

Team Meeting

An EC2 Cloud Pricing Model



Joseph Emeras


PCOG Research Group,
University of Luxembourg




EC2 Pricing

Depending on the region, instance type, options and pricing mode


- **On-Demand:** Fixed hourly rate. Most expensive / flexible
- **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
- **Spot:** Price bidding / flexible. Fixed duration option (1 / 6 hours)





Summary

- 1 **EC2 Versus HPC Comparison**
 - HPC Platform
 - EC2 Price Model
 - Model Evaluation
 - Price Model Application for Cost Comparison



From Platform TCO to Hourly Cost

CAPEX

- machines
- servers
- storage
- interconnect
- room equipment
- lifetime licenses
- building estimation

OPEX

- manpower
- energy
- support and yearly licenses

Method

- amortized CAPEX + OPEX → yearly TCO
- compute node hourly cost

	Node	CPUs	Memory GB	GPUs	Nb. Nodes	CPU Family	Disk GB	GFLOPS	Hourly Cost (\$)
Chaos	h-cluster1	12	24	0	32	westmere	250	108.48	0.428
	d-cluster1	12	24	0	16	westmere	250	108.48	0.439
	r-cluster1	32	1024	0	1	nehalem	250	289.28	1.814
	e-cluster1	16	32	0	16	sandybridge	250	281.60	0.433
	s-cluster1	16	32	0	16	sandybridge	250	281.60	0.433
Gala	gaia-[1-60]	12	48	0	60	westmere	256	108.48	0.453
	gaia-[61-62]	12	24	1792	2	westmere	256	108.48	0.641
	gaia-[63-72]	12	24	10240	10	westmere	256	108.48	0.599
	gaia-73	160	1024	0	1	sandybridge	256	2560.00	2.649
	gaia-74	32	1024	0	1	sandybridge	256	614.40	1.516
	gaia-[75-79]	16	64	12480	5	sandybridge	256	281.60	0.577
	gaia-[83-122]	12	48	0	40	westmere	256	140.64	0.344
	gaia-[123-154]	12	48	0	32	westmere	256	147.36	0.344



Building an EC2 Price Model

Goal


Construct a model that gives an instance price in function of its characteristics.

Why?

- For each HPC node, compute its *EC2 relative price*
- Compare EC2 versus HPC for same characteristics
- answer the question: which is the cheaper?

How?

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



Price Model Details

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.

Selected Criteria


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

Equation [1]





$$\text{Instance_Price} = \alpha.F + \beta.M + \gamma.D + \delta.G$$

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

Model	Types	Adj. R^2	P-Value
1st Gen.	m1, c1, m2, cg1	0.9999909	0e+00
2nd Gen.	cc2, m3, hi1	0.9999785	1e-07
3rd Gen.	hs1, cr1, g2, c3	1.0000000	0e+00
4th Gen.	i2, r3, c4	0.9998832	0e+00
5th Gen.	m4, d2	1.0000000	0.0e+00





Model Evaluation – On Demand Pricing





Price Comparison





Experimental Evaluation

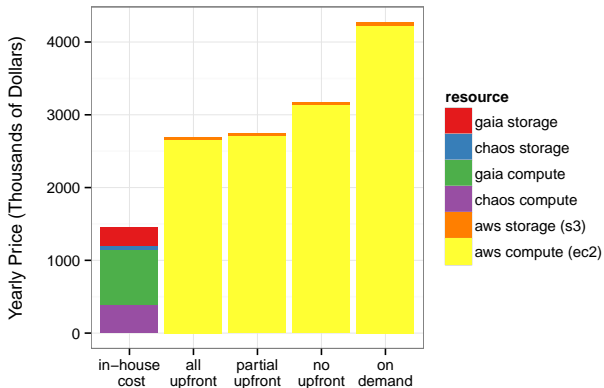
HPCG Benchmark

- compare *gaia*'s 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
gaia's efficiency improvement factor	2.5	2.4	2.3	3.2



Price Comparison vs. Performance



Thank you for your attention



Game

One free beer for the first that gives me the name of this cloud type...





Emeras, J., Varrette, S., and Bouvry, P.

Amazon elastic compute cloud (ec2) vs. in-house hpc platform: a cost analysis.

Submitted at CCGRID'16 (2016).