



August, 21 - 14:15h

Multi Region failover scenarios in the Cloud

How to turn the cloud into a safe place?



DevOps Track

Multi Region failover scenarios in the Cloud

ACQUIATM

Ricardo Amaro

Cloud Systems Engineer

Operations Team

ricardo.amaro@acquia.com



Hosting in...

the Cloud



Agenda

- Why the cloud?
- About Cloud Hosting
- Aws Regions
- The Failover
- DNS - CDN
- File replication
- Database replication
- Complete Setup
- Basic demo
- Other Software



Bonus agenda: Basic Demo

01

Multiregion-Shot recipe:

- 2 cloud servers (AWS us-east)
- 2 cloud servers (AWS eu-west)
- Ubuntu/Linux
- Gluster + Unison
- Lamp Stack
- Drupal
- Tungsten





Why the cloud?



Scaling Drupal

"It is not always easy to scale Drupal -- not because Drupal sucks, but simply because scaling the LAMP stack (including Drupal) takes no small amount of skill."

Dries Buytaert






in a Cloud
infrastructure...






..or moving
Infrastructure?





A site is a living thing,
always changing
and evolving!





A spike can ruin the day





but a bad setup
can ruin the
month



```
HTTP/1.1 503 Service Unavailable
Date: Mon, 04 May 2012 22:32:02 GMT
Server: Nginx
Cache-Control: max-age=1
Expires: Mon, 04 May 2012 22:32:03 GMT
Content-Length: 2615
Connection: close
Content-Type: text/html
```




Are you taking
chances?





Cloud computing

- **Infrastructure as a Service (IaaS)**
- Platform as a Service (PaaS)
- Software as a Service (SaaS)



Cloud: Regions

US West
(Northern California)

US West
(Oregon)

US East
(Northern Virginia)

South America
(São Paulo)

Europe
(Ireland)

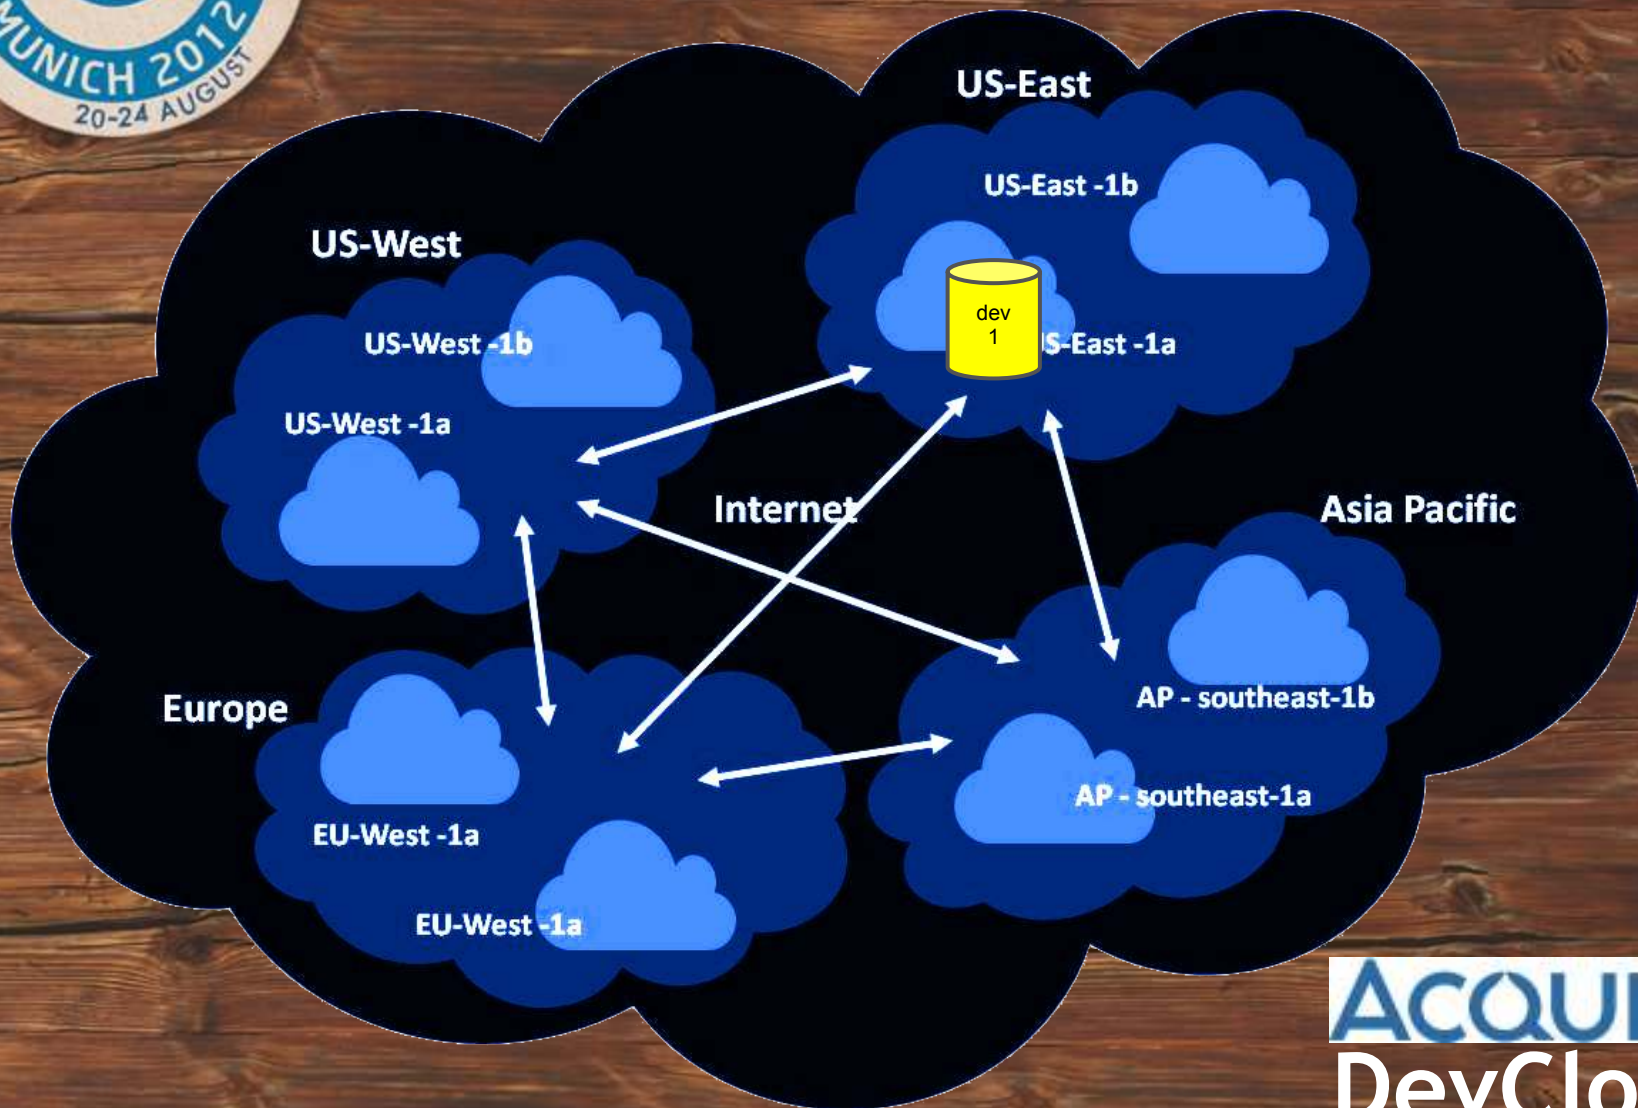
Asia Pacific
(Singapore)

