Using Amazon EC2 Cloud Spot Instances as cheap compute resource for compute intensive scientific data analysis

Gergely Debreczeni

2014.05.26 - VDASC telecon

About Amazon EC2

Amazon provides a wide range of various computing services.

http://aws.amazon.com

Almost any kind of configuration regarding, CPU, memory, storage and networking is possible.

Easy to register, video tutorial is helping you in the first steps.

Free trial account with 750 hour of free CPU usage / month during 1 year ! Just perfect for evaulation studies.

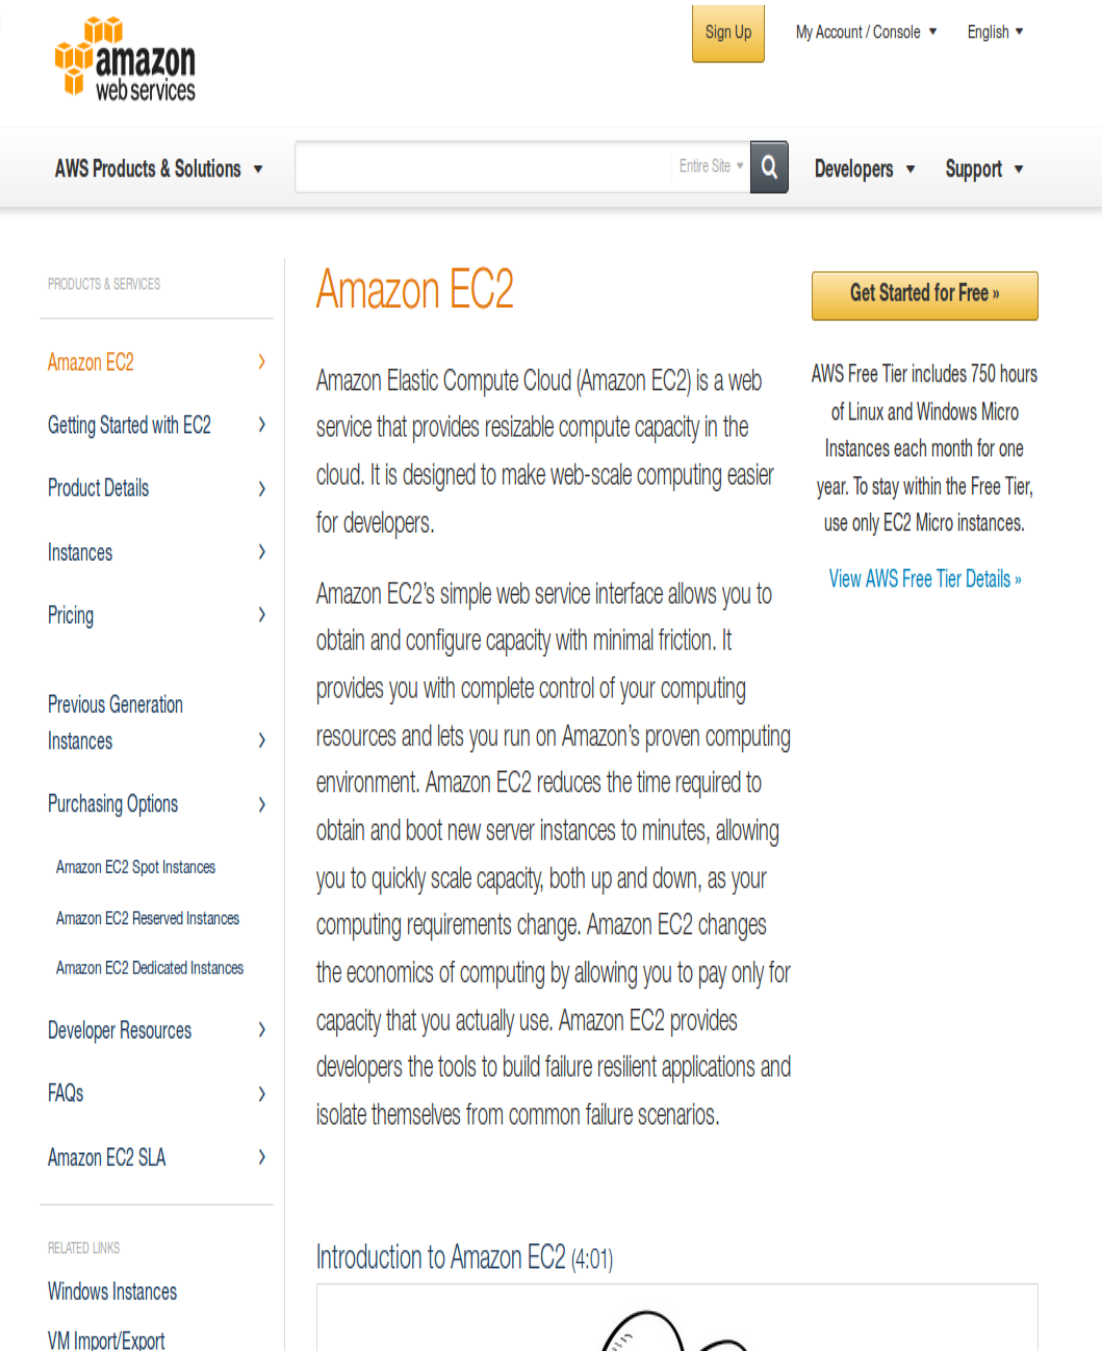

Instances - a few example

Instance Types Matrix

Instance Type	vCPU	Memory (GiB)	Storage (GB)	Networking Performance	Physical Processor	Clock Speed (GHz)	Intel® AES-NI	Intel® AVX†	Intel® Turbo	EBS OPT	Enhanced Networking
m3.medium	1	3.75	1 x 4 SSD	Moderate	Intel Xeon E5-2670*	2.6	Yes	Yes	Yes	-	-
m3.large	2	7.5	1 x 32 SSD	Moderate	Intel Xeon E5-2670	2.6	Yes	Yes	Yes	-	-
m3.xlarge	4	15	2 x 40 SSD	Moderate	Intel Xeon E5-2670	2.6	Yes	Yes	Yes	Yes	-
m3.2xlarge	8	30	2 x 80 SSD	High	Intel Xeon E5-2670	2.6	Yes	Yes	Yes	Yes	-
c3.large	2	3.75	2 x 16 SSD	Moderate	Intel Xeon E5-2670 v2	2.8	Yes	Yes	Yes	-	Yes
c3.xlarge	4	7.5	2 x 40 SSD	Moderate	Intel Xeon E5-2670 v2	2.8	Yes	Yes	Yes	Yes	Yes
c3.2xlarge	8	15	2 x 80 SSD	High	Intel Xeon E5-2670 v2	2.8	Yes	Yes	Yes	Yes	Yes
c3.4xlarge	16	30	2 x 160 SSD	High	Intel Xeon E5-2670 v2	2.8	Yes	Yes	Yes	Yes	Yes
c3.8xlarge	32	60	2 x 320 SSD	10 Gigabit	Intel Xeon E5-2670 v2	2.8	Yes	Yes	Yes	-	Yes

About spot instances

Spot instances are Amazon's „screen saver” solutions to sell their not used spare resources on a strongly reduced prices.

This was Amazon cloud is almost full all the time and applications that do not need on-demand immediate compute resources but can wait until computie prices are become sufficiently low can make very good use of it.

