

Evaluation of Cloud Providers for VPH Applications

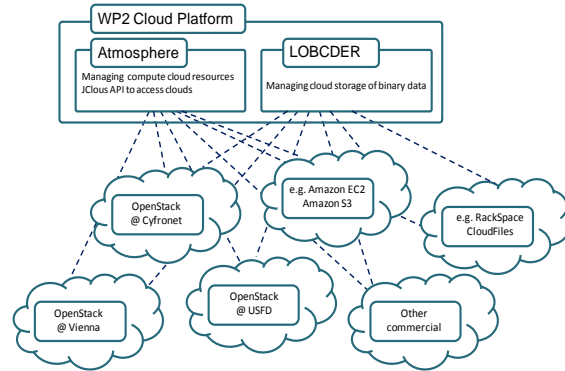


Marian Bubak¹, Marek Kasztelnik¹, Maciej Malawski¹, Jan Meizner¹, Piotr Nowakowski¹, Susheel Varma²
 (1) AGH University of Science and Technology, al. Mickiewicza 30, 30-059 Kraków, Poland
 (2) Dept of Cardiovascular Science, University of Sheffield, Sheffield S10 2RX, United Kingdom



Motivation

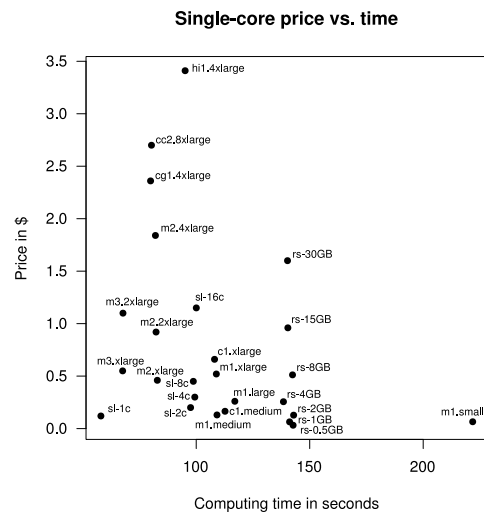
- VPH-Share biomedical applications
- Atmosphere hybrid cloud platform
 - Development and publishing of applications as Atomic Services
 - Dynamic resource allocation
 - Data access and security
- Need to procure resources from public cloud providers for the project



Evaluation criteria


- Compute service, Storage service (BLOB)
- jClouds library support for integration with Atmosphere
- European Economic Area (EEA) zoning due to policy
- Elasticity, on-demand access:
 - API access
 - Hourly or shorter billing
 - Published price
- VM Import/Export functionality
- Relational Database support

instance type	hourly price in \$	number of cores	RAM in GB	provider
m1.small	0.065	1	1.7	EC2
m1.medium	0.13	1	3.75	EC2
m1.large	0.26	2	7.5	EC2
m1.xlarge	0.52	4	15	EC2
m2.xlarge	0.46	2	17.1	EC2
m2.2xlarge	0.92	4	34.2	EC2
m2.4xlarge	1.84	8	68.4	EC2
c1.medium	0.165	2	1.7	EC2
c1.xlarge	0.66	8	7	EC2
hi1.4xlarge	3.41	8	60.5	EC2
cc2.8xlarge	2.7	16	60.5	EC2
cg1.4xlarge	2.36	8	22	EC2
m3.xlarge	0.55	4	15	EC2
m3.2xlarge	1.1	8	30	EC2
rs-0.5GB	0.032	1	0.5	RackSpace
rs-1GB	0.064	1	1	RackSpace
rs-2GB	0.128	2	2	RackSpace
rs-4GB	0.256	2	4	RackSpace
rs-8GB	0.512	4	8	RackSpace
rs-15GB	0.96	4	15	RackSpace
rs-30GB	1.6	8	30	RackSpace
sl-1c	0.12	1	1	SoftLayer
sl-2c	0.2	2	2	SoftLayer
sl-4c	0.3	4	4	SoftLayer
sl-8c	0.45	8	8	SoftLayer
sl-16c	1.15	16	16	SoftLayer




