Amazon Elastic Compute Cloud (EC2) vs. in-House HPC Platform

a Cost Analysis

J. Emeras, S. Varrette, P. Bouvry

Parallel Computing and Optimization Group (PCOG),

University of Luxembourg (UL), Luxembourg

Summary

Summary

Overview

- Plethora of Cloud Services
  - ↪ **Compute:** Elastic Compute Cloud (EC2) / Docker containers
  - ↪ **Storage:** Block Storage (EBS) / Simple Storage Service (S3) / Glacier (archiving)
  - ↪ **Networking:** Virtual Private Cloud (VPC) / Route 53 (DNS)
  - ↪ **Database:** Relational Database (RDS) / DynamoDB (NoSQL)
  - ↪ **Analytics:** Hadoop / Machine Learning...
  - ↪ etc.

⇒ Of main interest to mimic HPC facility: **EC2**, EBS, VPC

- Elastic Compute Cloud (EC2) Instances
  - ↪ Grouped by **Region** and **Availability Zone** (AZ)
  - ↪ VMs organized in *families*
  - ↪ Pay per VM according to time used

EC2 VMs Organization

Instance Family	Instance Type	Processor Microarchitecture	Introduction Date
General Purpose	m1	Xeon Family	2006-08-23
	m3	Ivy Bridge-EP	2012-10-31
	t2	Xeon Family	2014-07-01
	m4	Haswell-EP	2015-06-11
Memory Optimized	m2	Xeon Family	2010-02-22
	cr1	Sandy Bridge-EP	2013-01-21
	r3	Ivy Bridge-EP	2014-04-10
Compute Optimized	c1	Xeon Family	2008-08-08
	cc1	Nehalem-EP	2010-07-13
	cc2	Sandy Bridge-EP	2011-11-14
	c3	Ivy Bridge-EP	2013-11-14
	c4	Haswell-EP	2014-11-13
Storage Optimized	hi1	Xeon Family	2012-07-18
	hs1	Sandy Bridge-EP	2012-12-21
	i2	Ivy Bridge-EP	2013-12-20
Dense Storage	d2	Haswell-EP	2015-03-30
GPU	cg1	Nehalem-EP	2010-11-14
	g2	Sandy Bridge-EP	2013-11-05
Micro	t1	Xeon Family	2009-10-26

EC2 Pricing

- Depends on the region, instance type, options and pricing mode
- 3 payment modes:
  1. **On-Demand:** Fixed hourly rate. Most expensive / flexible
  2. **Reserved:** Fixed hourly rate with several upfront options. Contracts on 1 or 3 years term. Less expensive than OD / not flexible
     - ✓ no upfront
     - ✓ partial upfront
     - ✓ all upfront
  3. **Spot:** Price bidding / flexible. Fixed duration option (1 / 6 hours)

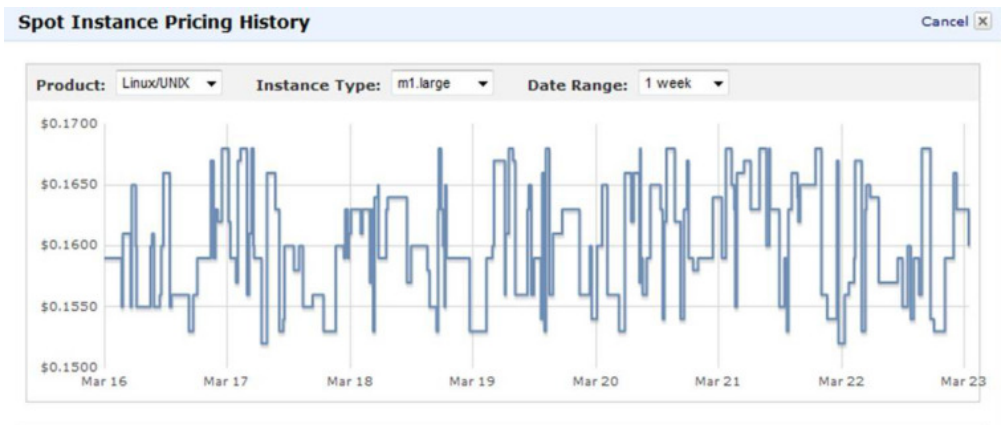

EC2 for High Performance Computing

- **Useful Features**
  - ↪ Enhanced networking (SR-IOV)
  - ↪ Placement groups
  - ↪ EBS-optimized (increased throughput)
  - ↪ Dedic. instances (no multi-tenancy)
- **Possible Instance Families**
  - ↪ Compute: *c3*, *c4*
  - ↪ Memory: *r3*
  - ↪ Storage: *i2*
  - ↪ Dense Storage: *d2*
  - ↪ GPU: *g2* (no SR-IOV)

EC2 for High Performance Computing

- **Useful Features**
  - ↪ Enhanced networking (SR-IOV)
  - ↪ Placement groups
  - ↪ EBS-optimized (increased throughput)
  - ↪ Dedic. instances (no multi-tenancy)

- **Possible Instance Families**
  - ↪ Compute: *c3*, *c4*
  - ↪ Memory: *r3*
  - ↪ Storage: *i2*
  - ↪ Dense Storage: *d2*
  - ↪ GPU: *g2* (no SR-IOV)

On-Demand Instance Prices

Linux | RHEL | SLES | Windows | Windows with SQL Standard | Windows with SQL Web | Windows with SQL Enterprise

Region: EU (Ireland)

	vCPU	ECU	Memory (GiB)	Instance Storage (GB)	Linux/UNIX Usage
General Purpose - Current Generation					
t2.micro	1	Variable	1	EBS Only	$0.014 per Hour
t2.small	1	Variable	2	EBS Only	$0.028 per Hour
t2.medium	2	Variable	4	EBS Only	$0.056 per Hour
t2.large	2	Variable	8	EBS Only	$0.112 per Hour
m4.large	2	6.5	8	EBS Only	$0.139 per Hour
m4.xlarge	4	13	16	EBS Only	$0.278 per Hour
m4.2xlarge	8	26	32	EBS Only	$0.556 per Hour
m4.4xlarge	16	53.5	64	EBS Only	$1.112 per Hour
m4.10xlarge	40	124.5	160	EBS Only	$2.78 per Hour

Computing Performance

- Measured in **ECUs** (black-box)
  - ↪ EC2 Compute Unit
  - ↪ currently $\simeq$ 1 Xeon core @ 1GHz