Prices are changing all time, just like in the stock market, upper (and lower) limits can be set by the user when to allow his/her image to be executed.

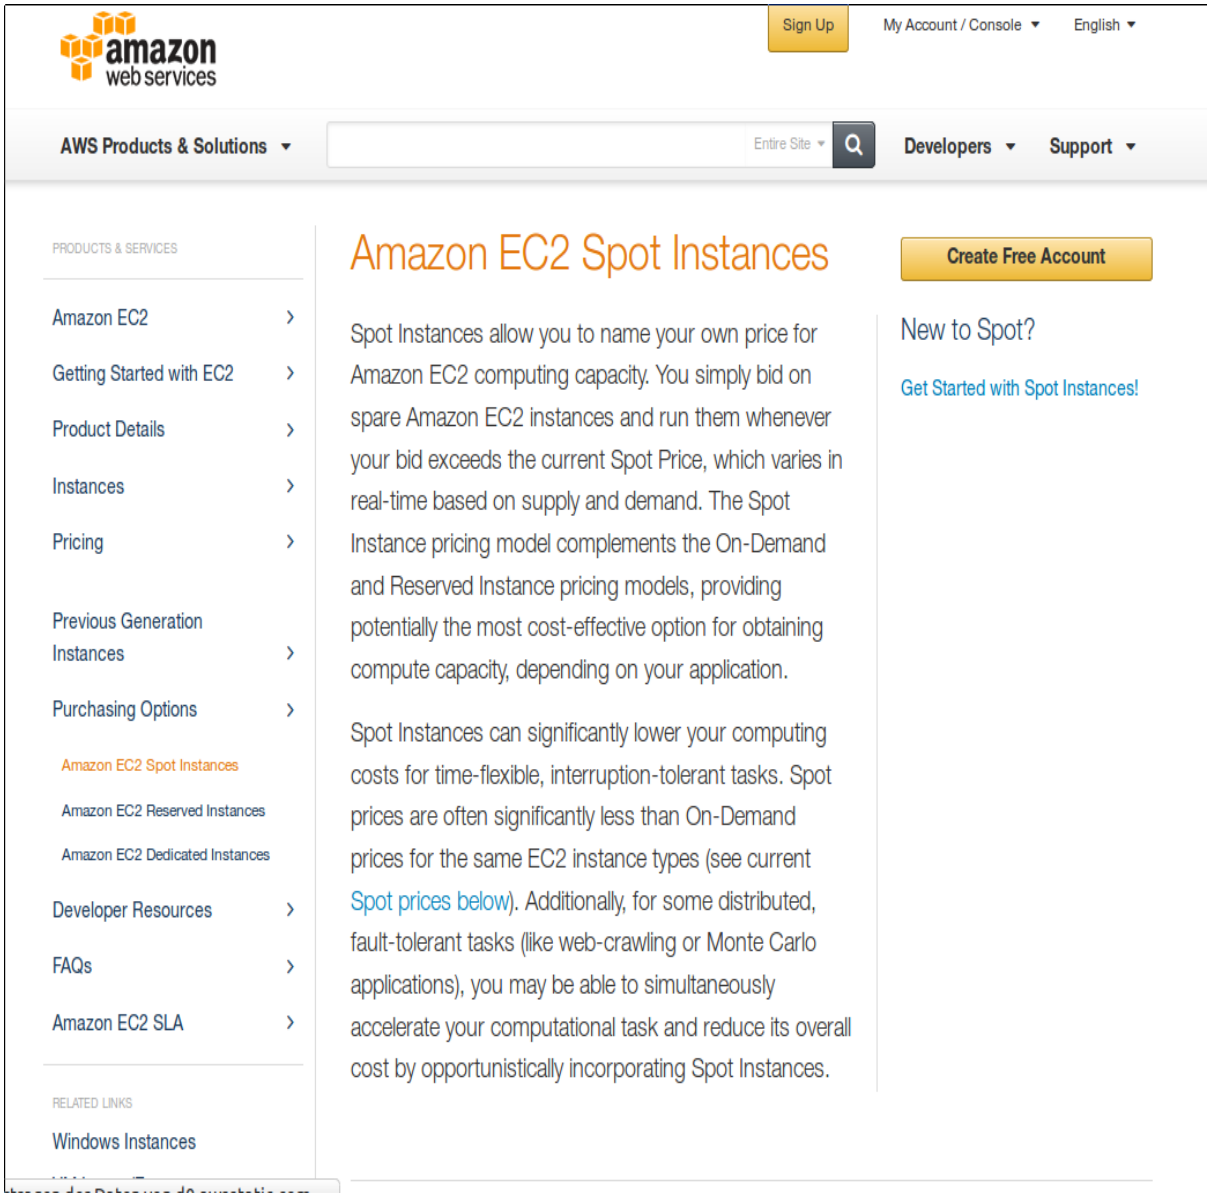

Pricing of spot instances

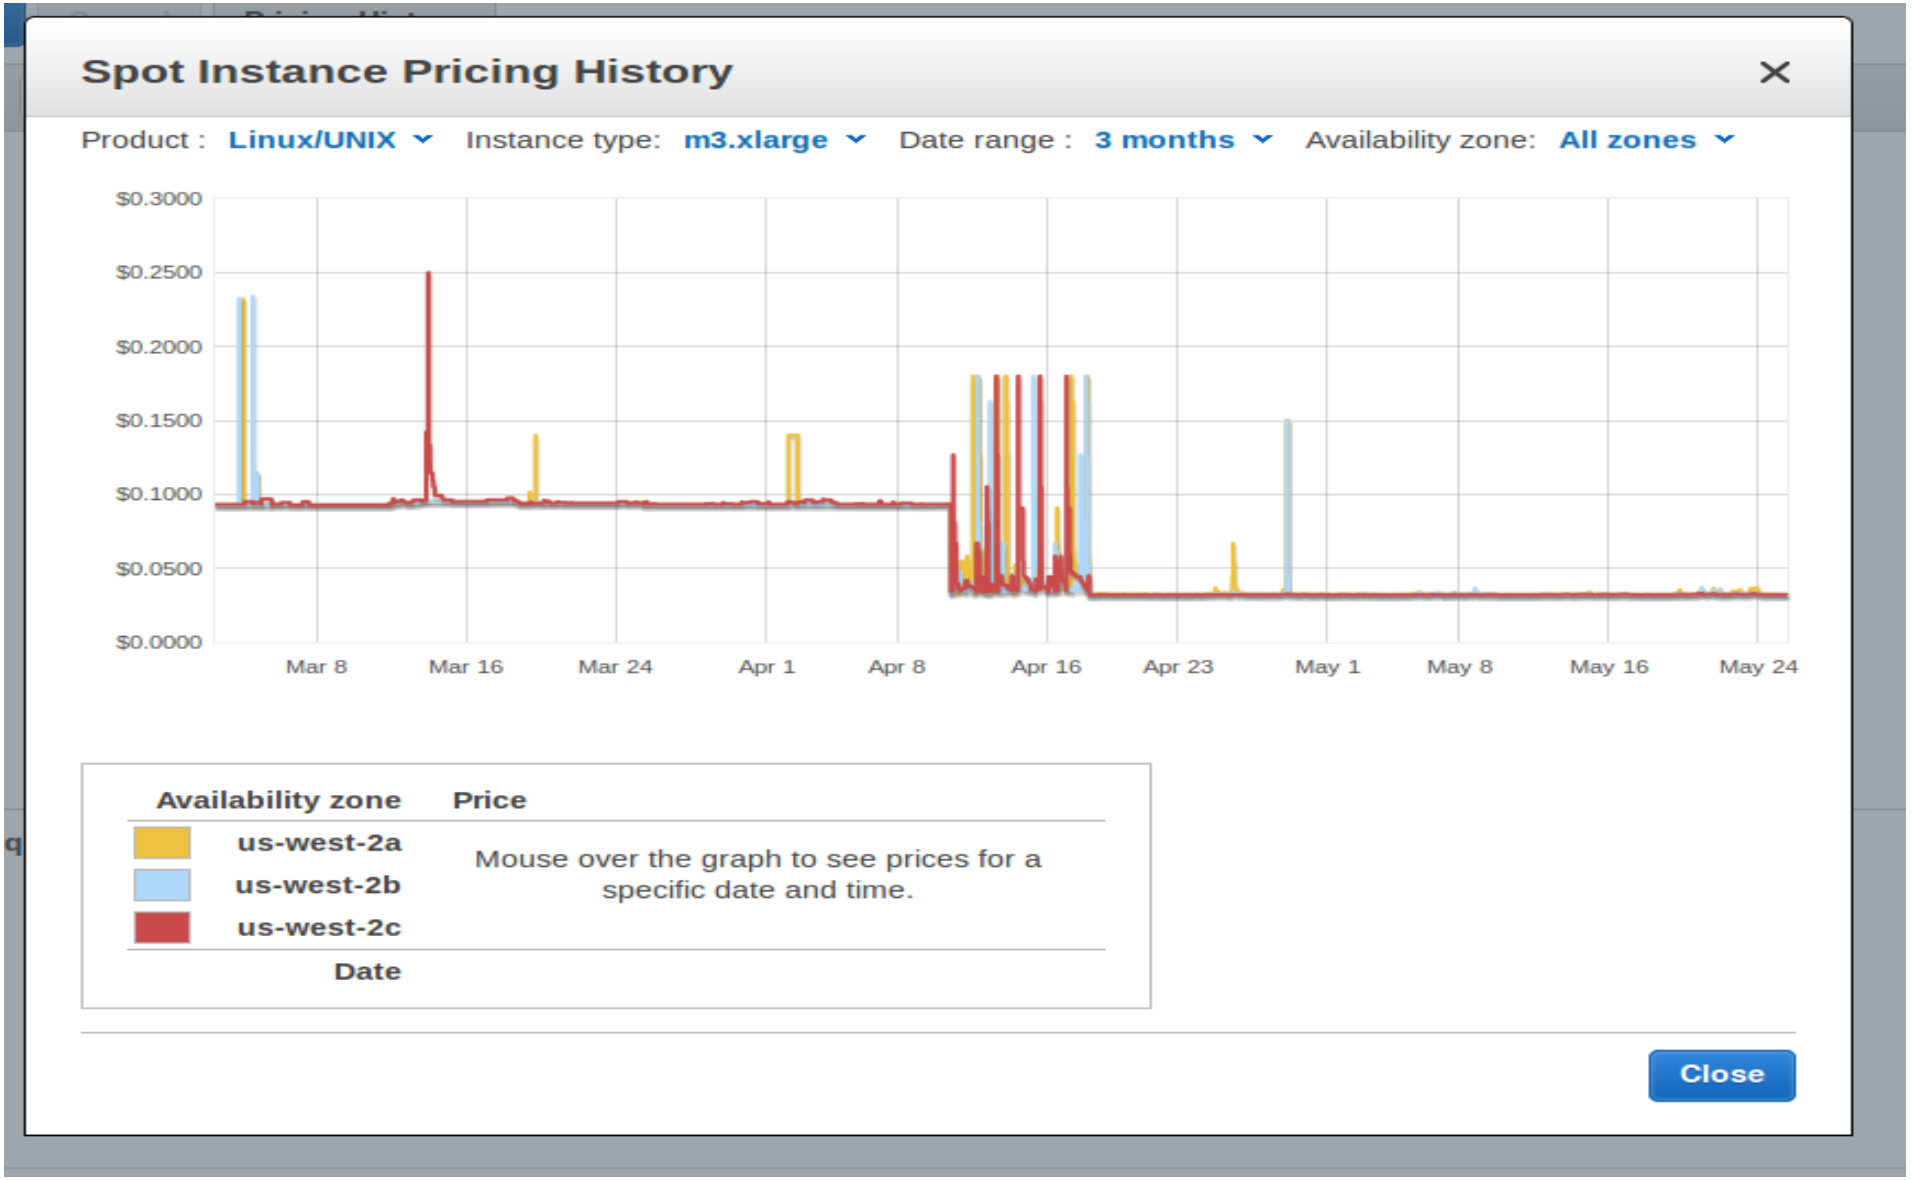

Pricing of spot instances

In various configurations the price / core price more or less the same. An ideal configuration is m3.xlarge with 4 CPU and 15 GB memory:

0.008 $ / hour = 0.192 $ / day = 0.0192 $ / HS06.day = **0.014 Euro / HS06.day**

this is approximately the half what we pay for our dedicated resources !

Of course we have to benchmark it before claiming anything definit !

Using Amazon EC2 Spot instances can save us 40-50 KEuro / year !

