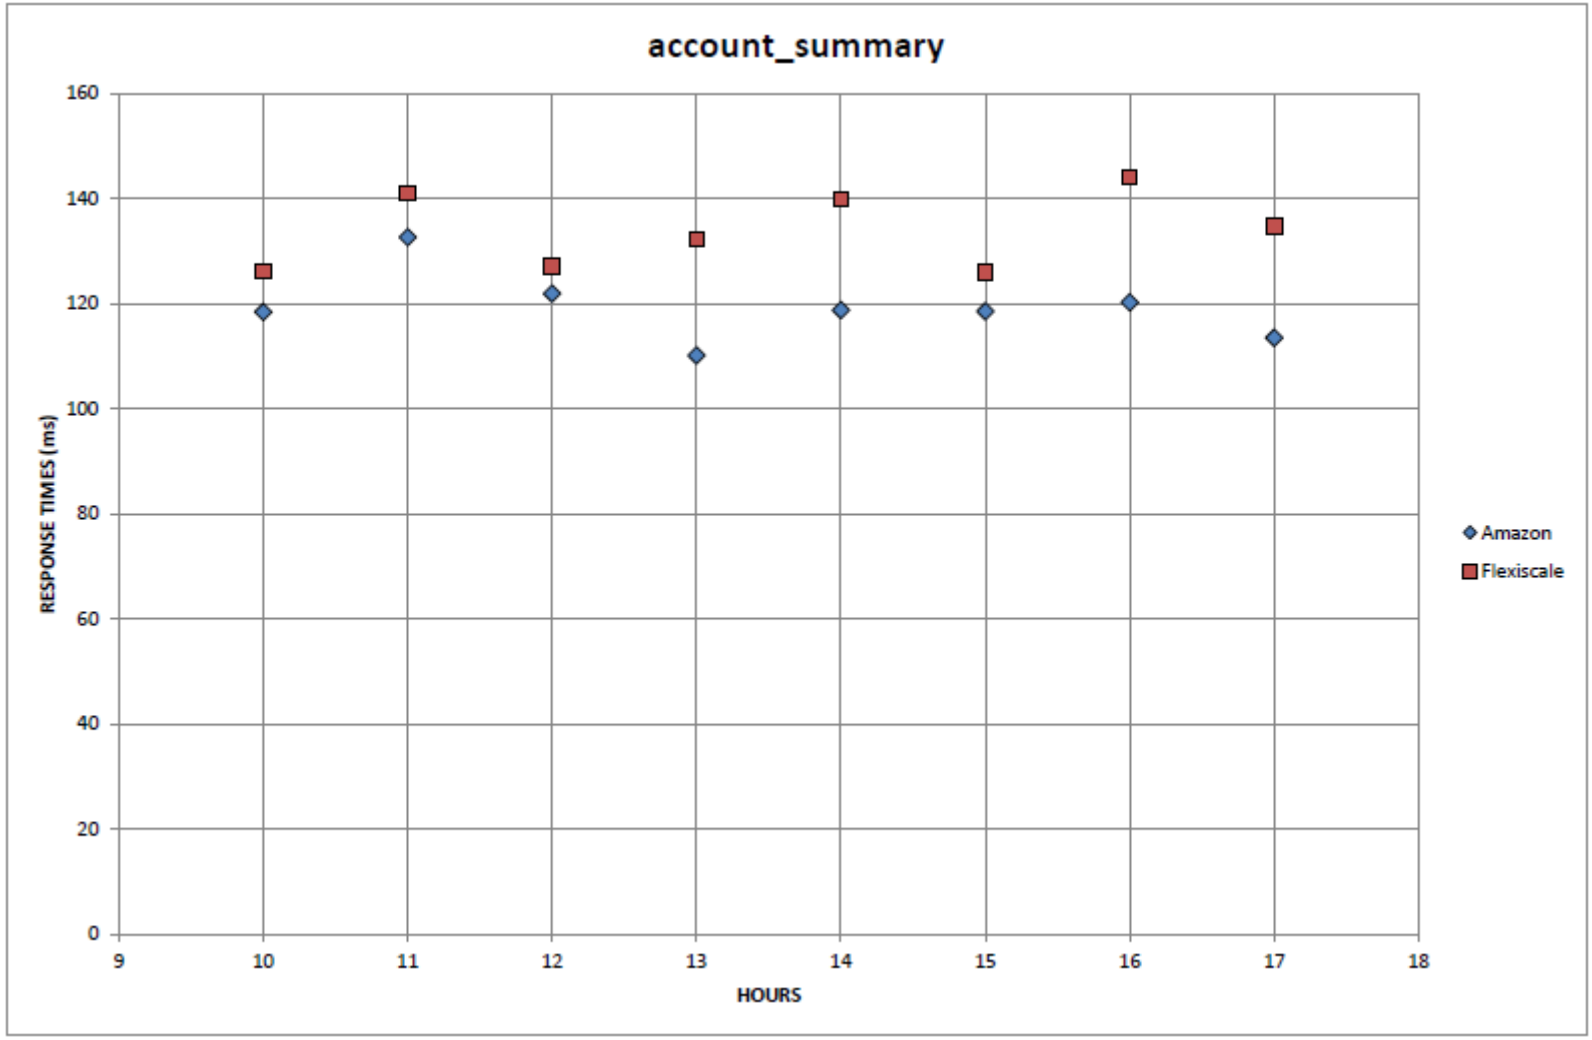

MODAClouds

Cloud Providers Comparison - Results

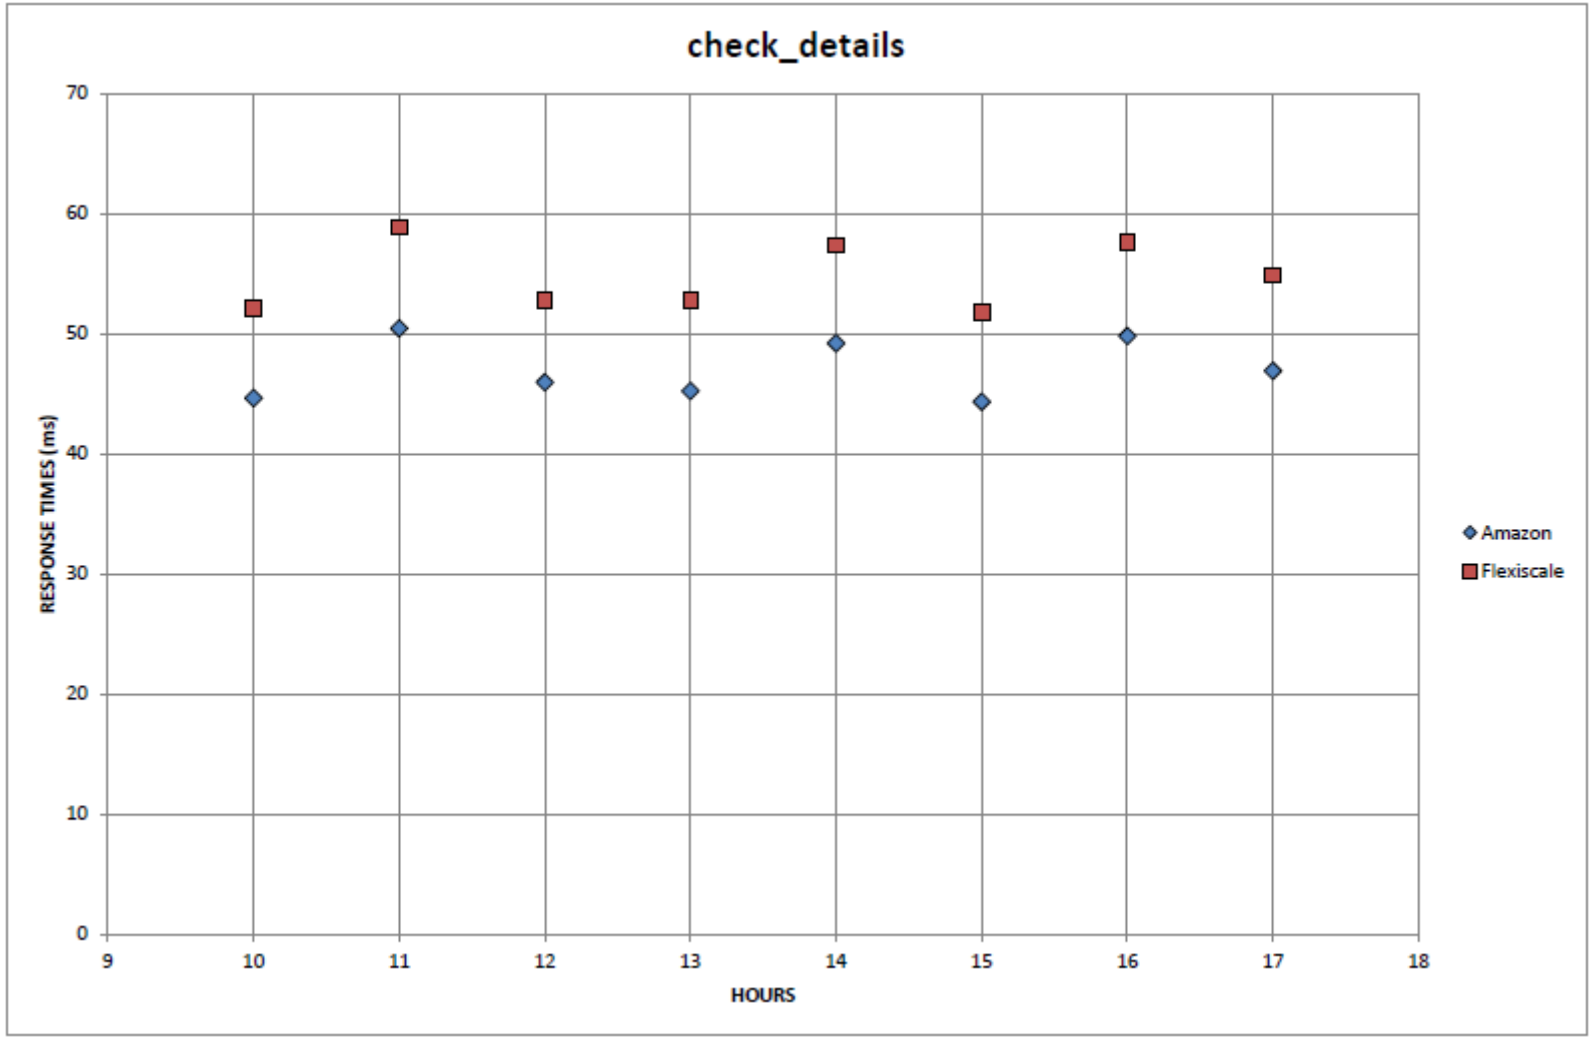