EC2 for High Performance Computing

- **Useful Features**
  - ↪ Enhanced networking (SR-IOV)
  - ↪ Placement groups
  - ↪ EBS-optimized (increased throughput)
  - ↪ Dedic. instances (no multi-tenancy)

- **Possible Instance Families**
  - ↪ Compute: *c3*, *c4*
  - ↪ Memory: *r3*
  - ↪ Storage: *i2*
  - ↪ Dense Storage: *d2*
  - ↪ GPU: *g2* (no SR-IOV)

On-Demand Instance Prices

Linux | RHEL | SLES | Windows | Windows with SQL Standard | Windows with SQL Web | Windows with SQL Enterprise

Region: EU (Ireland)

	vCPU	ECU	Memory (GiB)	Instance Storage (GB)	Linux/UNIX Usage
General Purpose - Current Generation					
t2.micro	1	Variable	1	EBS Only	$0.014 per Hour
t2.small	1	Variable	2	EBS Only	$0.028 per Hour
t2.medium	2	Variable	4	EBS Only	$0.056 per Hour
t2.large	2	Variable	8	EBS Only	$0.112 per Hour
m4.large	2	6.5	8	EBS Only	$0.139 per Hour
m4.xlarge	4	13	16	EBS Only	$0.278 per Hour
m4.2xlarge	8	26	32	EBS Only	$0.556 per Hour
m4.4xlarge	16	53.5	64	EBS Only	$1.112 per Hour
m4.10xlarge	40	124.5	160	EBS Only	$2.78 per Hour

Computing Performance

- Measured in **ECUs** (black-box)
  - ↪ EC2 Compute Unit
  - ↪ currently $\simeq$ 1 Xeon core @ 1GHz

$\Rightarrow$ ECU vs *traditional* HPC metrics?

EC2 Computing Performances

- **Strong linear relationship between ECUs and GFLOPs**
  - ↪ **adj. R^2=0.9**
  - ↪ Other explanatory attributes (e.g. processor generation).