Management console

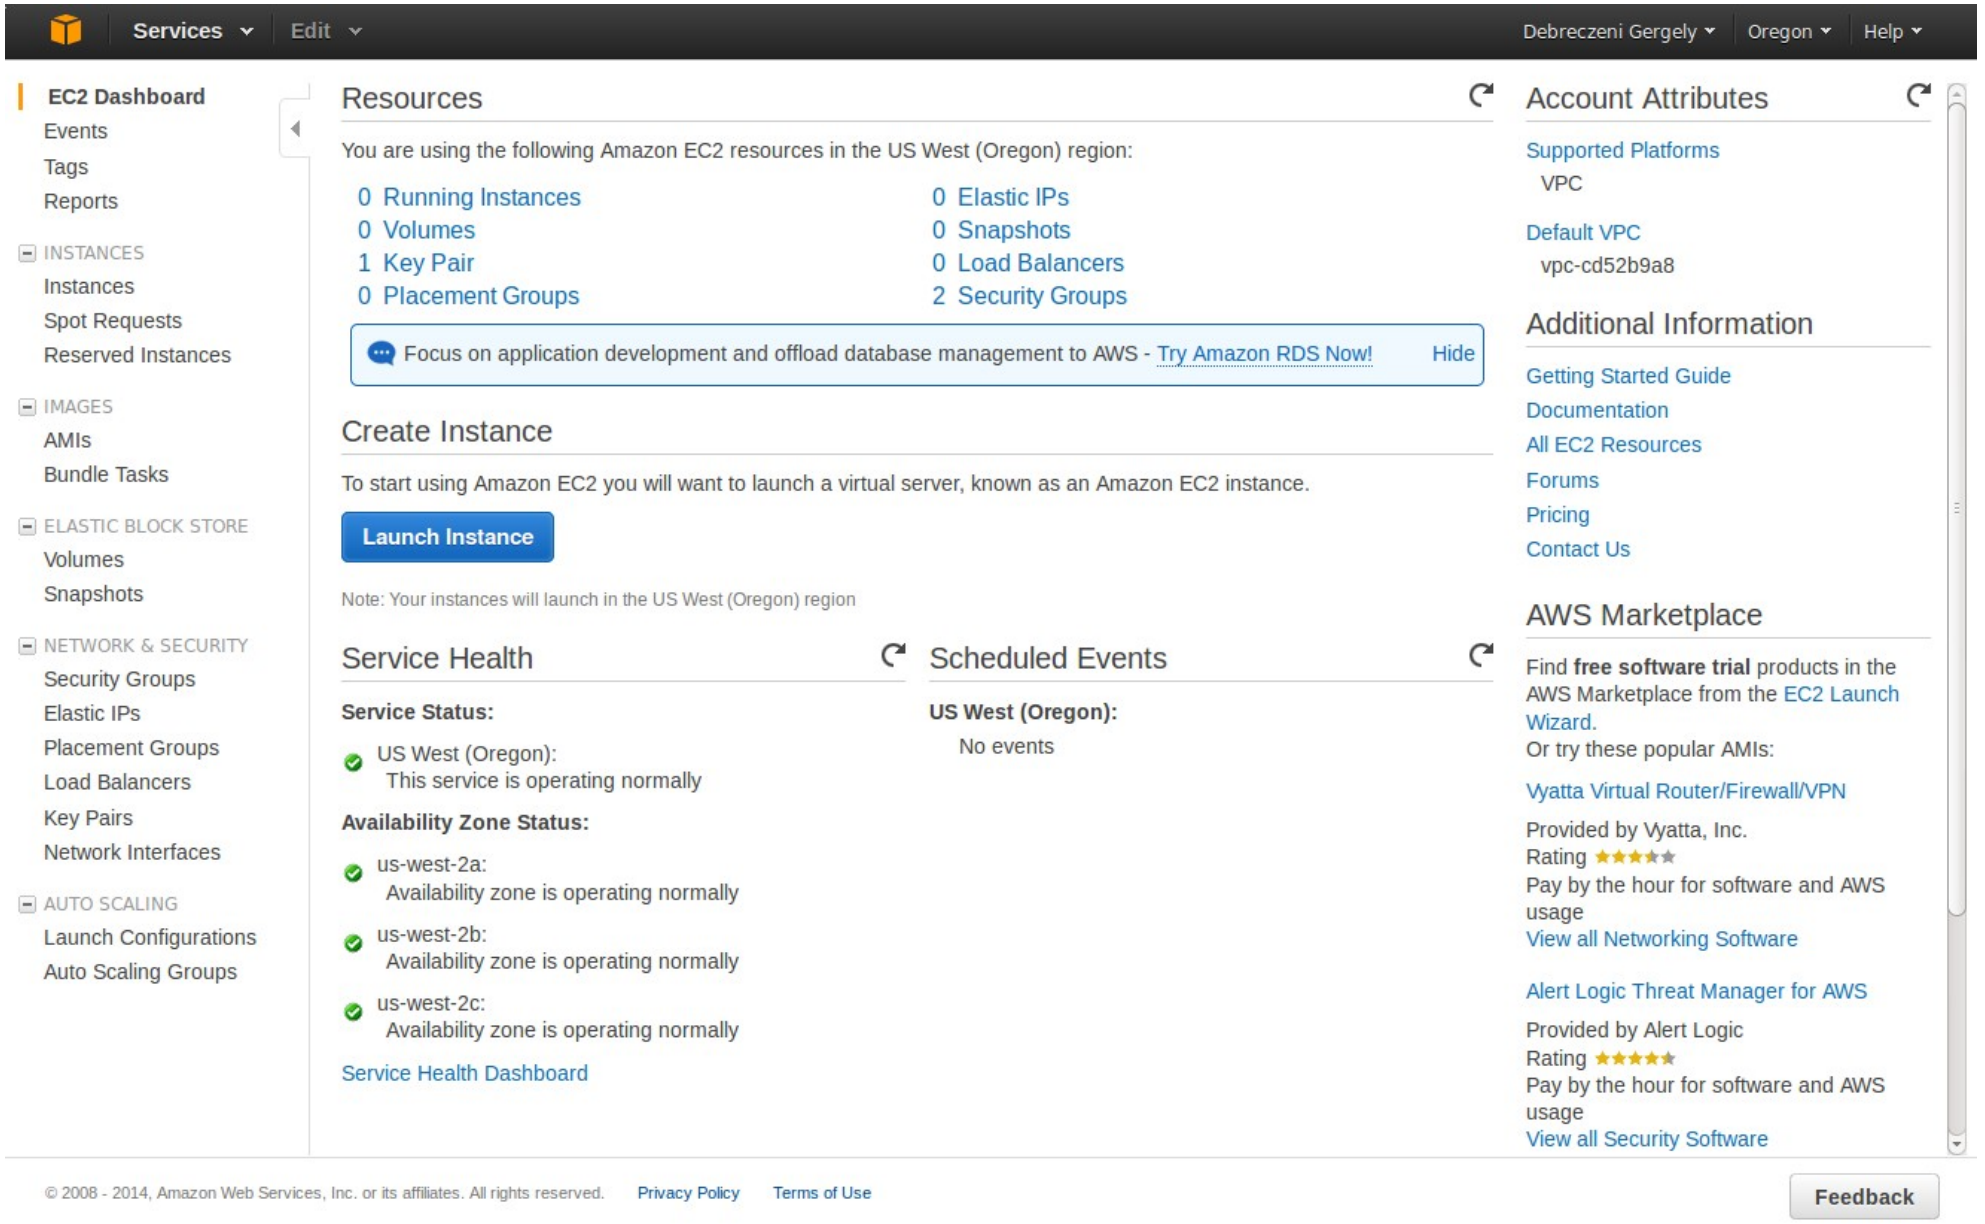

Launching

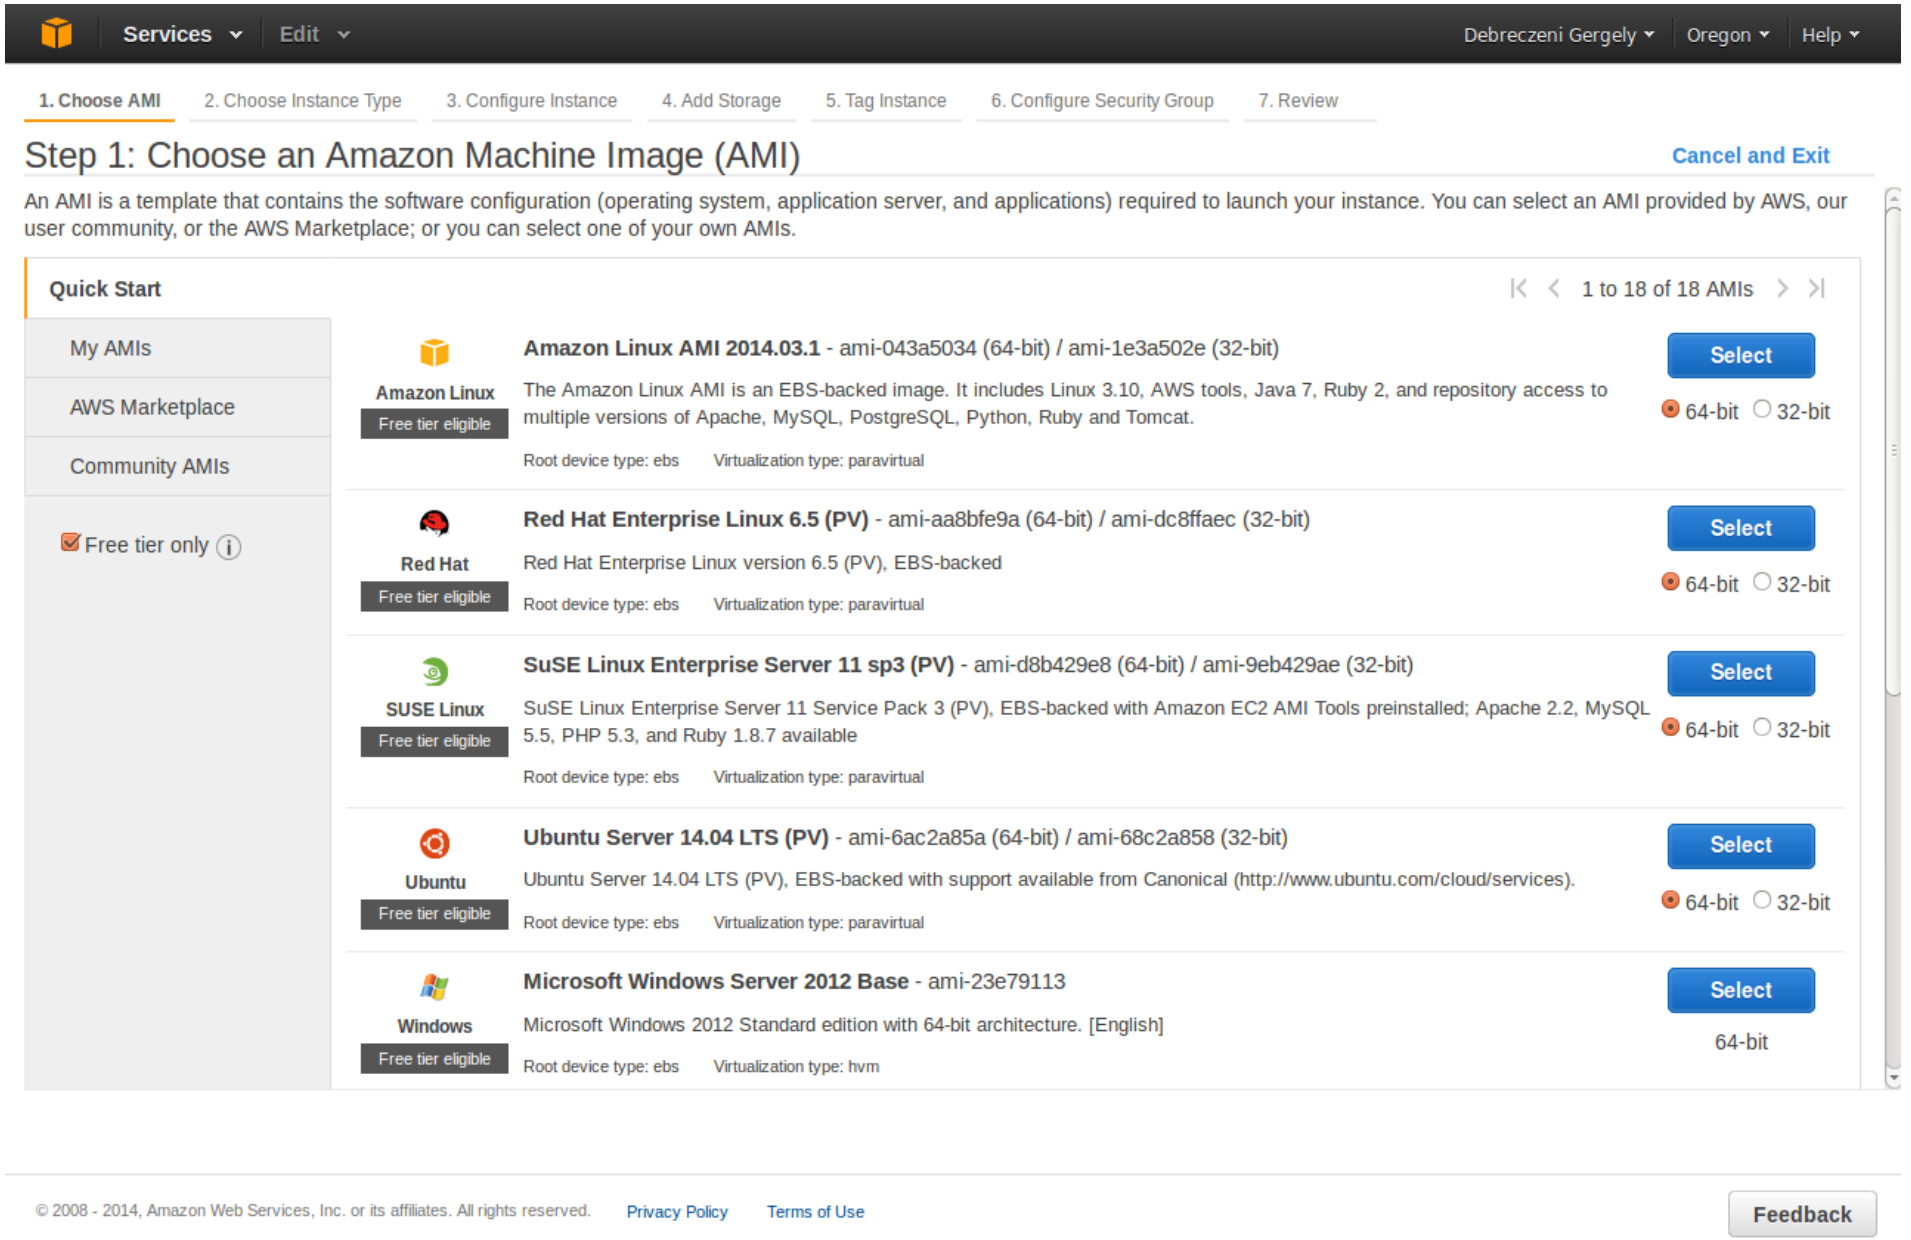
Services
Edit
Debreczeni Gergely
Oregon
Help
1. Choose AMI
2. Choose Instance Type
3. Configure Instance
4. Add Storage
5. Tag Instance
6. Configure Security Group
7. Review
Step 1: Choose an Amazon Machine Image (AMI)
Cancel and Exit
An AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. You can select an AMI provided by AWS, our user community, or the AWS Marketplace; or you can select one of your own AMIs.
Quick Start
1 to 18 of 18 AMIs
My AMIs
AWS Marketplace
Community AMIs
Free tier only
Amazon Linux AMI 2014.03.1 - ami-043a5034 (64-bit) / ami-1e3a502e (32-bit)
Amazon Linux
Free tier eligible
The Amazon Linux AMI is an EBS-backed image. It includes Linux 3.10, AWS tools, Java 7, Ruby 2, and repository access to multiple versions of Apache, MySQL, PostgreSQL, Python, Ruby and Tomcat.
Root device type: ebs Virtualization type: paravirtual
Select
64-bit 32-bit
Red Hat Enterprise Linux 6.5 (PV) - ami-aa8bfe9a (64-bit) / ami-dc8ffaec (32-bit)
Red Hat
Free tier eligible
