

Développement logiciel pour le Cloud (TLC)

7. Infrastructure-as-a-Service

Guillaume Pierre

Université de Rennes 1

Fall 2012

<http://www.globule.org/~gpierre/>

- 1 Introduction
- 2 Installing a Web Application in AWS
- 3 Network Management
- 4 Failover
- 5 Monitoring
- 6 Conclusion

- EBS: Elastic Block Store (a place to store AMIs)
- EC2: Elastic Compute Cloud
- S3: data storage
- Elastic IP addresses
- Elastic load balancing
- AutoScale
- Etc.

How do they all fit together?

Table of Contents

- 1 Introduction
- 2 Installing a Web Application in AWS**
- 3 Network Management
- 4 Failover
- 5 Monitoring
- 6 Conclusion

Select an AMI

Request Instances Wizard

Cancel


The bookmark that was activated refers to the AMI below. Please review...

AMI Details

Image Id: ami-1aad5273
Owner:
Manifest: 099720109477/ebs/ubuntu-images/ubuntu-natty-11.04-amd64-server-20110426
Platform:  Ubuntu
Architecture: x86_64
Root Device Type: ebs

Attached Block Devices

| Device Name | Volume Size |
|-------------|-------------|
| /dev/sda1 | 8 GB |

[Continue](#) 

Choose the instance types

Request Instances Wizard

Cancel X

CHOOSE AN AMI **INSTANCE DETAILS** CREATE KEY PAIR CONFIGURE FIREWALL REVIEW

Provide the details for your instance(s). You may also decide whether you want to launch your instances as "on-demand" or "spot" instances.

Number of Instances: **Availability Zone:**

Instance Type:

Launch Instances

EC2 Instances let you pay for compute capacity by the hour with no long term commitments. This transforms what are commonly large fixed costs into much smaller variable costs.

Request Spot Instances

Launch Instances Into Your Virtual Private Cloud

[< Back](#) [Continue >](#)

Instance types at AWS

| Type | CPU cores | Memory | Storage | Price |
|------------|------------|----------|----------|--------------|
| t1.micro | up to 1 | 613 MiB | — | \$0.020/hour |
| m1.small | 1 (slow) | 1.7 GiB | 160 GB | \$0.065/hour |
| m1.medium | 1 (medium) | 3.75 GiB | 410 GB | \$0.130/hour |
| m1.large | 2 (medium) | 7.5 GiB | 850 GB | \$0.260/hour |
| m1.xlarge | 4 (medium) | 15 GiB | 1,690 GB | \$0.520/hour |
| m3.xlarge | 4 (fast) | 15 GiB | — | \$0.580/hour |
| m3.2xlarge | 8 (fast) | 30 GiB | — | \$1.160/hour |

- Also:

- ▶ High-memory instances
- ▶ High-CPU instances
- ▶ Cluster compute instances (with dedicated networking)
- ▶ Cluster GPU instances
- ▶ High I/O instances

Request Instances Wizard

[Cancel](#)[CHOOSE AN AMI](#)**INSTANCE DETAILS**[CREATE KEY PAIR](#)[CONFIGURE FIREWALL](#)[REVIEW](#)**Number of Instances:** 1**Availability Zone:** No Preference

Advanced Instance Options

Here you can choose a specific [kernel](#) or [RAM disk](#) to use with your instances. You can also choose to enable CloudWatch Detailed Monitoring or enter data that will be available from your instances once they launch.

Kernel ID:

Use Default

RAM Disk ID:

Use Default

Monitoring: Enable CloudWatch detailed monitoring for this instance
(additional charges will apply)**User Data:** as text as file base64 encoded**Termination Protection:** Prevention against accidental termination.**Shutdown Behavior:**

Stop

Choose the behavior when the instance is shutdown from within the instance.