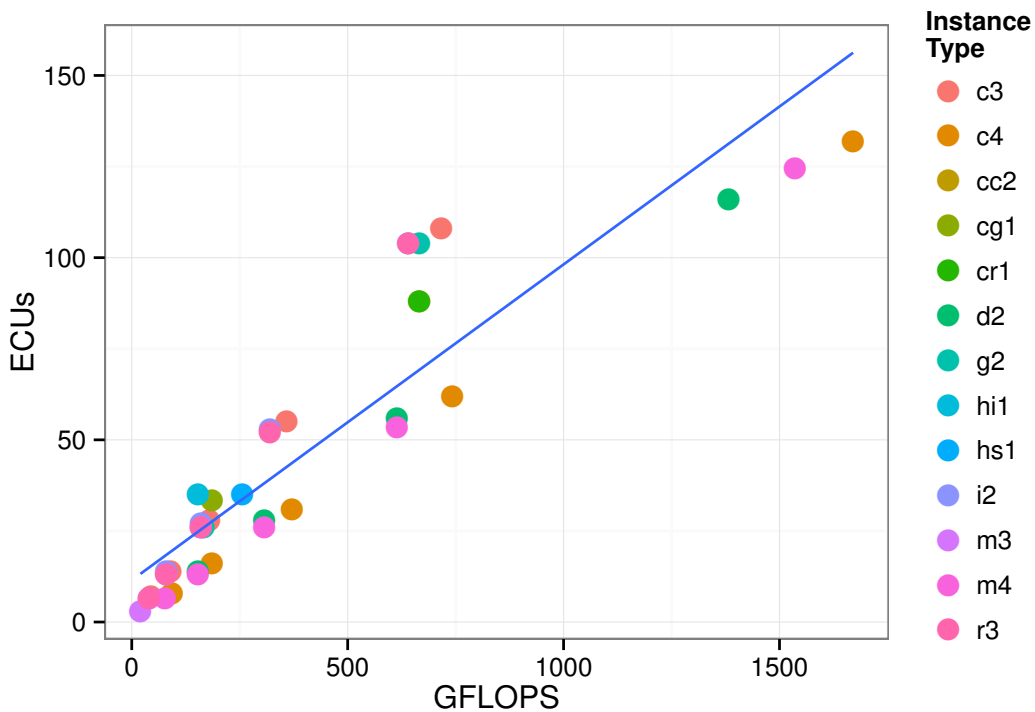

Summary

High Performance Computing @ UL

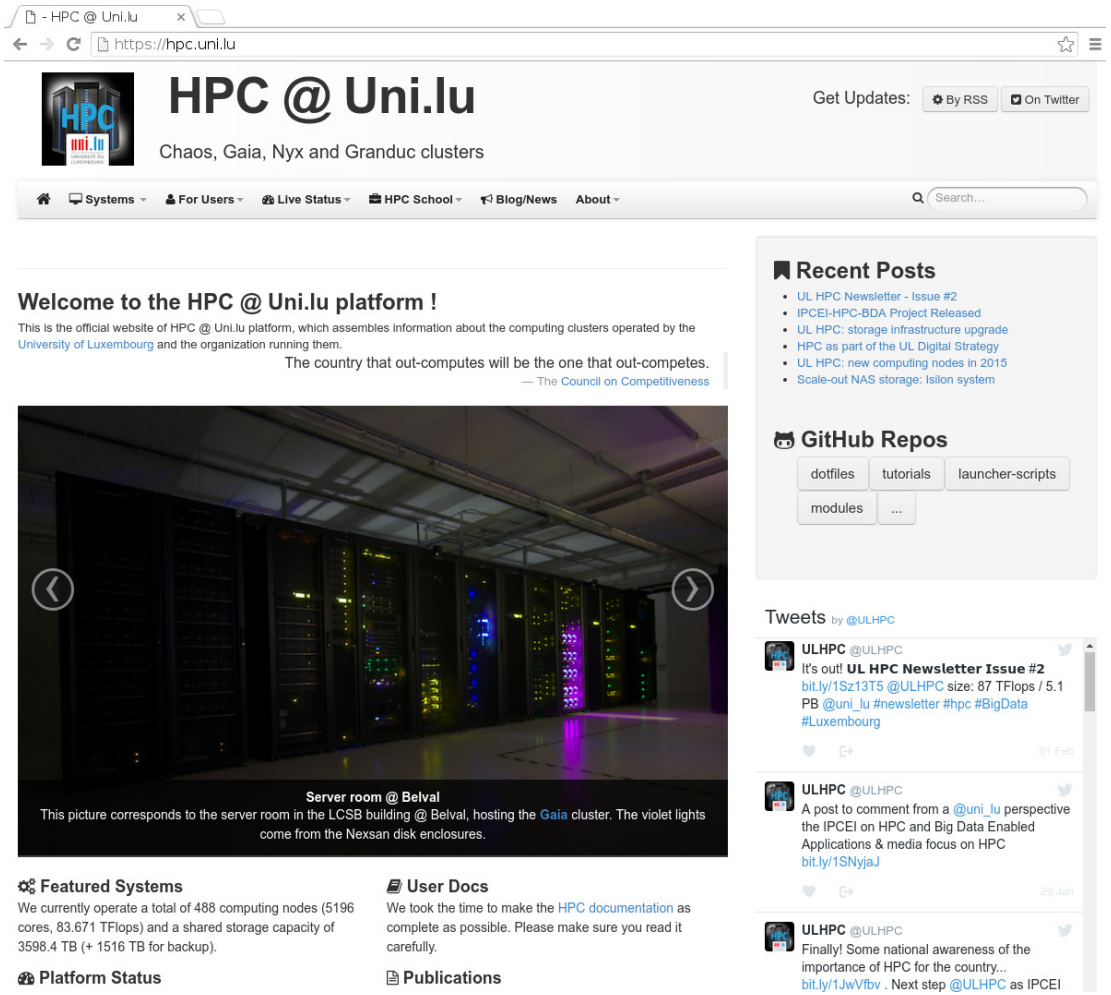

Key numbers

- 344 users
- 98 servers
- 492 nodes
  - ↪ 5300 cores
  - ↪ **85.543 TFlops**
- **5354.4 TB**
- 4 sysadmins
- 2 sites
  - ↪ Kirchberg
  - ↪ Belval

http://hpc.uni.lu

Local HPC Platform Investment

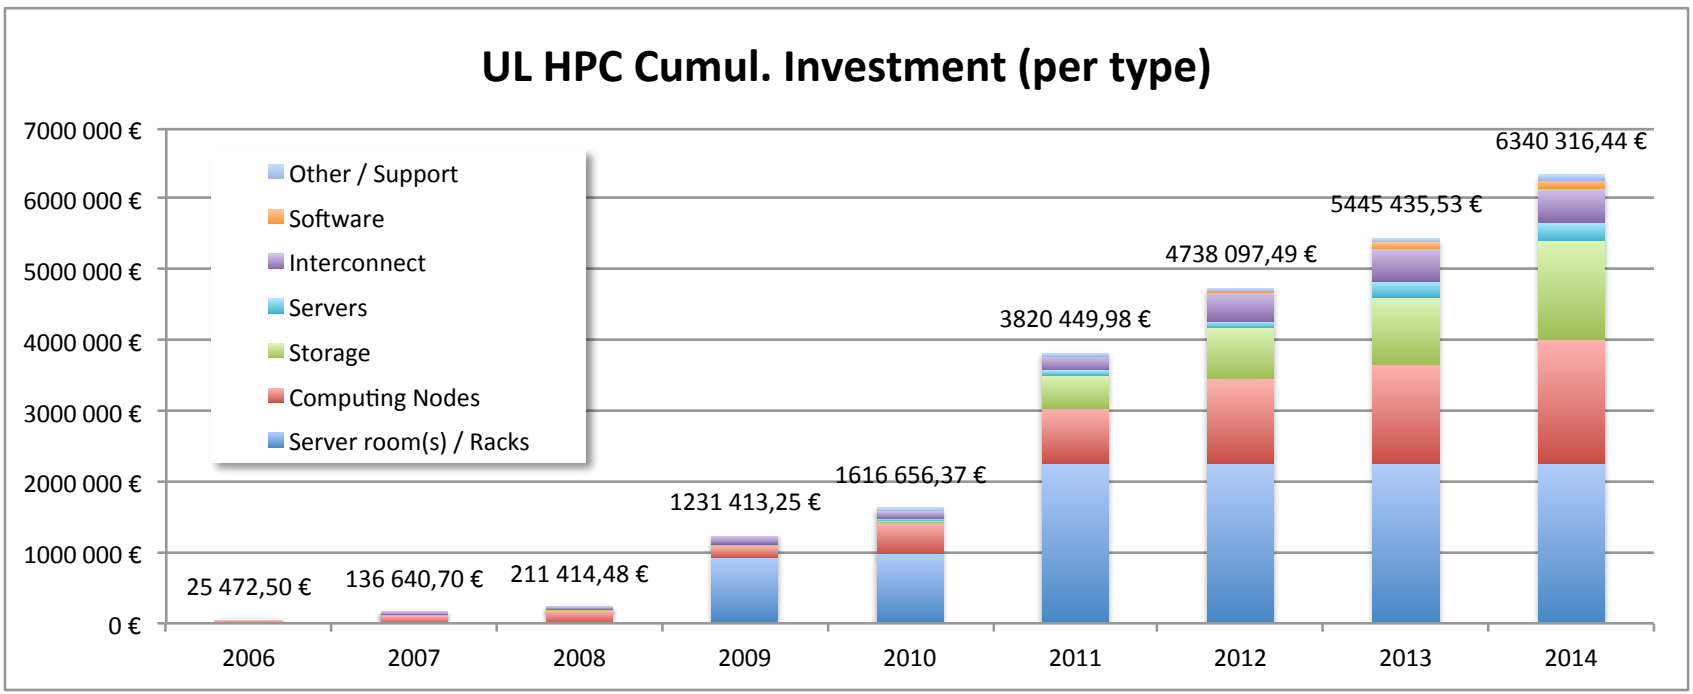

From Platform TCO to Hourly Cost (1)

- **CAPEX**
  - ↪ machines
  - ↪ servers
  - ↪ storage
  - ↪ interconnect
  - ↪ room equipment
  - ↪ lifetime licenses
  - ↪ building estimation
- **OPEX**
  - ↪ manpower
  - ↪ energy
  - ↪ support
  - ↪ yearly licenses

From Platform TCO to Hourly Cost (1)