MODAClouds

Cloud Providers Comparison - Results

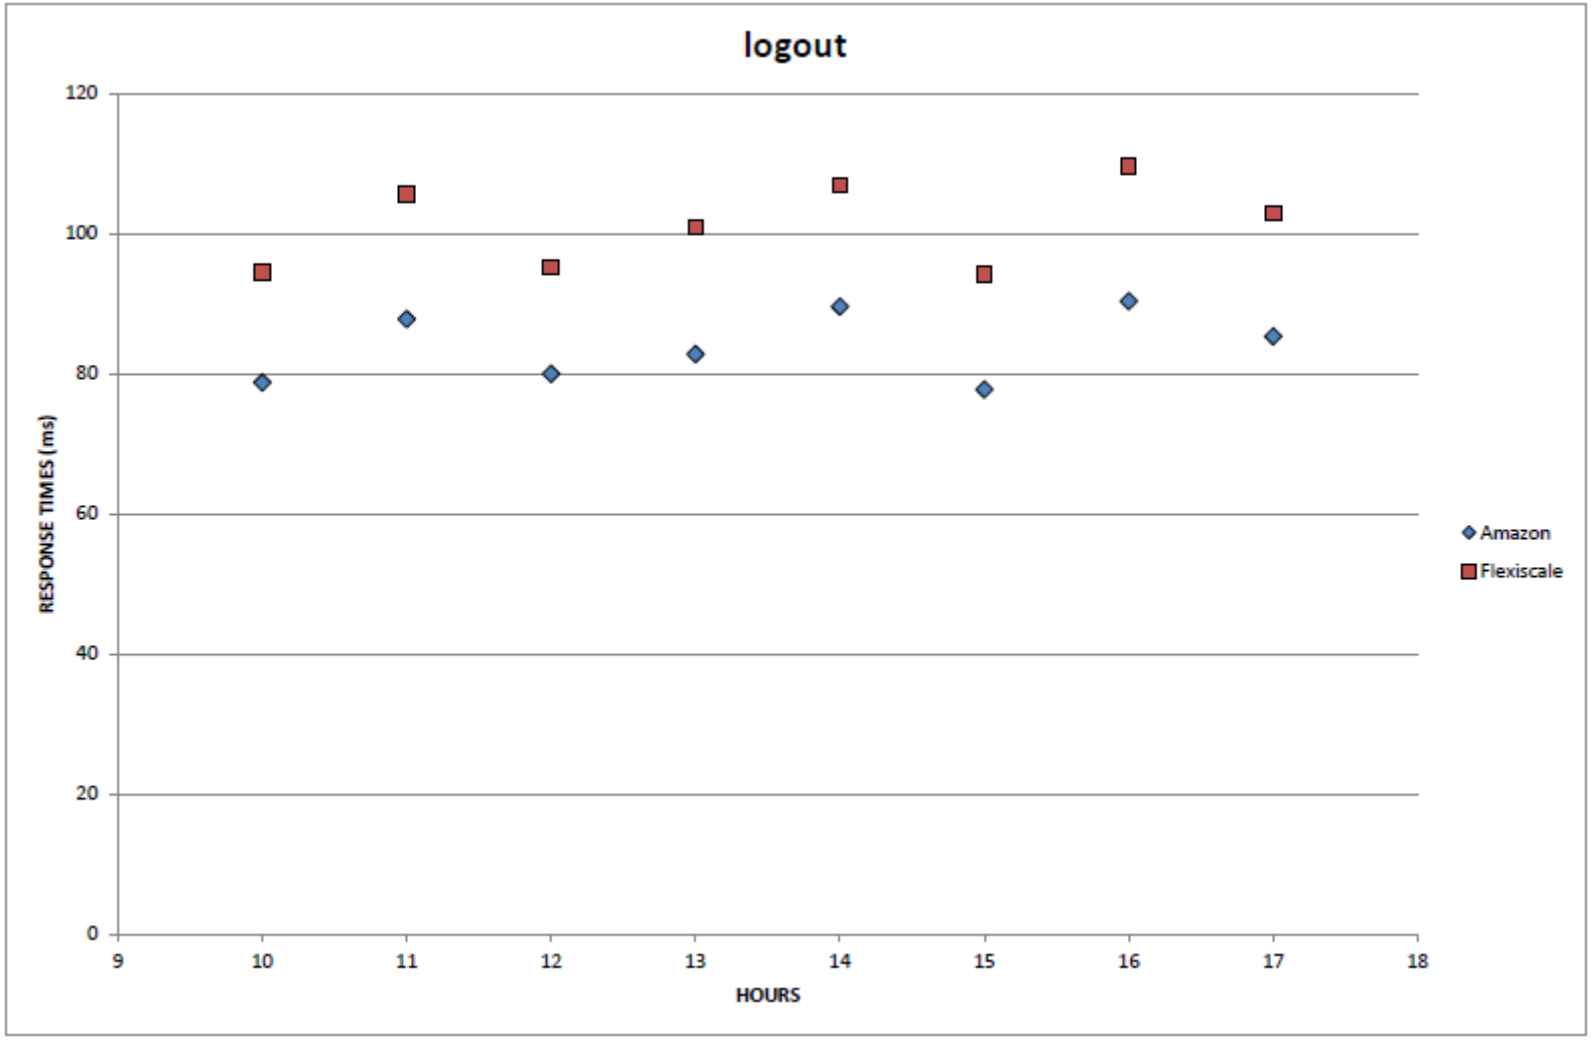

MODAClouds

Cloud Providers Comparison - Results