[Back](#)[Continue](#)

Request Instances Wizard

[Cancel](#)

CHOOSE AN AMI

INSTANCE DETAILS

CREATE KEY PAIR

CONFIGURE FIREWALL

REVIEW

Add tags to your instance to simplify the administration of your EC2 infrastructure. A form of metadata, tags consist of a case-sensitive key/value pair, are stored in the cloud and are private to your account. You can create user-friendly names that help you organize, search, and browse your resources. For example, you could define a tag with key = Name and value = Webserver. You can add up to 10 unique keys to each instance along with an optional value for each key. For more information, go to [Using Tags](#) in the *EC2 User Guide*.

| Key (127 characters maximum) | Value (255 characters maximum) | Remove |
|------------------------------|--|--------|
| Name | <input type="text" value="Webserver"/> | |
| <input type="text"/> | <input type="text"/> | |

[Add another Tag.](#) (Maximum of 10)

[< Back](#)

[Continue >](#)

Request Instances Wizard

Cancel 

CHOOSE AN AMI

INSTANCE DETAILS

CREATE KEY PAIR

CONFIGURE FIREWALL

REVIEW

Public/private key pairs allow you to securely connect to your instance after it launches. To create a key pair, enter a name and click **Create & Download your Key Pair**. You will then be prompted to save the private key to your computer. Note, you only need to generate a key pair once - not each time you want to deploy an Amazon EC2 instance.

Choose from your existing Key Pairs

Create a new Key Pair

1. Enter a name for your key pair:* (e.g., jdoekey)

2. Click to create your key pair:*



Create & Download your Key Pair

Save this file in a place you will remember. You can use this key pair to launch other instances in the future or visit the Key Pairs page to create or manage existing ones.

Proceed without a Key Pair

[< Back](#)

Continue 

Don't lose your keypair! You cannot retrieve it afterwards Source:

<http://www.jasongrimes.org>

Create a security group

Request Instances Wizard Cancel


CHOOSE AN AMI INSTANCE DETAILS CREATE KEY PAIR **CONFIGURE FIREWALL** REVIEW

Security groups determine whether a network port is open or blocked on your instances. You may use an existing security group, or we can help you create a new security group to allow access to your instances using the suggested ports below. Add additional ports now or update your security group anytime using the Security Groups page.

Choose one or more of your existing Security Groups

Create a new Security Group

Group Name

Group Description
 This field is required


Inbound Rules

Create a new rule:

Port range:
(e.g., 80 or 49152-65535)

Source:
(e.g., 192.168.2.0/24, sg-47ad482e, or 1234567890/default)

| TCP | Port (Service) | Source | Action |
|-----|----------------|-----------|--------|
| | 80 (HTTP) | 0.0.0.0/0 | Delete |
| | 443 (HTTPS) | 0.0.0.0/0 | Delete |
| | 22 (SSH) | 0.0.0.0/0 | Delete |


< Back Continue 

Request Instances Wizard

Cancel

CHOOSE AN AMI INSTANCE DETAILS CREATE KEY PAIR CONFIGURE FIREWALL **REVIEW**

Please review the information below, then click **Launch**.

AMI:  Ubuntu AMI ID ami-1aad5273 (x86_64) [Edit AMI](#)

Number of Instances: 1

Availability Zone: No Preference

Instance Type: Micro (t1.micro)

Instance Class: On Demand [Edit Instance Details](#)

Monitoring: Disabled **Termination Protection:** Disabled

Tenancy: Default


Kernel ID: Use Default **Shutdown Behavior:** Stop

RAM Disk ID: Use Default

User Data: [Edit Advanced Details](#)

Key Pair Name: awskey [Edit Key Pair](#)

Security Group(s): sg-2084d049 [Edit Firewall](#)

[< Back](#) [Launch](#) 

SSH access to your instances

```
chmod 400 ~/aws/awskey.pem # Make the private key readable only by you  
ssh -i ~/aws/awskey.pem ubuntu@your.ec2.hostname
```

- Update already-installed packages:

```
$ sudo apt-get update
```

```
$ sudo apt-get upgrade -y
```


- Install a LAMP stack:

```
$ sudo apt-get install apache2 libapache2-mod-php5  
php5-mysql mysql-server
```

- Then, install your LAMP application as usual...