- **CAPEX**
  - ↪ machines
  - ↪ servers
  - ↪ storage
  - ↪ interconnect
  - ↪ room equipment
  - ↪ lifetime licenses
  - ↪ building estimation
- **OPEX**
  - ↪ manpower
  - ↪ energy
  - ↪ support
  - ↪ yearly licenses

Used Method

- Amortized CAPEX + OPEX → yearly TCO
- Permits to **compute node hourly cost**

From Platform TCO to Hourly Cost (2)

- **Obtained Results for the UL HPC platform**

	Node	CPUs	RAM GB	GPUs	#Nodes	CPU Family	GFLOPS	**Hourly Cost ($)**
Chaos	h-cluster1	12	24	0	32	westmere	108.48	0.428
	d-cluster1	12	24	0	16	westmere	108.48	0.439
	r-cluster1	32	1024	0	1	nehalem	289.28	1.814
	e-cluster1	16	32	0	16	sandybridge	281.60	0.433
	s-cluster1	16	32	0	16	sandybridge	81.60	0.433
Gaia	gaia-[1-60]	12	48	0	60	westmere	108.48	0.453
	gaia-[61-62]	12	24	1792	2	westmere	108.48	0.641
	gaia-[63-72]	12	24	10240	10	westmere	108.48	0.599
	gaia-73	160	1024	0	1	sandybridge	2560.00	2.649
	gaia-74	32	1024	0	1	sandybridge	614.40	1.516
	gaia-[75-79]	16	64	12480	5	sandybridge	281.60	0.577
	gaia-[83-122]	12	48	0	40	westmere	140.64	0.344
	gaia-[123-154]	12	48	0	32	westmere	147.36	0.344

From Platform TCO to Hourly Cost (2)

- **Obtained Results for the UL HPC platform**

	Node	CPUs	RAM GB	GPUs	#Nodes	CPU Family	GFLOPS	**Hourly Cost ($)**
Chaos	h-cluster1	12	24	0	32	westmere	108.48	0.428
	d-cluster1	12	24	0	16	westmere	108.48	0.439
	r-cluster1	32	1024	0	1	nehalem	289.28	1.814
	e-cluster1	16	32	0	16	sandybridge	281.60	0.433
	s-cluster1	16	32	0	16	sandybridge	81.60	0.433
Gaia	gaia-[1-60]	12	48	0	60	westmere	108.48	0.453
	gaia-[61-62]	12	24	1792	2	westmere	108.48	0.641
	gaia-[63-72]	12	24	10240	10	westmere	108.48	0.599
	gaia-73	160	1024	0	1	sandybridge	2560.00	2.649
	gaia-74	32	1024	0	1	sandybridge	614.40	1.516
	gaia-[75-79]	16	64	12480	5	sandybridge	281.60	0.577
	gaia-[83-122]	12	48	0	40	westmere	140.64	0.344
	gaia-[123-154]	12	48	0	32	westmere	147.36	0.344

$\Rightarrow$ How to compare/match with Cloud offers ?

Summary

Methods and Objectives

- **Goal:** Build a new model for *EC2 relative instance price*
  - ↪ ... in function of *all* its HPC characteristics

- **Obj.**: fair comparison between EC2 & in-house HPC facility
  - ↪ for each HPC node, compute its *EC2 relative price*
  - ↪ compare EC2 versus HPC for same characteristics
  - ↪ answer the question: **which one is the cheaper?**

Methods and Objectives

- **Goal:** Build a new model for *EC2 relative instance price*
  - ↪ ... in function of *all* its HPC characteristics

- **Obj.**: fair comparison between EC2 & in-house HPC facility
  - ↪ for each HPC node, compute its *EC2 relative price*
  - ↪ compare EC2 versus HPC for same characteristics
  - ↪ answer the question: **which one is the cheaper?**

- **Method used**: Multiple linear regression
  - ↪ Automated bidirectional stepwise selection
  - ↪ Selection on both criteria and instance types
  - ↪ Selection criterion: adjusted R^2 shrinkage

Proposed Price Model Details

Selected Criteria

- HPC characteristics:
  - ↪ GFLOPs: F
  - ↪ Memory (GB): M
  - ↪ Disk size (GB): D
  - ↪ Nb. GPUs: G

Equation [1]

$$Instance\_Price = \alpha.F + \beta.M + \gamma.D + \delta.G$$

With $\alpha, \beta, \gamma, \delta$ for a given model generation and pricing mode.

Proposed Price Model Details

Selected Criteria

- HPC characteristics:
  - ↪ GFLOPs: F
  - ↪ Memory (GB): M
  - ↪ Disk size (GB): D
  - ↪ Nb. GPUs: G

Equation [1]

$$Instance\_Price = \alpha.F + \beta.M + \gamma.D + \delta.G$$

With $\alpha, \beta, \gamma, \delta$ for a given model generation and pricing mode.

Model	Types	**GFLOPS** (α)	**MemGiB** (β)	**DiskGiB** (γ)	**GPUs** (δ)	Adj. R^2	P-Value
1st Gen.	m1, c1, m2, cg1	0.0039522	0.0061130	0.0000670	0.0015395	0.9999909	0e+00
2nd Gen.	cc2, m3, hi1	-0.0035266	0.0355353	0.0007284	0.0000000	0.9999785	1e-07
3rd Gen.	hs1, cr1, g2, c3	0.0017209	0.0106101	0.0000655	0.0001644	1.0000000	0e+00
4th Gen.	i2, r3, c4	0.0009952	0.0081883	0.0007605	0.0000000	0.9998832	0e+00
5th Gen.	m4, d2	0.0000000	0.0173750	0.0000342	0.0000000	1.0000000	0e+00

Note: The linearity of the model works for instances released at the same period and is broken for new instance releases. Leads to different model generations.

Model Evaluation – On Demand Pricing

- Evaluation of Proposed EC2 Price Models Fittings (with error rates)
  - ↪ ... against actual On Demand instance prices.