Red Hat Enterprise Linux version 6.5 (PV), EBS-backed
Root device type: ebs Virtualization type: paravirtual
Select
64-bit 32-bit
SuSE Linux Enterprise Server 11 sp3 (PV) - ami-d8b429e8 (64-bit) / ami-9eb429ae (32-bit)
SUSE Linux
Free tier eligible
SuSE Linux Enterprise Server 11 Service Pack 3 (PV), EBS-backed with Amazon EC2 AMI Tools preinstalled; Apache 2.2, MySQL 5.5, PHP 5.3, and Ruby 1.8.7 available
Root device type: ebs Virtualization type: paravirtual
Select
64-bit 32-bit
Ubuntu Server 14.04 LTS (PV) - ami-6ac2a85a (64-bit) / ami-68c2a858 (32-bit)
Ubuntu
Free tier eligible
Ubuntu Server 14.04 LTS (PV), EBS-backed with support available from Canonical (http://www.ubuntu.com/cloud/services).
Root device type: ebs Virtualization type: paravirtual
Select
64-bit 32-bit
Microsoft Windows Server 2012 Base - ami-23e79113
Windows
Free tier eligible
Microsoft Windows 2012 Standard edition with 64-bit architecture. [English]
Root device type: ebs Virtualization type: hvm
Select
64-bit
© 2008 - 2014, Amazon Web Services, Inc. or its affiliates. All rights reserved.
Privacy Policy
Terms of Use
Feedback