Save your customized AMI

- Amazon provides standard Amazon Machine Images (AMIs)
 - ▶ Standard systems: Linux, Windows, etc.
 - ▶ You probably want an AMI which already contains the software you will need (web server, libraries, etc)
- Once you customized a VM instance you can create a custom AMI



- You can also create a new AMI from scratch
 - ▶ But that's complex and rarely necessary

Table of Contents

- 1 Introduction
- 2 Installing a Web Application in AWS
- 3 Network Management**
- 4 Failover
- 5 Monitoring
- 6 Conclusion

- Problem: each time you start a VM instance it gets a new IP address
 - ▶ Can we always start VMs at the same address?
- Solution: **Elastic IP addresses**
 - ▶ Each user can request up to 5 elastic IPs
 - ▶ You “own” the IP address independently from the VM instances
 - ▶ You can attach any elastic IP address to any VM instance
 - ★ Even across multiple availability zones
 - ▶ Elastic IPs are free when they are in use, \$0.01/hour otherwise

- Create a new elastic IP address:

```
$ ec2-allocate-address  
ADDRESS 75.101.137.243
```

- Attach an elastic IP to an instance:

```
$ ec2-associate-address -i i-07612d6e 75.101.137.243  
ADDRESS 75.101.137.243 i-07612d6e
```

- Register the elastic IP in DNS:

```
db.example.com.    CNAME    ec2-75-101-137-243.compute-1.amazonaws.com.
```

Table of Contents

- 1 Introduction
- 2 Installing a Web Application in AWS
- 3 Network Management
- 4 Failover**
- 5 Monitoring
- 6 Conclusion

- Individual VM instances may fail
 - ▶ Because the hardware machine failed
 - ▶ Or the guest OS crashed
 - ▶ Etc.
- Entire data centers may fail
 - ▶ Power outage
 - ▶ Hurricane
 - ▶ Etc.
- Availability zones should not fail at the same time
 - ▶ Entirely independent EC2 systems...
 - ▶ Availability zone failures happen roughly once per year

 **Instagram Support**
@InstagramHelp

[Follow](#) 

We're currently experiencing technical difficulties and we're working to correct the issues. Thanks for your patience

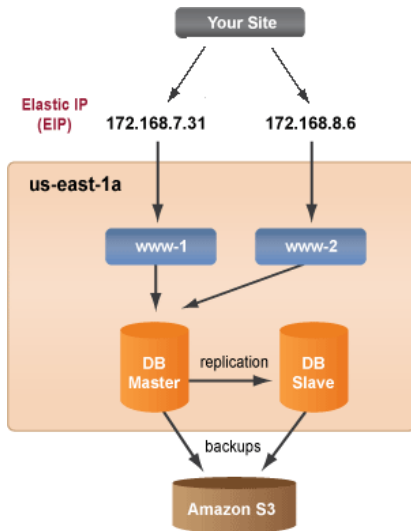
[Reply](#) [Retweet](#) [Favorite](#)

27,956 RETWEETS **1,581** FAVORITES

11:16 PM - 29 Jun 12 via Twitter for Mac · [Embed this Tweet](#)

Storm ⇒ power cut ⇒ availability zone failure ⇒ Instagram failure

Designing against VM instance failures



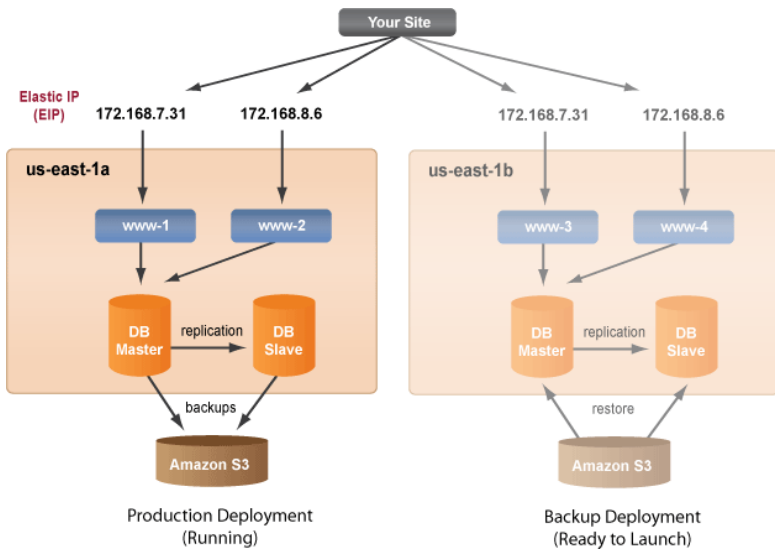
Source: <http://support.rightscale.com>

Designing against Availability Zone failures

| Architecture | Cost | Downtime | Comment |
|--------------|--------|------------------|---|
| Basic | Low | At least 5-10 mn | EIPs with 1 Avail. Zone and a backup deployment (ready to launch) |
| Intermediate | Medium | Minimal | EIPs with 2 Avail. Zones (few instances in zone 2) |
| Advanced | High | None | EIPs with 2 Avail. Zones (duplicate setup in zone 2) |