1st Instance Generation

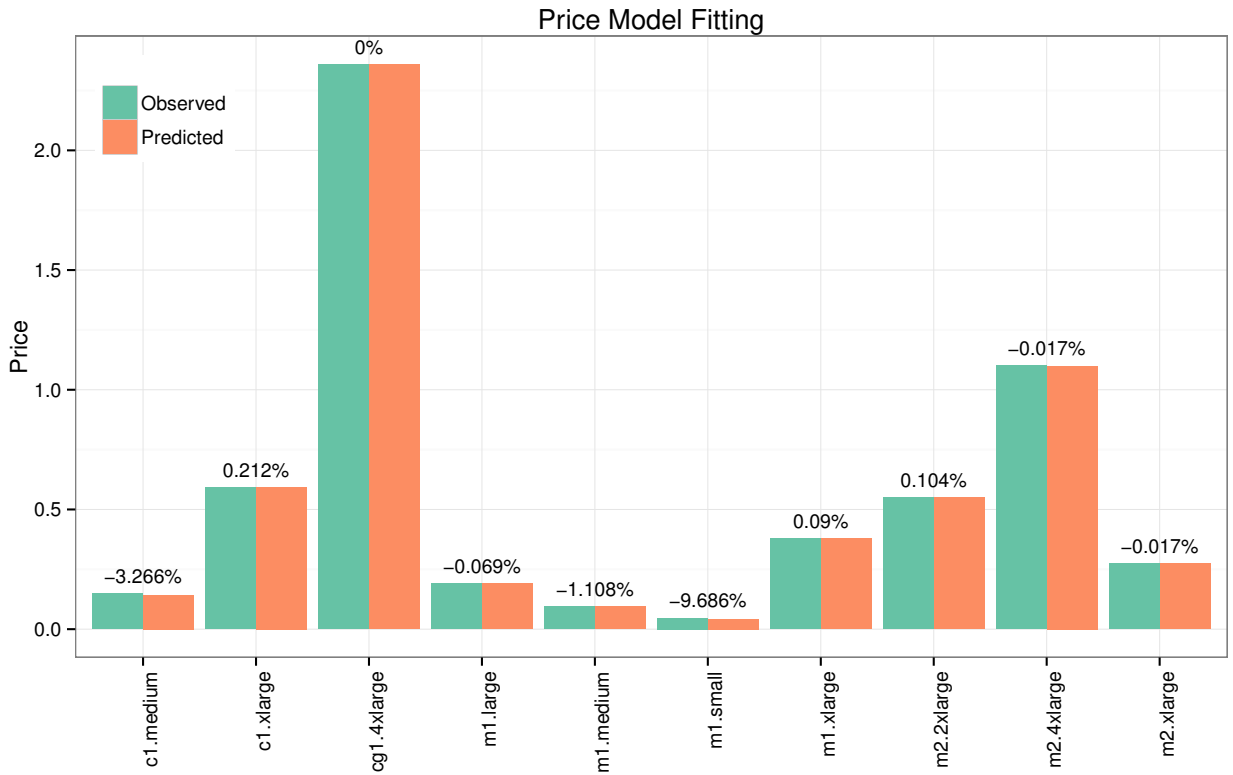

Model Evaluation – On Demand Pricing

- Evaluation of Proposed EC2 Price Models Fittings (with error rates)
  - ↪ ... against actual On Demand instance prices.

2nd Instance Generation

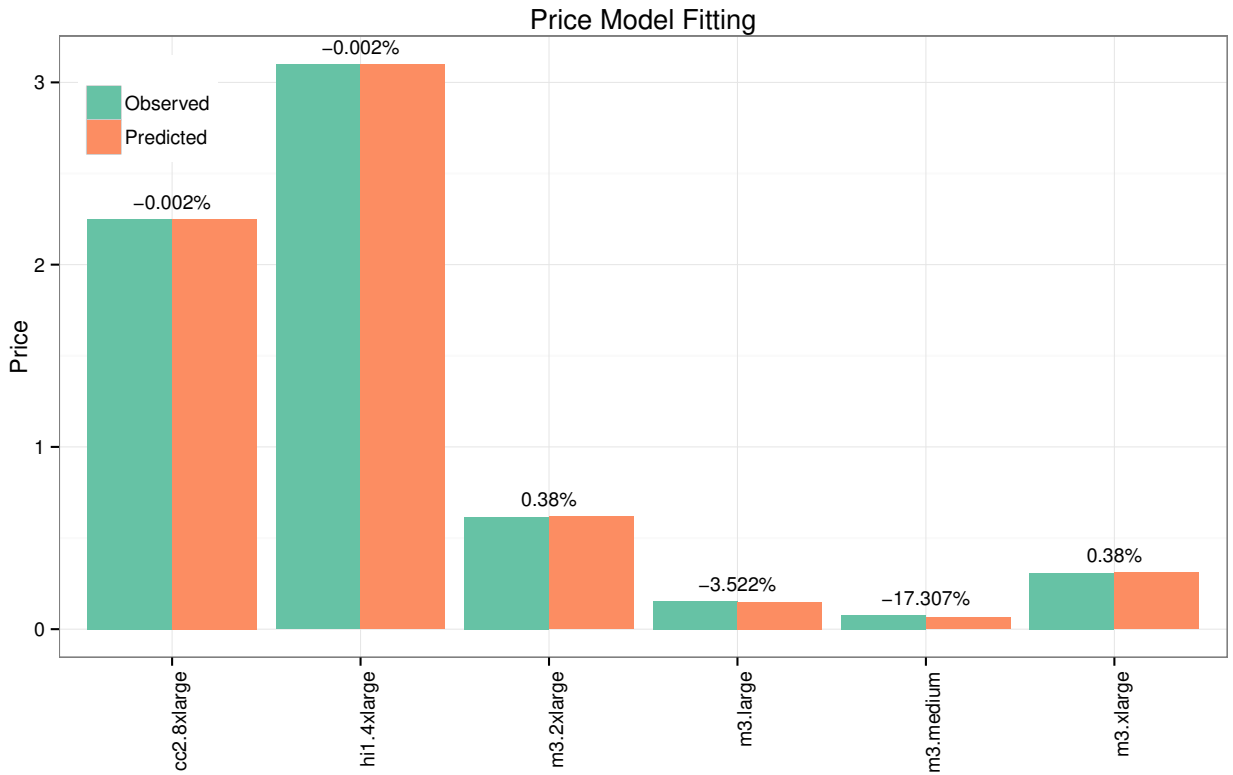

Model Evaluation – On Demand Pricing

- Evaluation of Proposed EC2 Price Models Fittings (with error rates)
  - ↪ ... against actual On Demand instance prices.