Asia Pacific
(Tokyo)



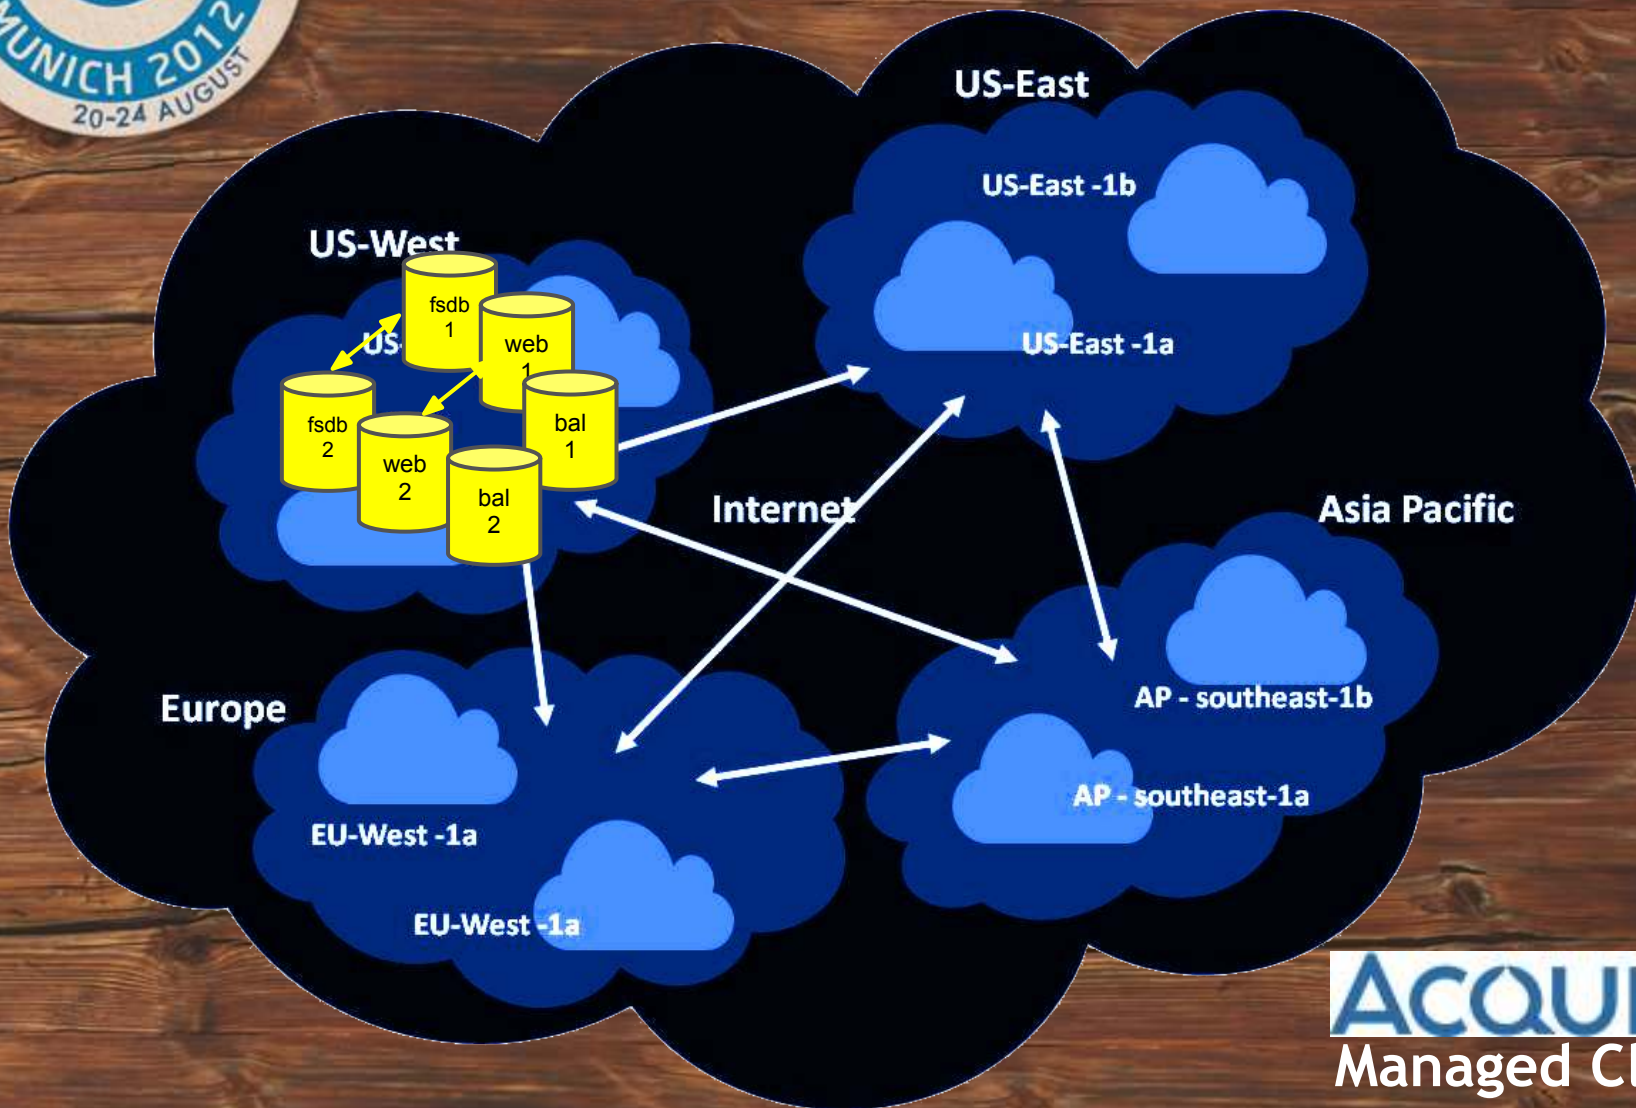


Availability Zones & Regions



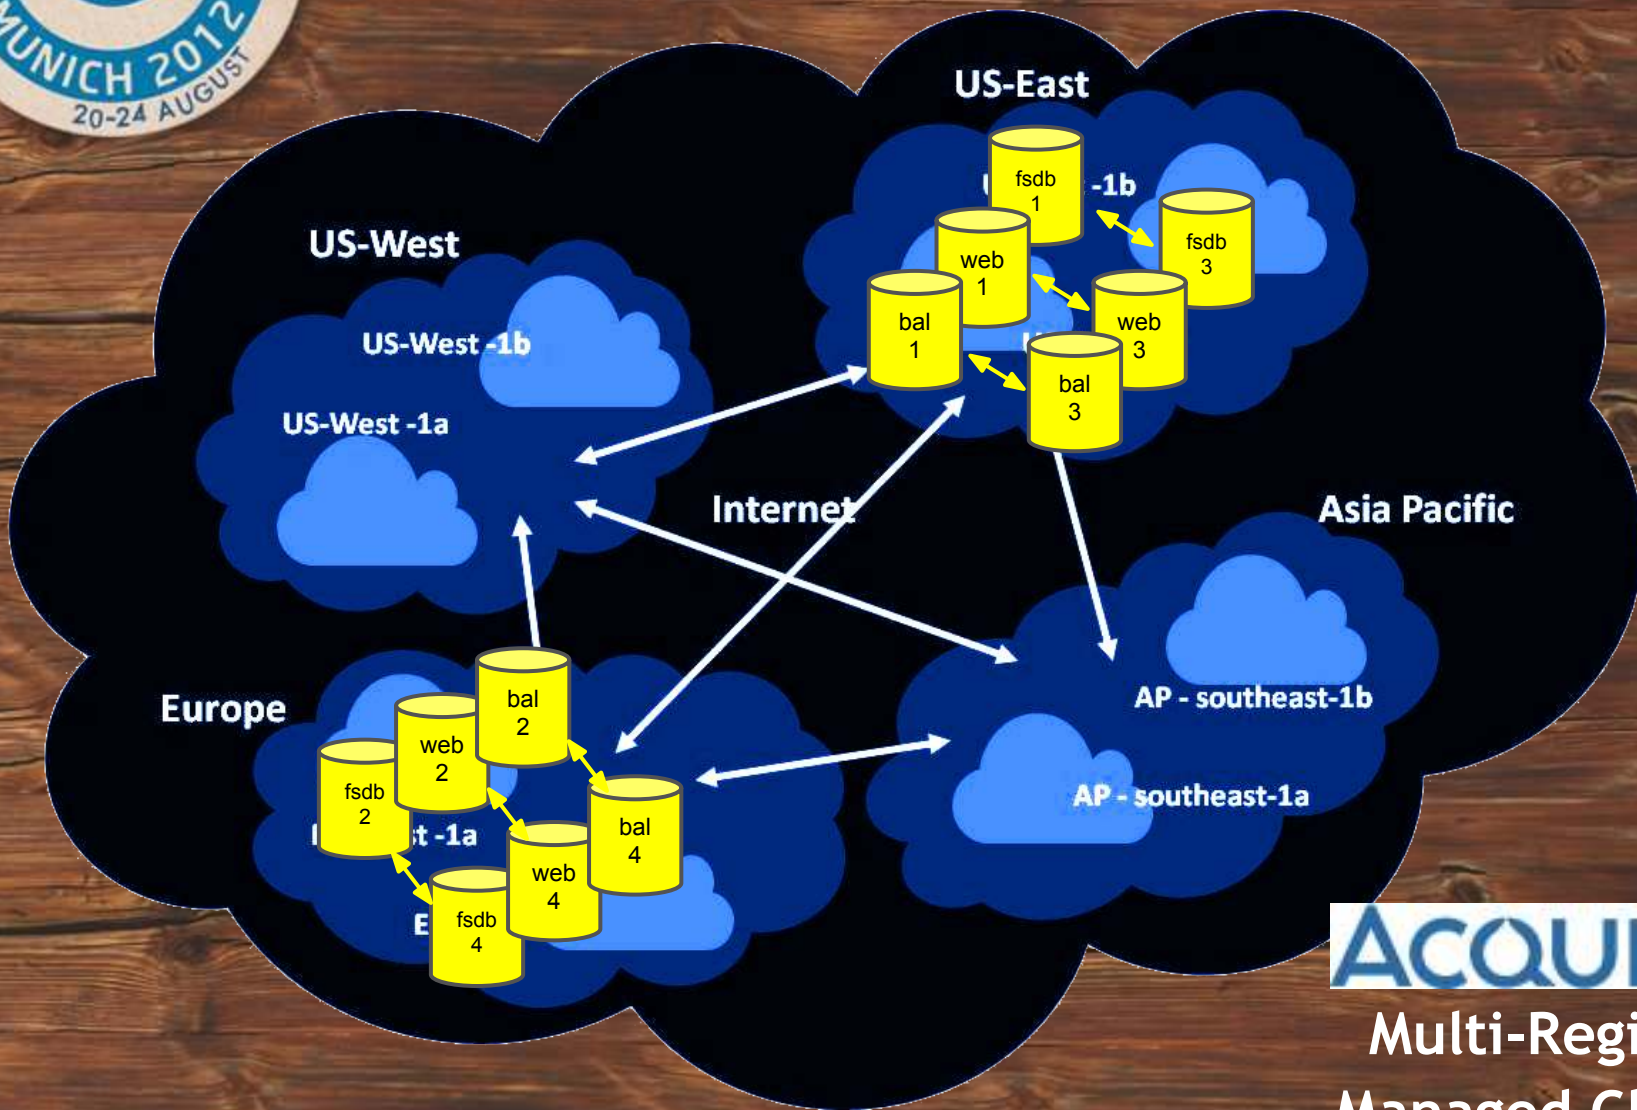


Availability Zones & Regions





Availability Zones & Regions



AcquiaTM

**Multi-Region
Managed Cloud**



The Cloud Providers



Cloud providers

Amazon Ec2
 Reliacloud
 GoDaddy
 Softlayer
 OpSource
 Rackspace
 GoGrid
 Nephoscale
 Hosting.com