Launching

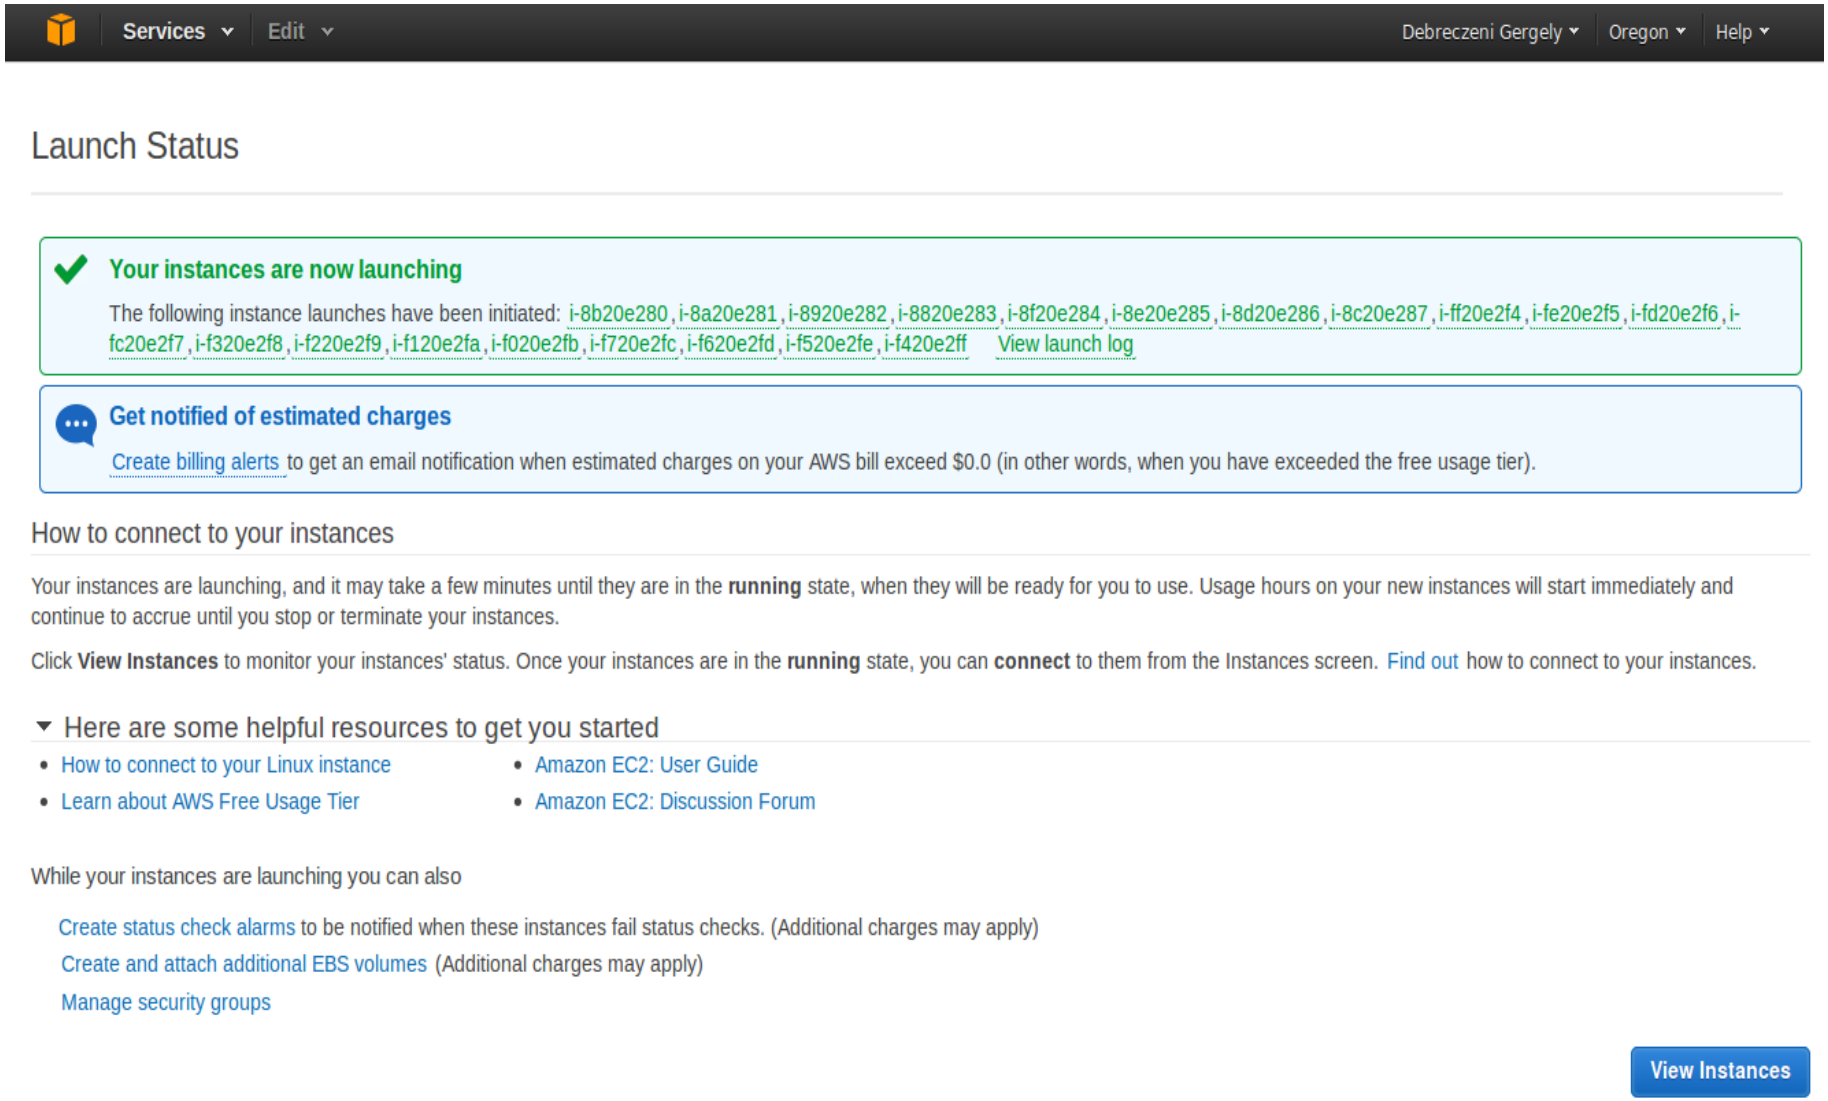
Services
Edit
Debreczeni Gergely
Oregon
Help
Launch Status
Your instances are now launching
The following instance launches have been initiated: i-8b20e280, i-8a20e281, i-8920e282, i-8820e283, i-8f20e284, i-8e20e285, i-8d20e286, i-8c20e287, i-ff20e2f4, i-fe20e2f5, i-fd20e2f6, i-fc20e2f7, i-f320e2f8, i-f220e2f9, i-f120e2fa, i-f020e2fb, i-f720e2fc, i-f620e2fd, i-f520e2fe, i-f420e2ff View launch log
Get notified of estimated charges
Create billing alerts to get an email notification when estimated charges on your AWS bill exceed $0.0 (in other words, when you have exceeded the free usage tier).
How to connect to your instances
Your instances are launching, and it may take a few minutes until they are in the running state, when they will be ready for you to use. Usage hours on your new instances will start immediately and continue to accrue until you stop or terminate your instances.
Click View Instances to monitor your instances' status. Once your instances are in the running state, you can connect to them from the Instances screen. Find out how to connect to your instances.
Here are some helpful resources to get you started
How to connect to your Linux instance
Learn about AWS Free Usage Tier
Amazon EC2: User Guide
Amazon EC2: Discussion Forum
While your instances are launching you can also
Create status check alarms to be notified when these instances fail status checks. (Additional charges may apply)
Create and attach additional EBS volumes (Additional charges may apply)
Manage security groups
View Instances