3rd Instance Generation

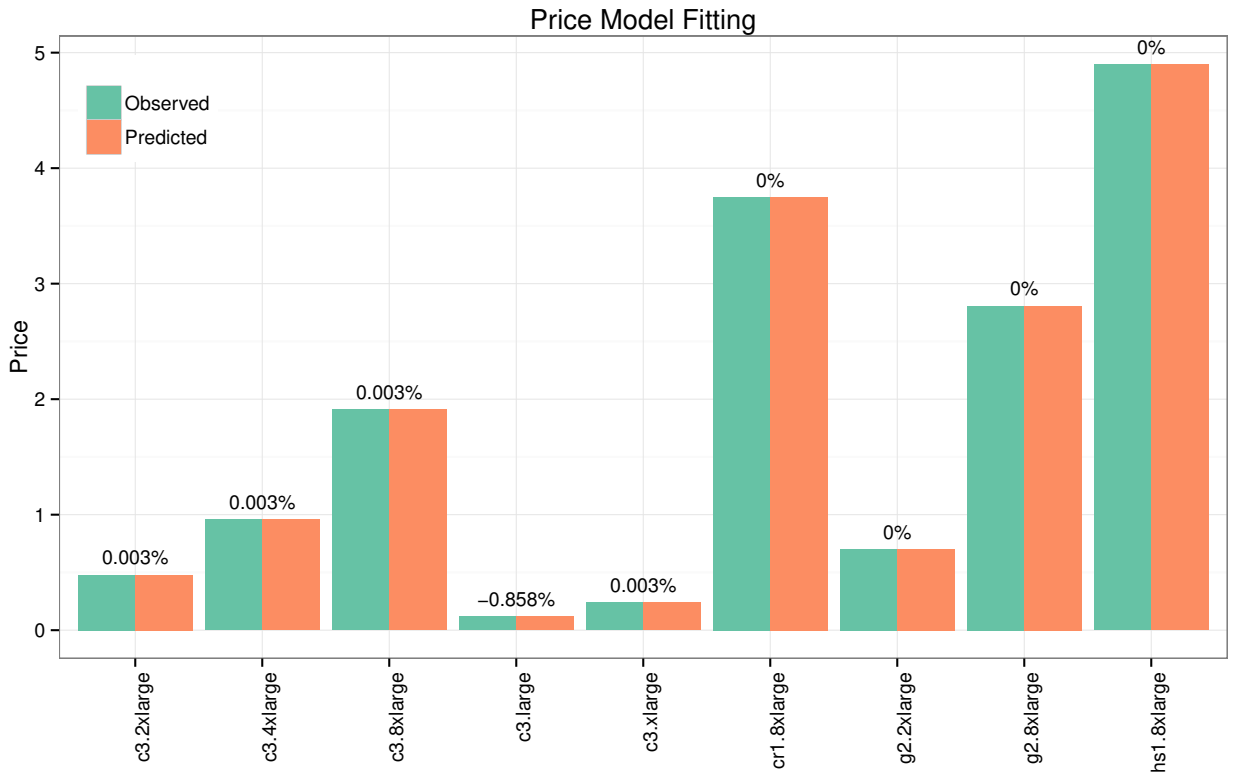

Model Evaluation – On Demand Pricing

- Evaluation of Proposed EC2 Price Models Fittings (with error rates)
  - ↪ ... against actual On Demand instance prices.

4th Instance Generation

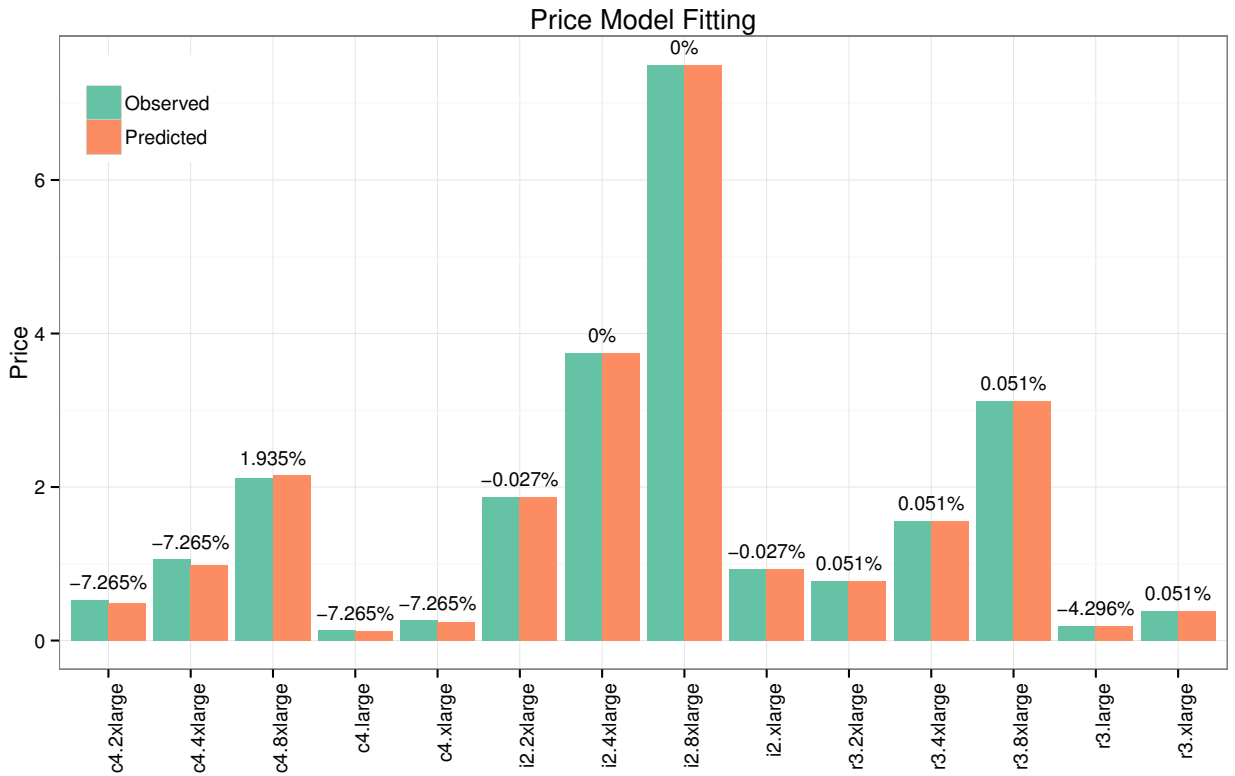

Summary

Hourly Price Comparison

- New cost model permits to estimate *EC2 equivalent price*
  - ↪ ... for **any** computing node configuration