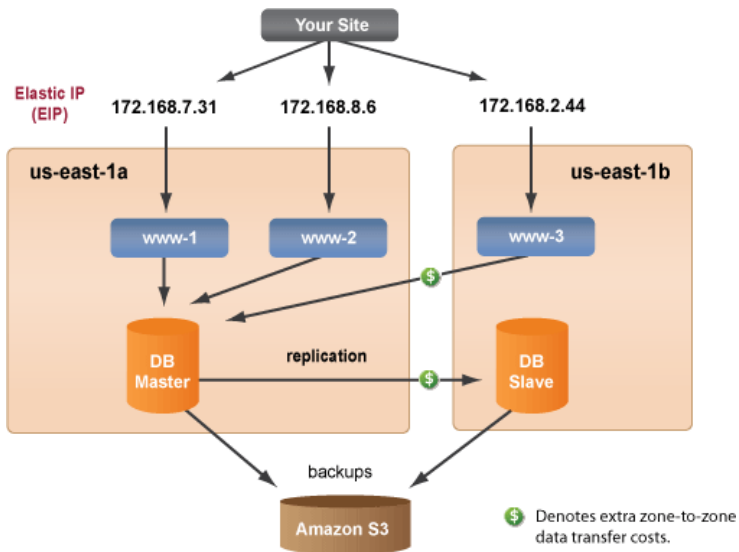
Source: <http://support.rightscale.com>

Basic multi-zone setup




Source: <http://support.rightscale.com>

Intermediate multi-zone setup



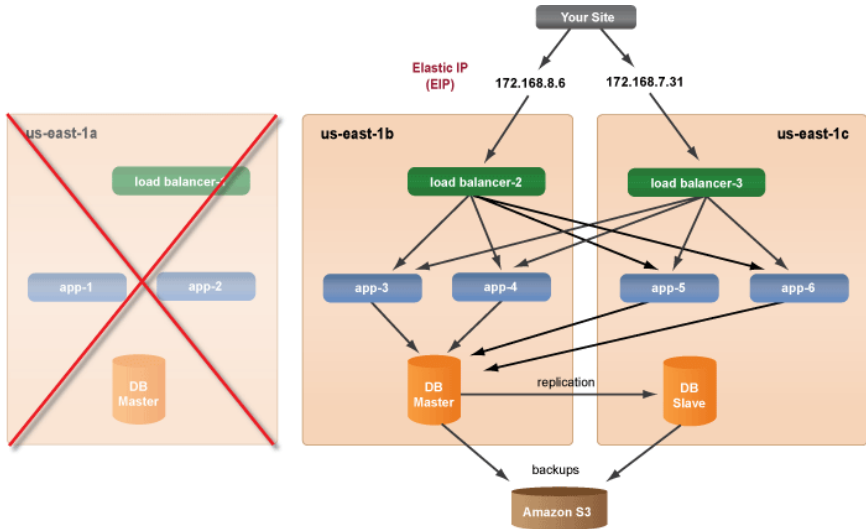
Source: <http://support.rightscale.com>

Advanced multi-zone setup



Source: <http://support.rightscale.com>

Advanced multi-zone setup after a zone failure



Source: <http://support.rightscale.com>

Table of Contents

- 1 Introduction
- 2 Installing a Web Application in AWS
- 3 Network Management
- 4 Failover
- 5 Monitoring**
- 6 Conclusion

Monitoring your running instances with CloudWatch

Request Instances Wizard

Cancel

CHOOSE AN AMI **INSTANCE DETAILS** CREATE KEY PAIR CONFIGURE FIREWALL REVIEW


Number of Instances: 1
Availability Zone: No Preference

Advanced Instance Options

Here you can choose a specific **kernel** or **RAM disk** to use with your instances. You can also choose to enable CloudWatch Monitoring or enter data that will be available from your instances once they launch.