A	B	N	O	P
Provider	Pricing	Instance Types	Data Transfer out (/GB)	Data Transfer in (/GB)
Amazon EC2	Pay-as-you-go	12	0.12	0
BitRefinery		Configurable	0	0
GoDaddy		5	0	0
GoGrid	Pay-as-you-go or Monthly	1	0.29	0
Hosting.com	Monthly	Configurable	0	0
NephoScale	Pay-as-you-go or Monthly	4	6	0.13
OpSource	Pay-as-you-go or Monthly	No	4	Configurable
Rackspace	Pay-as-you-go	No	8	8
ReliaCloud	Monthly	No	5	5
Softlayer	Pay-as-you-go or Monthly	No	6	13
Terremark*	Pay-as-you-go	No	5	Configurable

Source: <http://www.techrepublic.com/blog/datacenter/11-cloud-iaas-providers-compared/5285>



Dedicated Drupal Cloud Hosting



Acquia

01

Dedicated Drupal Hosting

Acquia™



Acquia Managed Cloud: By The Numbers

**600 customers
on 2300 Servers**

**7.7 BILLION+ requests
served per month**

**1.4 BILLION+ page views
per month**

**133+ terabytes of content
delivered per month**

**90M
requests in
one day**

Large Media &
Entertainment
organization

**73M
requests in
one day**

Global political
magazine

**51M
requests in
one day**

Large news
organization

Acquia™ Cloud

Application

Drupal™

LAMP Stack



ubuntu®



Virtual Infrastructure
Network
Physical Media
Physical Security

Security

Availability

Compliance

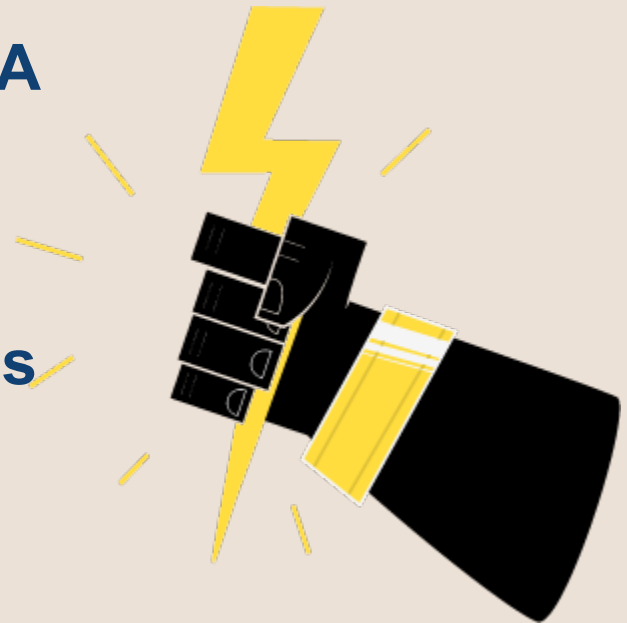


Pantheon

New dedicated Drupal Cloud Hosting.

- Recently raised \$5M in Series A financing, led by the Foundry Group.
- Targeting on SaaS (Software as a Service) based on Drupal.

Promising!





Omega8.cc

Hosting startup was born in 2009.

- Platform is based on AEGIR and supports it.
- Actively contributing to Barracuda and Octopus installers.
- Flag on speed and performance.



Follow the discussion on:

<http://groups.drupal.org/node/208123>



The Clouds Outages

The hard journey of hosting



**A very large
number of
companies
don't have a
disaster
recovery plan**





21 April 2011

Major outage causing a huge lack of confidence on AWS.

Providers like Rackspace take the opportunity to promote their services.



December 2011

Mass preventive relaunches due to instance degradation.



June 2012

Massive outage damaging hundreds of data volumes and instances.

US Virginia's datacenter hit by a storm.




Plan for failure!