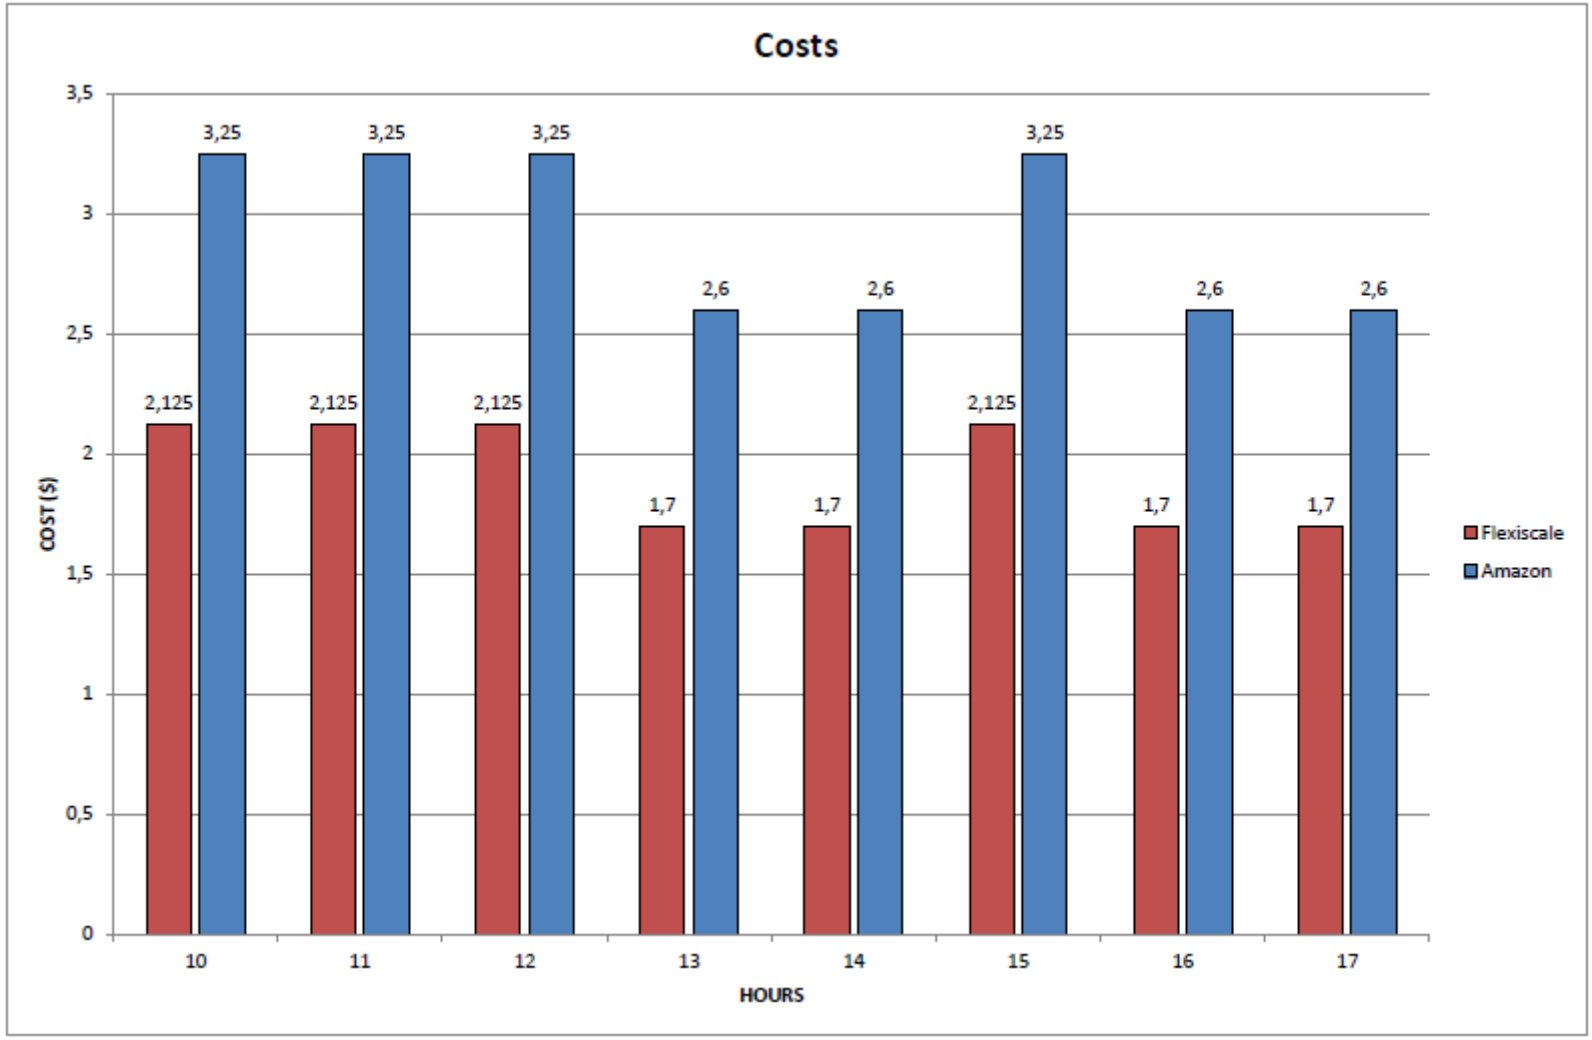

MODAClouds

Cloud Providers Comparison - Results

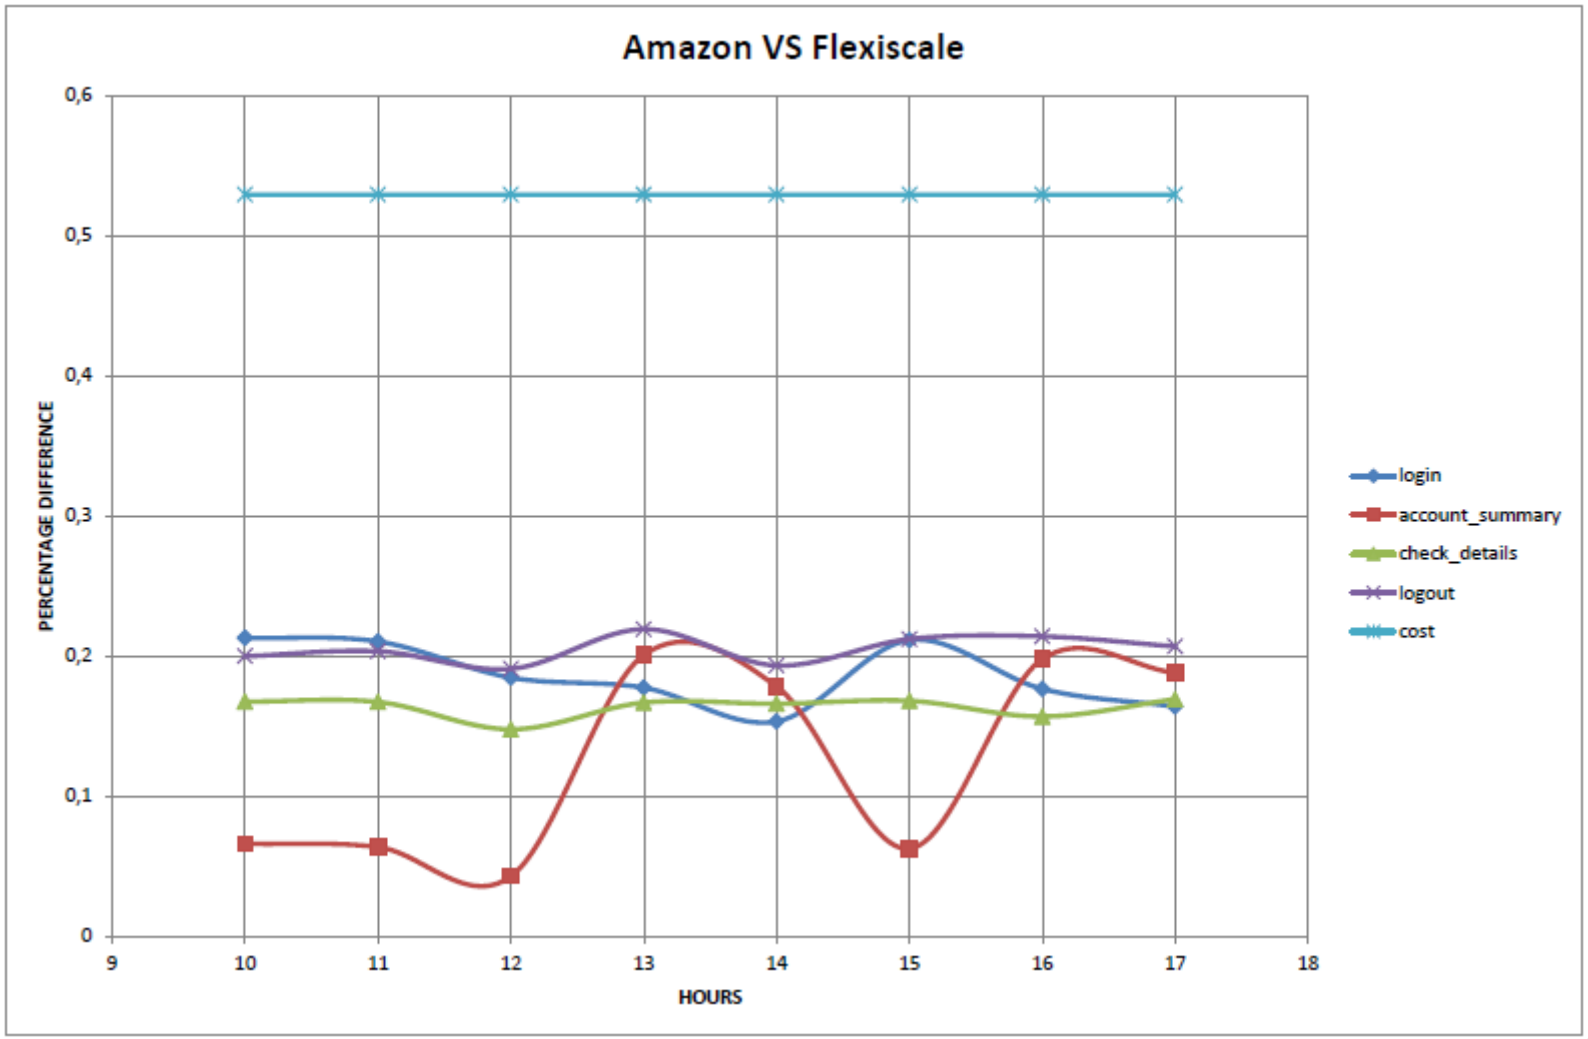

MODAClouds

Conclusions and Future Work

- Model-driven approach for the design of Cloud applications, taking into account non-functional requirements like costs and performance
- Meta-models have been integrated with the existing performance and cost evaluation tools extending their analysis capabilities to Cloud systems
- Include system availability in design-time analysis
- Develop a local search to support design-time exploration and costs minimization

MODAClouds

Thanks for your attention...