Provider	EEA Zoning	jClouds API Support	BLOB storage support	Per-hour instance billing	API Access	Published price	VM Image Import/Export	Relational DB support	Score
1 Amazon AWS	1	1	1	1	1	1	0	1	27
2 Rackspace	1	1	1	1	1	1	0	1	27
3 SoftLayer	1	1	1	1	1	1	0	0	25
4 CloudSigma	1	1	0	1	1	1	1	0	18
5 ElasticHosts	1	1	0	1	1	1	1	0	18
6 Serverlove	1	1	0	1	1	1	1	0	18
7 GoGrid	1	1	0	1	1	1	0	0	15
8 Terremark ecloud	1	1	0	1	1	0	1	0	13
9 RimuHosting	1	1	0	0	1	1	0	1	12
10 Stratogen	1	1	0	0	1	0	1	0	8
11 Bluelock	1	1	0	0	1	0	0	0	5
12 Fujitsu GCP	1	1	0	0	1	0	0	0	5
13 BitRefinery	0	0	0	0	0	1	0	1	0
14 BrightBox	1	0	0	1	1	1	1	0	0
15 BT Global Services	1	0	0	0	1	0	1	0	0
16 Carpathia Hosting	1	0	0	0	0	0	1	0	0
17 City Cloud	1	0	0	1	1	1	0	0	0
18 Claris Networks	0	0	0	1	0	0	0	0	0
19 Coderio	0	0	0	1	1	1	0	0	0
20 CSC	1	0	0	0	0	0	1	0	0
21 Datapipe	1	0	0	1	1	0	0	0	0
22 e24cloud	1	0	0	1	0	1	0	0	0
23 eApps	0	0	0	0	0	1	0	0	0
24 FlexiScale	1	0	0	1	1	1	1	0	0
25 Google GCE	1	0	1	1	1	1	0	1	0
26 Green House Data	0	0	0	0	1	0	1	0	0
27 Hosting.com	0	0	0	0	0	1	1	1	0
28 HP Cloud	0	1	1	1	1	1	1	1	0
29 IBM SmartCloud	0	0	1	1	1	1	0	1	0
30 IJ GIO	0	0	0	0	0	0	0	0	0
31 iland cloud	1	0	0	1	0	0	1	0	0
32 Internap	0	0	1	1	1	1	0	0	0
33 Joyent	0	0	0	1	1	1	0	0	0
34 LunaCloud	1	0	1	1	1	1	0	0	0
35 Oktawave	1	0	1	1	1	1	0	1	0
36 Openhosting.co.uk	1	0	0	0	0	1	0	0	0
37 Openhosting.com	0	1	0	1	1	1	1	0	0
38 OpSource	1	0	1	1	1	1	1	0	0
39 ProfitBricks	1	0	0	1	1	1	0	0	0
40 Qube	1	0	0	0	0	1	0	0	0
41 ReliaCloud	0	0	0	0	0	0	0	0	0
42 SaavisDirect	0	0	1	1	0	1	0	0	0
43 SkaliCloud	0	1	0	1	1	1	1	0	0
44 Teklinks	0	0	0	0	0	0	0	0	0
45 Terremark vcloud	0	1	0	1	1	1	1	0	0
46 Tier 3	0	0	0	0	0	1	0	0	0
47 Umbec	1	0	0	1	1	1	1	0	0
48 VPS.net	1	0	0	0	1	1	1	0	0
49 Windows Azure	1	0	1	1	1	1	0	1	0


Instance prices



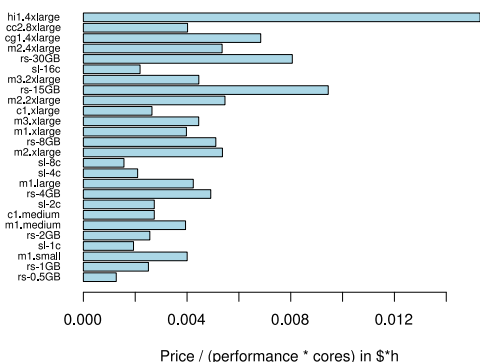
Computing time



Single-core price / performance



Multi-core price / performance



Observations

- For sequential applications, instance type is not highly relevant
- Cheapest instances are useful for high-throughput computing
- Single-core instances give better granularity for auto-scaling
- SL and second-generation EC2 instances m3.large and m3.xlarge give the best sequential performance
- m3.x are not as efficient for multi-core jobs – the speedup using their virtual cores is not linear
- RAM and Disk requirements introduce additional constraints

References

- P. Nowakowski, T. Bartynski, T. Gubala, D. Harezlak, M. Kasztelnik, M. Malawski, J. Meizner, M. Bubak: Cloud Platform for Medical Applications, *eScience* 2012
- M. Malawski, T. Gubala, M. Bubak, Component-based approach for programming and running scientific applications on grids and clouds, *International Journal of High Performance Computing Applications*, vol. 26, no. 3, pp. 275–295, August 2012
- M. Malawski, M. Kuzniar, P. Wójcik, M. Bubak: How to Use Google App Engine for Free Computing. *IEEE Internet Computing* 17(1): 50-59 (2013)
- M. Malawski, K. Figiela, J. Nabrzyski: Cost Minimization for Computational Applications on Hybrid Cloud Infrastructures, *FGCS* 2013
- M. Malawski, G. Juve, E. Deelman, J. Nabrzyski: Cost- and deadline-constrained provisioning for scientific workflow ensembles in IaaS clouds. *SC* 2012