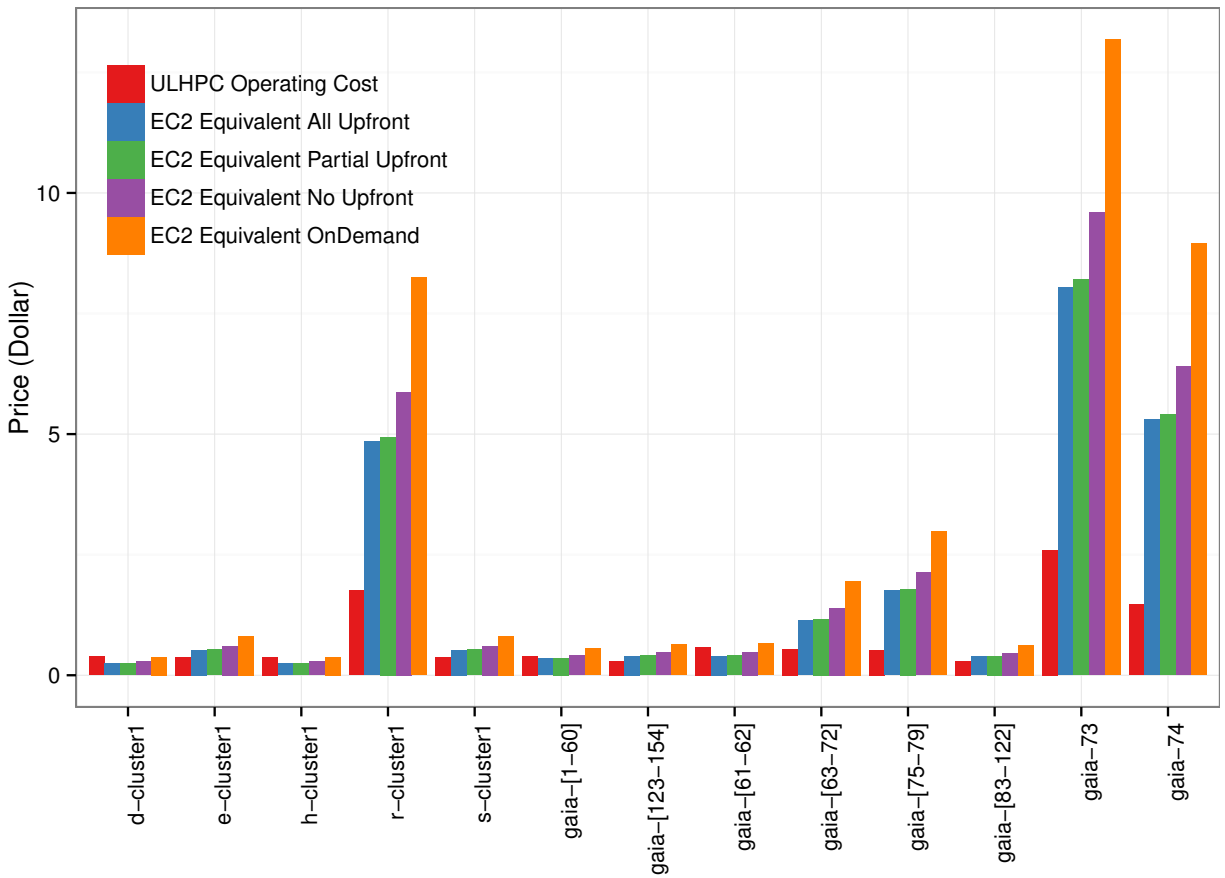

Hourly Price Comparison (by resource classes)

- Assuming definition of 3 HPC *resources classes*

Class	Description
`Normal`	Regular HPC resource
`BigMem`	Regular HPC resource with huge RAM ($\geq$ 1024 GB)
`BigSMP`	SMP node ($\geq$ 16 sockets) with a huge RAM ($\geq$ 1024 GB)

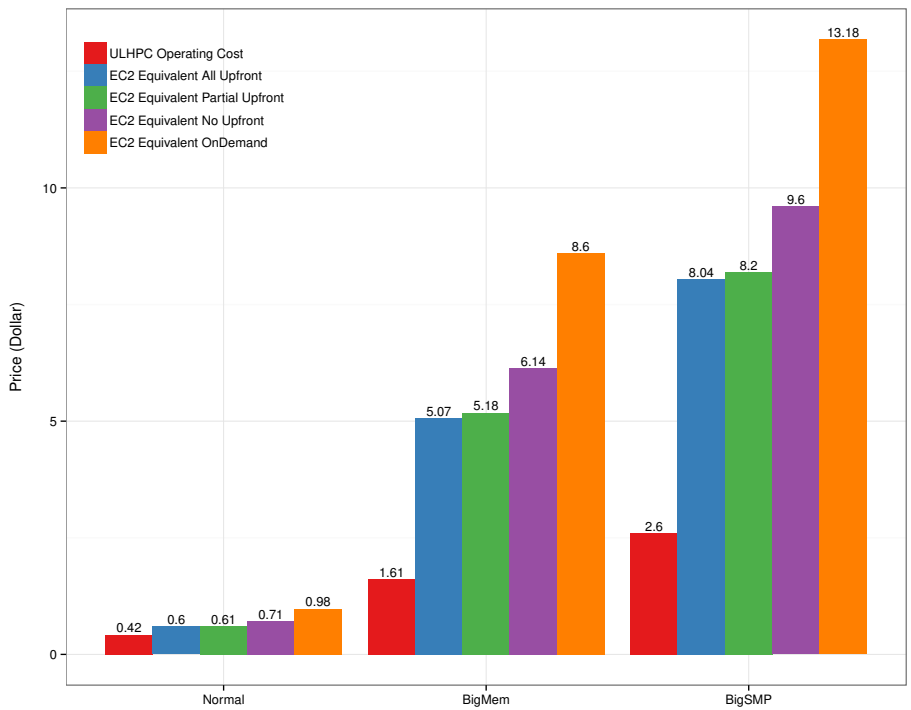

Perf. Comparison – HPC workload

Pure Computing Performance Comparison

- Based on HPCG benchmark http://www.hpcg-benchmark.org/
- compare in-house *gaia*'s cluster obtained performance with EC2
  - ↪ use instances that match the most *gaia* nodes characteristics
  - ↪ obtained score for **1024 cores**

	c3.4xlarge	*r3.4xlarge*	*r3.8xlarge*	*g2.8xlarge*
in-house *gaia*'s efficiency IMPROVEMENT factor	2.5	2.4	2.3	3.2