Production scenario



Acquia™



DRUPAL 

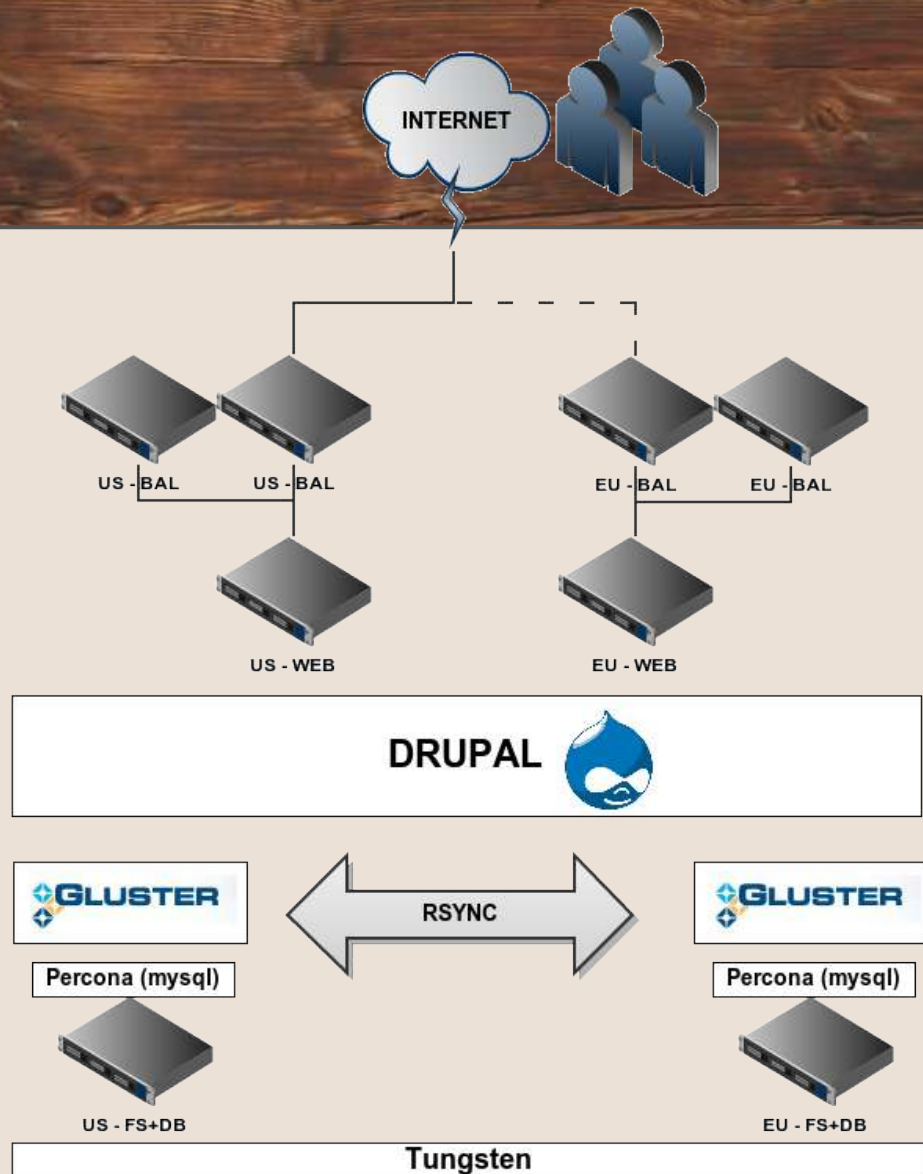




Staging Scenario



Staging Scenario

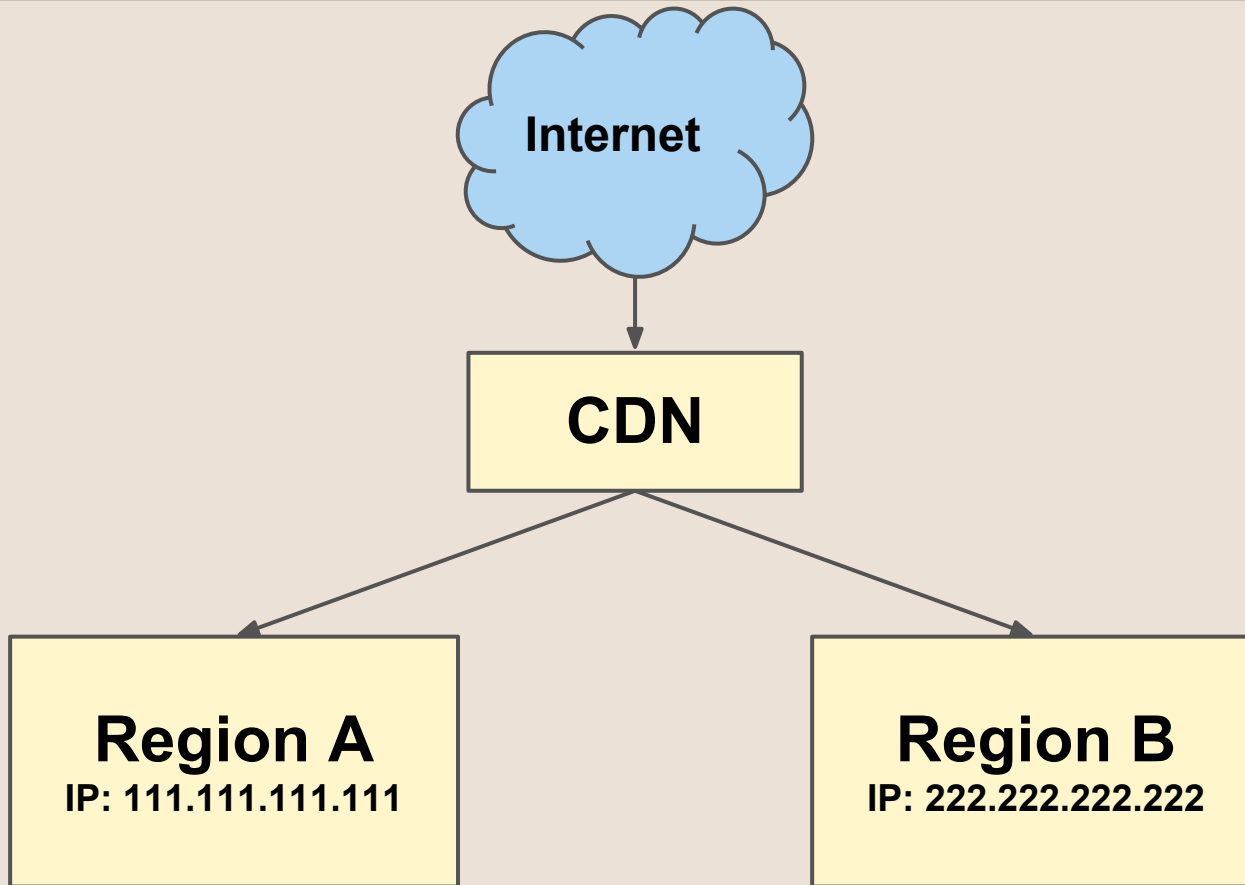




The Failover

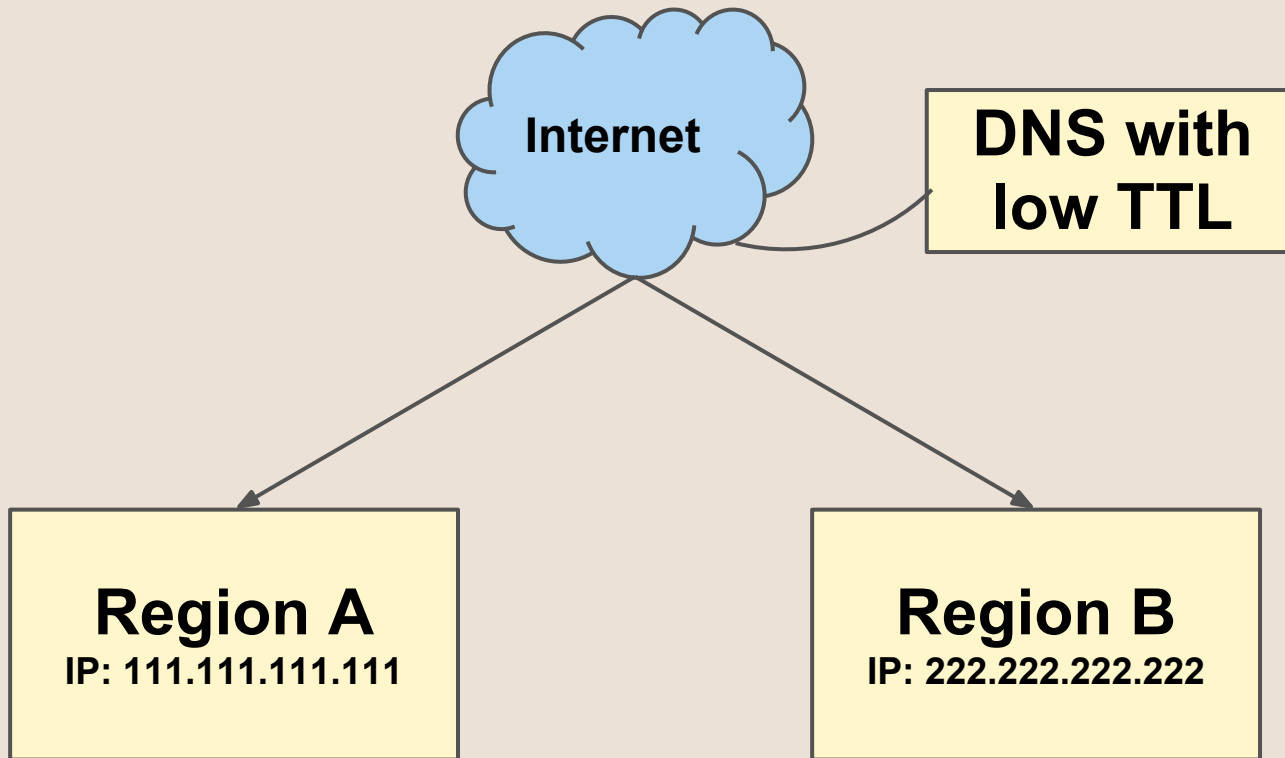


DNS / CDN





DNS / CDN





EIP / DNS

Always use Elastic IP's

- EIP's guarantee instance IP and names stick on restart.
- Use the EIP internal name for consistent DNS results.



The Replication Software



GlusterFS



GlusterFS is a scale-out NAS file system. It is free software, licensed under the GNU GPL v3. It aggregates various storage servers over Ethernet or Infiniband RDMA interconnect into one large parallel network file system. It has found a variety of applications including cloud computing, biomedical sciences and archival storage.