Status

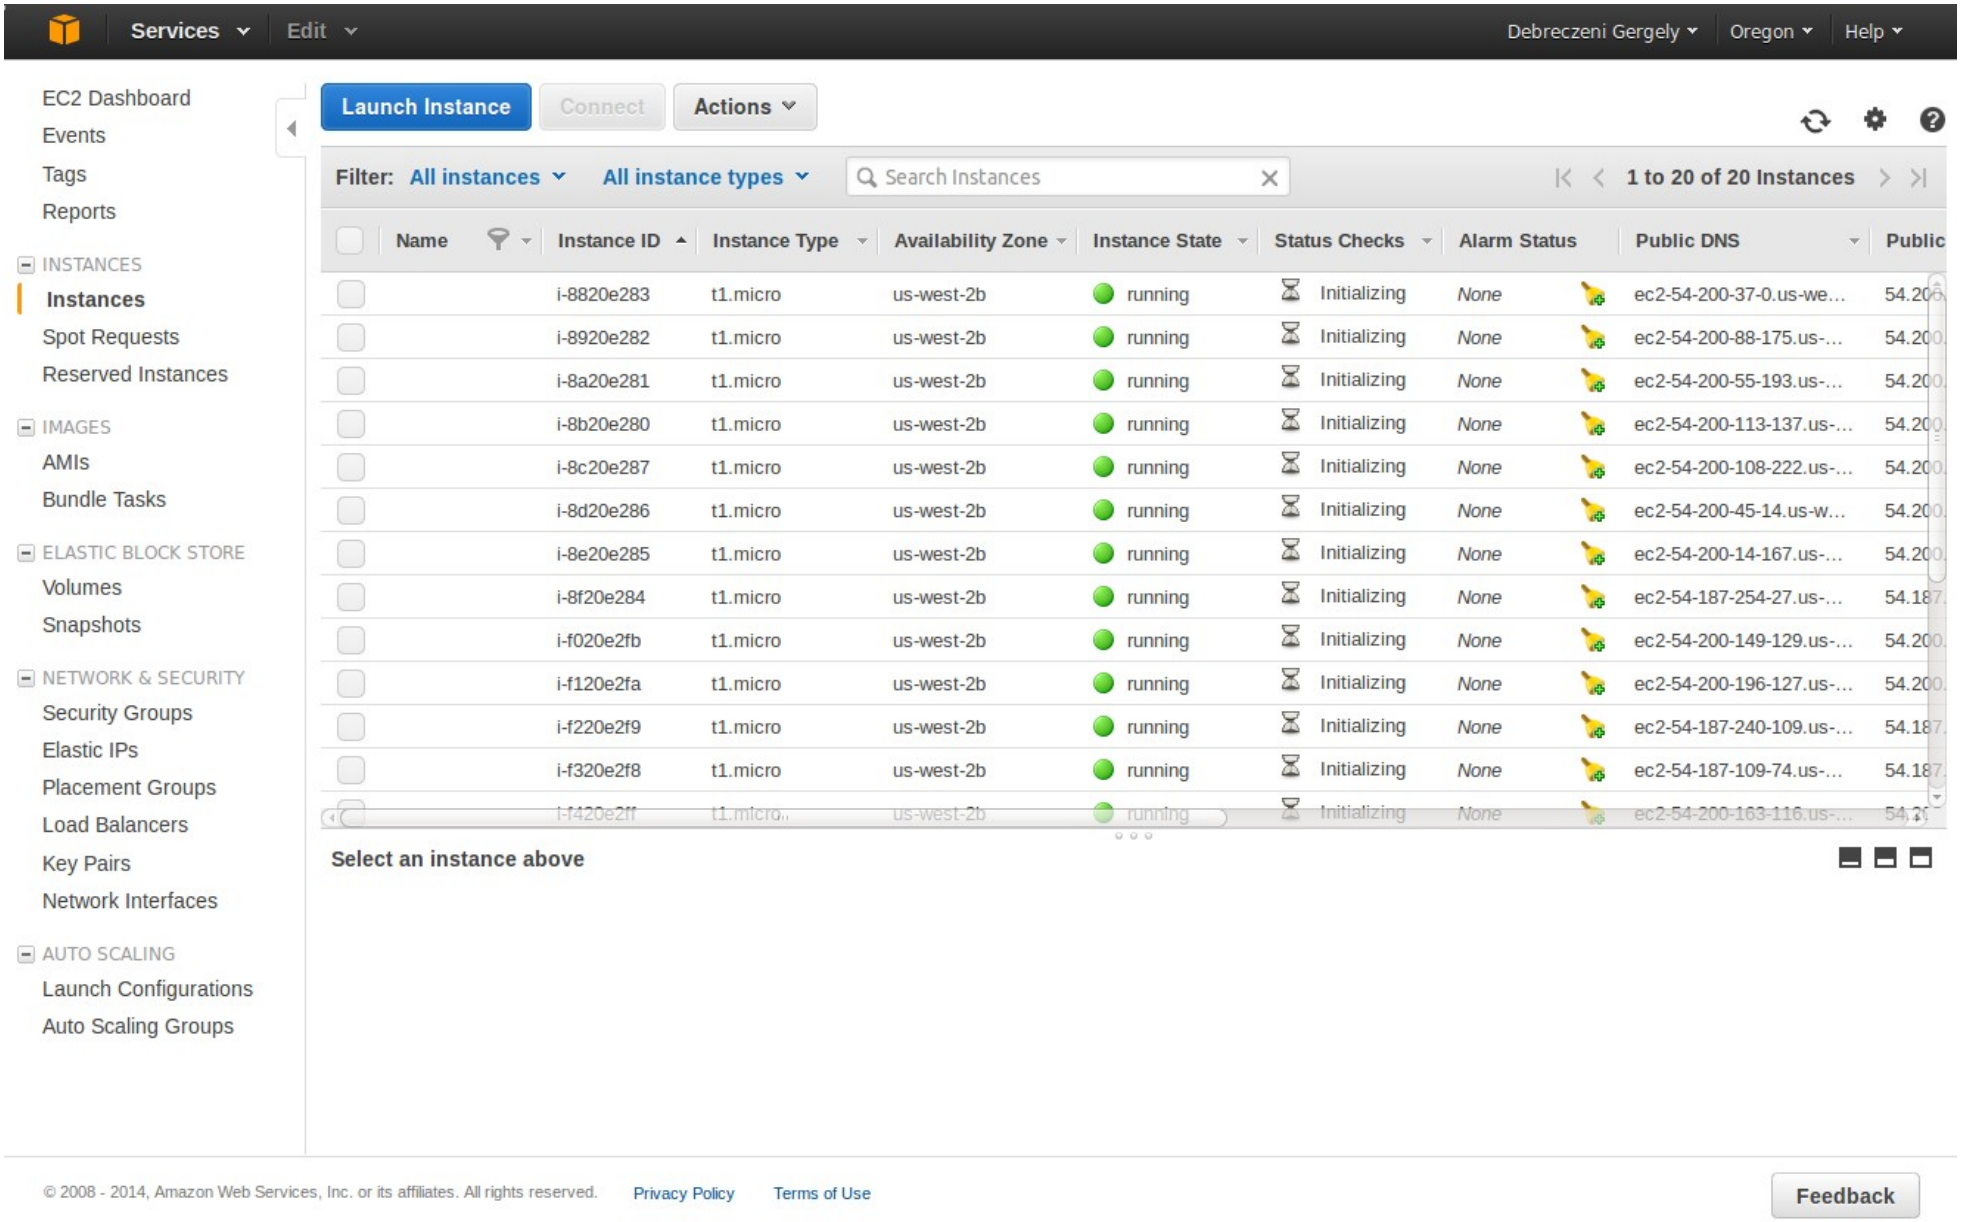
Services
Edit
Debreczeni Gergely
Oregon
Help
EC2 Dashboard
Events
Tags
Reports
INSTANCES
Instances
Spot Requests
Reserved Instances
IMAGES
AMIs
Bundle Tasks
ELASTIC BLOCK STORE
Volumes
Snapshots
NETWORK & SECURITY
Security Groups
Elastic IPs
Placement Groups
Load Balancers
Key Pairs
Network Interfaces
AUTO SCALING
Launch Configurations
Auto Scaling Groups
Launch Instance
Connect
Actions
Filter: All instances
All instance types
Search Instances
1 to 20 of 20 Instances
Name
Instance ID
Instance Type
Availability Zone
Instance State
Status Checks
Alarm Status
Public DNS
Public
i-8820e283 t1.micro us-west-2b running Initializing None ec2-54-200-37-0.us-we... 54.20
i-8920e282 t1.micro us-west-2b running Initializing None ec2-54-200-88-175.us-... 54.20
i-8a20e281 t1.micro us-west-2b running Initializing None ec2-54-200-55-193.us-... 54.20
i-8b20e280 t1.micro us-west-2b running Initializing None ec2-54-200-113-137.us-... 54.20
i-8c20e287 t1.micro us-west-2b running Initializing None ec2-54-200-108-222.us-... 54.20
i-8d20e286 t1.micro us-west-2b running Initializing None ec2-54-200-45-14.us-w... 54.20
i-8e20e285 t1.micro us-west-2b running Initializing None ec2-54-200-14-167.us-... 54.20
i-8f20e284 t1.micro us-west-2b running Initializing None ec2-54-187-254-27.us-... 54.187
i-f020e2fb t1.micro us-west-2b running Initializing None ec2-54-200-149-129.us-... 54.20
i-f120e2fa t1.micro us-west-2b running Initializing None ec2-54-200-196-127.us-... 54.20
i-f220e2f9 t1.micro us-west-2b running Initializing None ec2-54-187-240-109.us-... 54.187
i-f320e2f8 t1.micro us-west-2b running Initializing None ec2-54-187-109-74.us-... 54.187
Select an instance above
© 2008 - 2014, Amazon Web Services, Inc. or its affiliates. All rights reserved.
Privacy Policy
Terms of Use
Feedback