Toward a Yearly Price Comparison

New Price Model Applied to Real Cluster Usage

- Real HPC job usage extracted from Batch Scheduler Logs
  - ↪ collect all jobs resource allocations
  - ↪ refine cost model with actual performance
  - ↪ apply refined cost model to each job
  - ↪ sum up for year 2014
- includes storage costs

Toward a Yearly Price Comparison

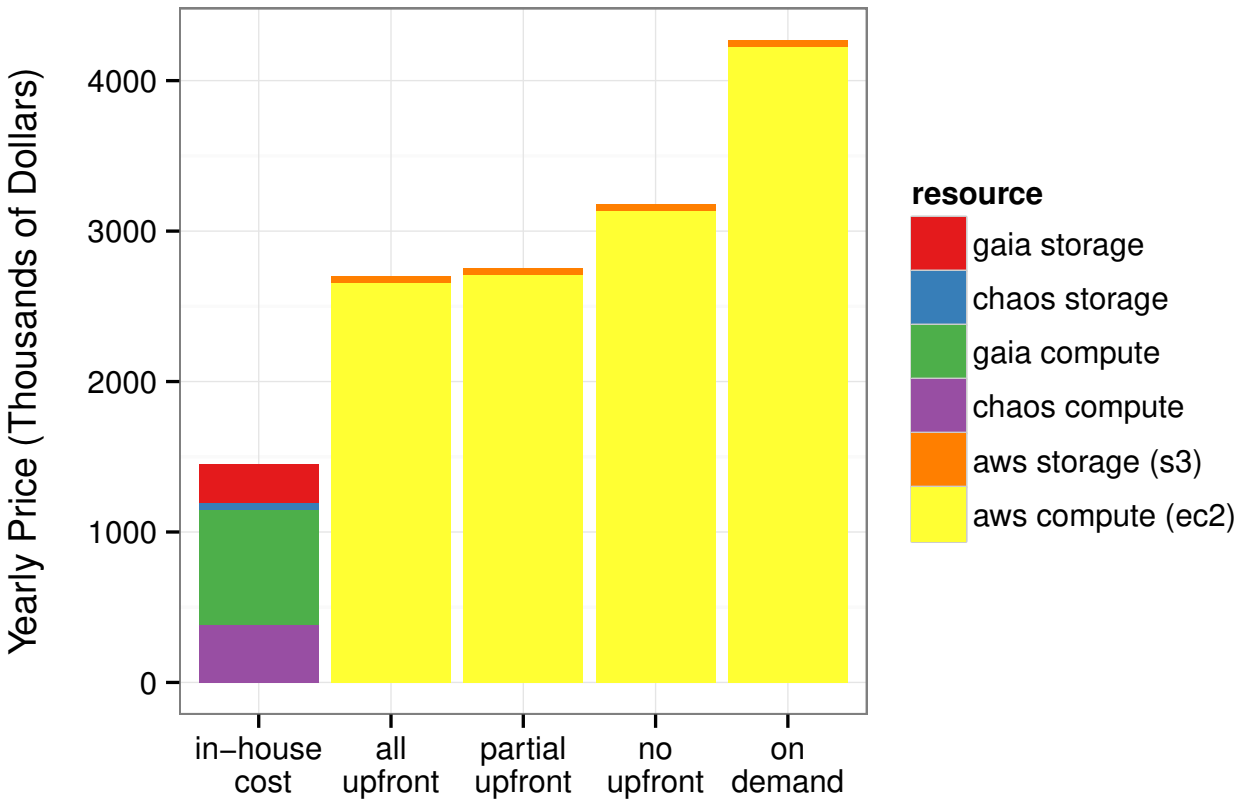

Summary

Conclusion & Perspectives

- **In this talk...**
  - ↪ Cost-effectiveness of CC platforms vs. in-house HPC facility
  - ↪ TCO analysis of a medium-size academic HPC facility
    - ✓ $\simeq$ 350 active users, 5000 cores, 4 sysadmins
  - ↪ novel price model applied to the main Cloud IaaS provider
    - ✓ flexible model that relies on inherent HPC performance metrics
  - ↪ **accurate** cost analysis **based on real HPC usage**

Conclusion & Perspectives

- **In this talk...**
  - ↪ Cost-effectiveness of CC platforms vs. in-house HPC facility
  - ↪ TCO analysis of a medium-size academic HPC facility
    - ✓ $\simeq$ 350 active users, 5000 cores, 4 sysadmins
  - ↪ novel price model applied to the main Cloud IaaS provider
    - ✓ flexible model that relies on inherent HPC performance metrics
  - ↪ **accurate** cost analysis **based on real HPC usage**

$\Rightarrow$ advocates in general in favor of the acquisition of an in-house HPC facility

Conclusion & Perspectives

Lessons Learned

- Deciding where to run your workload is complex
  - ↪ Depends on the performance needs and the workload itself
    - ✓ highly variable or more stable?
  - ↪ Depends on the users' *awareness* of system usage optimization.
- Despite what the Cloud providers advertises
  - ↪ Scale out is complicated...
    - ✓ And not necessarily because of your application scalability.
  - ↪ Do not neglect the cost of experimental setup time

- **Open Perspectives**
  - ↪ extend our analysis over Spot instance (Amazon) and Azure offer
  - ↪ integrate communication cost (to/from cloud storage)

Submit and Join IEEE CloudCom'16 !!!

http://2016.cloudcom.org

Full Papers:	**July 15th, 2016**
Short Papers/Demos:	August 10th, 2016
Notification of acceptance:	August 31th, 2016
Conference Date:	Dec. 12th →15h, 2016

Contacts:

Sebastien.Varrette@uni.lu
Pascal.Bouvry@uni.lu

uni.lu
UNIVERSITÉ DU LUXEMBOURG

Thank you for your attention...

Questions?

J. Emeras, S. Varrette, and P. Bouvry.
Amazon Elastic Compute Cloud (EC2) vs. in-House HPC Platform: a Cost Analysis.
In *Proc. of the 9th IEEE Intl. Conf. on Cloud Computing (CLOUD 2016)*, San Francisco, USA, June 2016. IEEE Computer Society.

Contacts:
mail: firstname.lastname@uni.lu