Kernel ID: Use Default ▾

RAM Disk ID: Use Default ▾

Monitoring: Enable CloudWatch Monitoring for this instance
(additional charges will apply) 

User Data:

base64 encoded

◀ Back Continue ▶

Metrics monitored by CloudWatch

- CPU Utilization (%)
- Memory Utilization (%)
- Network Out Utilization (MB)
- Memory Used (MB)
- Memory Available (MB)
- Swap Utilization (%)
- Swap Used (MB)
- Disk Space Utilization (%)
- Disk Space Used (GB)
- Disk Space Available (GB)
- Etc.

Metrics monitored by CloudWatch

- CPU Utilization (%)
- Memory Utilization (%)
- Network Out Utilization (MB)
- Memory Used (MB)
- Memory Available (MB)
- Swap Utilization (%)
- Swap Used (MB)
- Disk Space Utilization (%)
- Disk Space Used (GB)
- Disk Space Available (GB)
- Etc.
- You can also **create your own metric...**

Granularity: refreshed every 5 minutes

AutoScaling

Your CloudWatch Alarms

Create Alarm Modify Delete Show/Hide Refresh Help

Viewing: All alarms 1 to 2 of 2 Items

| | State | Name | Threshold |
|-------------------------------------|-------|------------------|---------------------------------------|
| <input checked="" type="checkbox"/> | OK | scale-up-alarm | httpd-busyworkers > 10 for 10 minutes |
| <input type="checkbox"/> | ALARM | scale-down-alarm | httpd-busyworkers < 9 for 10 minutes |

1 Alarm selected

Alarm: scale-up-alarm

Description Metric History

State Details:
State changed to ✔ OK' at 2012/11/05 16:36 UTC.
Reason: Threshold Crossed: 1 datapoint (9.0) was not greater than the threshold (10.0).

Description:
-

Threshold:
httpd-busyworkers > 10 for 10 minutes

Actions:
in ALARMSend message to topic "AutoScaling"
state - (juan.domenech@gmail.com)
Use policy "scale-up-prueba (Add 1 instance)" for group "grupo-prueba"

scale-up-alarm

httpd-busyworkers > 10

| Date/Time | Value |
|------------|-------|
| 11/4 16:00 | 0 |
| 11/5 00:00 | 0 |
| 11/5 08:00 | 0 |
| 11/5 16:00 | 85 |

Table of Contents

- 1 Introduction
- 2 Installing a Web Application in AWS
- 3 Network Management
- 4 Failover
- 5 Monitoring
- 6 Conclusion**

- Cloud computing may seem intimidating
 - ▶ Lots of different systems and APIs
 - ▶ Business concerns (vendor lock-in, pricing strategies etc.)
 - ▶ Low-level programming
 - ▶ Security
 - ▶ Performance

- Cloud computing may seem intimidating
 - ▶ Lots of different systems and APIs
 - ▶ Business concerns (vendor lock-in, pricing strategies etc.)
 - ▶ Low-level programming
 - ▶ Security
 - ▶ Performance

- Don't be scared!
 - ▶ Act rationally
 - ▶ Follow a few common-sense tips

Do not solve problems which do not exist

- Cloud platforms propose lots of sophisticated services
 - ▶ But they are not all useful for all applications
 - ▶ Many tools are designed for very specific use cases
- Do not solve problems before making sure that the problem exists!
 - ▶ Better start simple, and expand later if necessary

- The more you know your application the easier it is to address its needs
 - ▶ What does the application do?
 - ▶ Which data does it use?
- Study the non-functional properties carefully
 - ▶ What kind of queries does the application issue to its database?
 - ▶ How computationally intensive is each request's processing?
 - ▶ How frequent is each request?
 - ▶ How much scalability do we need?
 - ▶ How much fault-tolerance do we need?
 - ▶ Cost analysis

Choose the right tool for the right task

- Different Cloud systems offer different types of functionality
 - ▶ Choose your platform carefully!
 - ▶ Moving from one platform to another is doable but complex and costly
- Do not be afraid to mix and match
 - ▶ Example: start with a PaaS environment, add a few extra IaaS resources
 - ▶ Some people even mix and match services from different clouds (AppEngine + AWS S3)

- **Standardization**

- ▶ Ongoing efforts on the IaaS layer (APIs, VM image formats, etc.)
- ▶ Live VM migration across different clouds

- **Service variety**

- ▶ Cloud operators will keep on inventing new types of services
- ▶ (Useful ones or not)

- **Price decrease**

- ▶ More competition \Rightarrow lower prices