...any questions?

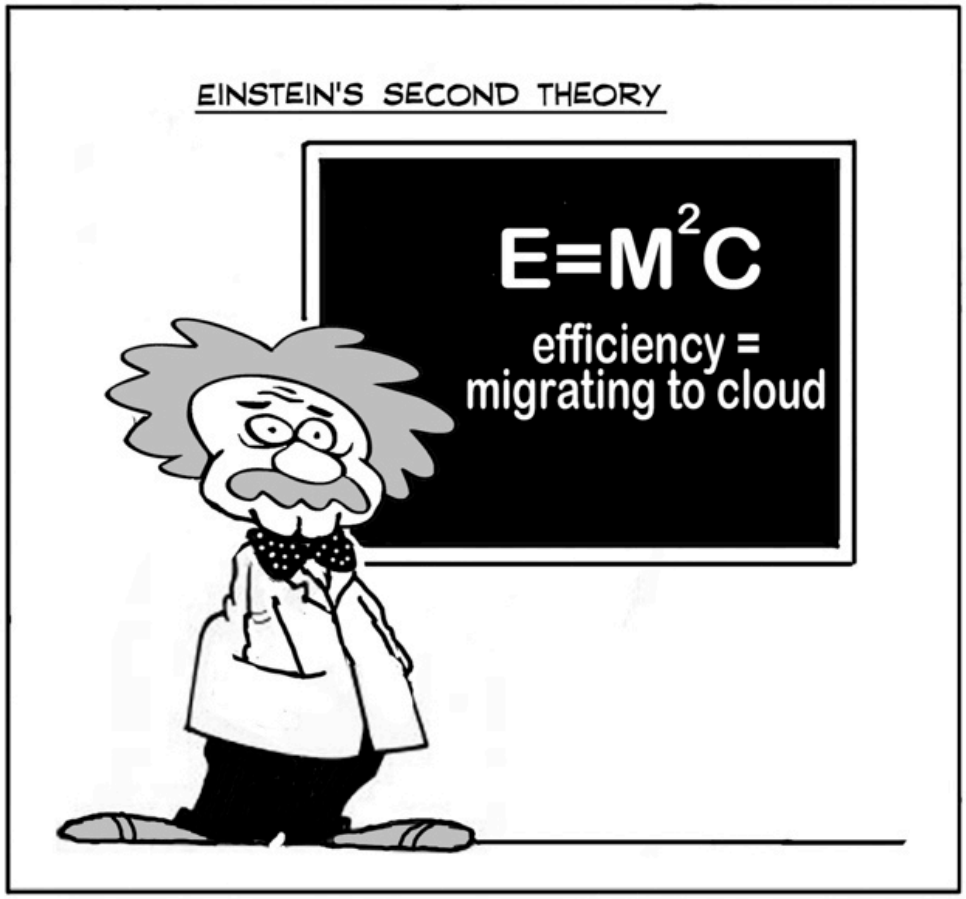

MODAClouds

References

- D. Ardagna, E. Di Nitto, D. Petcu, P. Mohagheghi, S. Mosser, P. Matthews, A. Gericke, C. Ballagny, F. D'Andria, C. Nechifor, C. Sheridan. **MODAClouds: A Model-Driven Approach for the Design and Execution of Applications on Multiple Clouds.** *MiSE 2012 Workshops Proceedings*
- D. Franceschelli, D. Ardagna, M. Ciavotta, E. Di Nitto. **SPACE4CLOUD: A Tool for System PerformAnce and Cost Evaluation of CLOUD Systems**. *Multi-Cloud 2013 Workshop Proceedings*. 27-34. Prague, Czech Republic
- G. Casale, M. Tribastone. **Modelling exogenous variability in cloud deployments.** *SIGMETRICS Perform. Eval. Rev. 40 (4) 73-82*
- G. Casale, M. Tribastone. **Fluid Analysis of Queueing in Two-Stage Random Environments**. *QEST 2011: 21-30*

MODAClouds