Login

```
gdebrecz@stan:/home/gdebrecz/Arbeits/Virgo/presentation/2014.05.26-VDASC-Amazon$
gdebrecz@stan:/home/gdebrecz/Arbeits/Virgo/presentation/2014.05.26-VDASC-Amazon$
gdebrecz@stan:/home/gdebrecz/Arbeits/Virgo/presentation/2014.05.26-VDASC-Amazon$ ssh -i /home/gdebrecz/.ssh/AmazonTestingKey.p
em 54.200.37.0 -l root
Please login as the user "ec2-user" rather than the user "root".


Connection to 54.200.37.0 closed.
gdebrecz@stan:/home/gdebrecz/Arbeits/Virgo/presentation/2014.05.26-VDASC-Amazon$ ssh -i /home/gdebrecz/.ssh/AmazonTestingKey.p
em 54.200.37.0 -l ec2-user
Last login: Sun May 25 09:37:02 2014 from 79.172.200.249
[ec2-user@ip-172-31-45-255 ~]$ sudo su
[root@ip-172-31-45-255 ec2-user]# cat /proc/cpuinfo
processor       : 0
vendor_id       : GenuineIntel
cpu family      : 6
model           : 45
model name      : Intel(R) Xeon(R) CPU E5-2650 0 @ 2.00GHz
stepping        : 7
cpu MHz         : 1795.672
cache size      : 20480 KB
fpu             : yes
fpu_exception   : yes
cpuid level     : 13
wp              : yes
flags           : fpu de tsc msr pae cx8 sep cmov pat clflush mmx fxsr sse sse2 ss ht syscall nx lm up rep_good unfair_spinloc
k pni pclmulqdq ssse3 cx16 pcid sse4_1 sse4_2 x2apic popcnt tsc_deadline_timer aes hypervisor lahf_lm
bogomips        : 3591.34
clflush size    : 64
cache_alignment : 64
address sizes   : 46 bits physical, 48 bits virtual
power management:

[root@ip-172-31-45-255 ec2-user]# exit
exit
[ec2-user@ip-172-31-45-255 ~]$ exit
logout
Connection to 54.200.37.0 closed.
gdebrecz@stan:/home/gdebrecz/Arbeits/Virgo/presentation/2014.05.26-VDASC-Amazon$ scrot -s login.png
```

Conclusion

Evaulation studies has to be done. But we cannot ignore this possibility.

- Benchmarking (To be done)
- Condor client joining the Pegasus pool (Done)
- Scalability (To be done)
- Input / Output transfer testing (To be done)