GlusterFS

SERVER 1:

```
$sudo mkdir /mnt/brick1
$sudo mkdir /mnt/gfs
$sudo apt-get -y install glusterfs-server
$sudo service glusterfs-server start
```

SERVER 2

```
$sudo mkdir /mnt/brick2
$sudo mkdir /mnt/gfs
$sudo apt-get -y install glusterfs-server
$sudo service glusterfs-server start
$sudo gluster peer probe SERVER-1
$sudo gluster peer status
```

SERVER 1:

```
$sudo gluster peer probe SERVER-1
$sudo gluster peer status
```




GlusterFS

```
$sudo gluster volume create volumel replica 2 ec2-23-20-23-207.compute-1.amazonaws.com:/mnt/brick1 ec2-184-72-153-229.compute-1.amazonaws.com:/mnt/brick2
$sudo gluster volume info
$sudo gluster volume start volumel
$sudo mount -t glusterfs /etc/glusterd/vols/volumel/volumel-fuse.vol /mnt/gfs/ -o rw,allow_other,default_permissions,max_read=131072,_netdev
```

SERVER 2:

```
$sudo mount -t glusterfs /etc/glusterd/vols/volumel/volumel-fuse.vol /mnt/gfs/ -o rw,allow_other,default_permissions,max_read=131072,_netdev
```

FSTAB (for boot mount)

```
/etc/glusterd/vols/volumel/volumel-fuse.vol /mnt/gfs/ fuse.
glusterfs rw,allow_other,default_permissions,max_read=131072 0
0
```



Video

P2- Gluster Install



GlusterFS

Prepare EBS volumes for Performance

- Network connected storage
 - 1Gb/sec assuming no noisy neighbors
- Per volume limit of ~20MB/sec for I/O => ~32KB
 - Slows to ~2MB/sec @ I/O ~2KB
- Always build mdadm arrays of 8 devices
 - ~80MB/sec with a single, 8 volume array.



Unison Replication

Configuration?! AH!



Unison Replication

Setup cron to run every x minutes

```
unison ssh://Server4//var/www /var/www \  
-rsync -times -prefer=newer \  
-auto -batch \  
-ignore "Path tmp/*" \  
> /tmp/unisonreplication.log 2>&1
```



Video

P3- Lamp Unison Drupal



Tungsten Replicator

Database replication



Tungsten Replicator

Tungsten Replicator is a high performance, open source, data replication engine for MySQL. It offers a set of features that surpass any open source replicator available today: global transaction IDs to support failover, flexible transaction filtering, extensible transaction metadata, sharding, multiple replication services per process...



**Always
Online**

Processing services command...

NAME	VALUE
-----	-----
appliedLastSeqno:	50844930
appliedLatency :	0.527
role :	master
serviceName :	fsdbmeshA
serviceType :	local
started :	true
state :	ONLINE

NAME	VALUE
-----	-----
appliedLastSeqno:	745199
appliedLatency :	0.431
role :	slave
serviceName :	fsdbmeshB
serviceType :	remote
started :	true
state :	ONLINE

NAME	VALUE
-----	-----
appliedLastSeqno:	16599
appliedLatency :	0.781
role :	slave
serviceName :	fsdbmeshC
serviceType :	remote
started :	true
state :	ONLINE

NAME	VALUE
-----	-----
appliedLastSeqno:	16619
appliedLatency :	4.678
role :	slave
serviceName :	fsdbmeshD
serviceType :	remote
started :	true
state :	ONLINE

Finished services command...



Tungsten Replicator

`sql_mode` on `settings.php`

(...)

```
"driver" => "mysql",
"database" => "jblogsrmr",
"username" => "jblogsrmr",
"password" => "zzzxzzzzxzxzzxzz",
"host" => "fsdbmesh-aaaa",
"port" => "$mysql_port",
"prefix" => "",
"init_commands" =>
array (
    "sql_mode" => "SET sql_mode=\\"REAL_AS_FLOAT,PIPES_AS_CONCAT,
IGNORE_SPACE,STRICT_TRANS_TABLES,STRICT_ALL_TABLES,NO_ZERO_IN_DATE,
NO_ZERO_DATE,ERROR_FOR_DIVISION_BY_ZERO,NO_AUTO_CREATE_USER\\" ,
),
```



Video

P4- Tungsten Install



Tungsten Replicator

AWS Performance tips:

Use arrays of Elastic Block Storage volumes

- Dramatically improves performance.
- Consider pre-warming the array
- Raid Arrays of 8 volumes gives better timings.
- Use ec2-consistent-snapshot for snapshots.

<http://alestic.com/2009/09/ec2-consistent-snapshot>

Use the largest instance size in each class

- Limits the potential impact of a noisy neighbor.
- m1.xlarge, m2.4xlarge



Video

P5- Installation end



"Do it yourself" Scenario.

Self-hosted + scalable site?



*Admiral Ackbar
Battle of Endor in Return of the Jedi*

or the smart choice:

Drupal™

Acquia™

With whatever you do...

Choose H.A. in...

the Cloud



Multi Region failover scenarios in the Cloud

Questions?

ricardo.amaro@acquia.com

twitter: [@ricardoamaro](https://twitter.com/ricardoamaro)



Multi Region failover scenarios in the Cloud

Locate this session on the
DrupalCon Munich website:

<http://munich2012.drupal.org/program>

Thank you!

ricardo.amaro@acquia.com

twitter: [@ricardoamaro](https://twitter.com/ricardoamaro)