



BARCELONA
DRUPALCON 2015

Drupal Extreme Scaling

Zequi Vázquez

DevOps Track

barcelona2015.drupal.org/sessions/drupal-extreme-scaling



Who's me?

- Zequi Vázquez
- DevOps & Backend
- PhD student
- Hacking & Security
- Rock'n'Roll (electric guitarist)
- Videogames
- Books





- 1 Introduction
- 2 The Project
- 3 Problems and Solutions
- 4 Demo
- 5 Conclusions



- 1 Introduction
- 2 The Project
- 3 Problems and Solutions
- 4 Demo
- 5 Conclusions



Drupal performance, scalability & availability


- Memcache
- Apc / Opcache
- Varnish
- Redundancy
- Choose your cache improvement module





Cloud computing and containers

- Elastic computing → resources grow on demand
- In theory, we achieve full scalability and availability. **False!**
- The hidden enemy: budget grows till infinity and beyond
- Containers (Docker) → Only for local environments?






- 1 Introduction
- 2 The Project
- 3 Problems and Solutions
- 4 Demo
- 5 Conclusions

Some base requirements


- Built in Drupal 7, as a multisite install
- At first, +30k sites. It needs to scale up to +100k sites
- Availability close to 99.999 %
- Of course, high performance at lowest possible cost
- Need to control site creation from external app
- Automated and not-disruptive deployments
- Ability to execute drush commands on all sites
- Migration from previous platform
- Three devs only - frontend, backend and devops






Seems scary, uh?

- Usual techs are not enough to tackle project
- We are “computer scientists”, aren’t we?
- God bless Open Source
- Stairway to ~~heaven~~ cloud
- Docker is more than local






BARCELONA
DRUPALCON 2015





Drupal over Nginx and PHP-fpm

- Nginx allow a real very flexible configuration
- PHP fpm increases performance significantly
- Mod Security with custom rules for Drupal






Manage sites on a multisite Drupal

- Node.js is asynchronous: execute something, and forget about it
- Use DynamoDB to store a list of deployed sites
- API to allow batch operations on sites list




Stateless Drupal container

- Memcache, MySQL, MongoDB as external services
- S3fs module → all files on S3
- Emails sent through PostMark
- Monitorization using NewRelic




The Cluster: Apache Mesos and Marathon

- Mesos abstracts the resources and creates a cluster
- Two masters orchestrated by Zookeeper, one or more workers
- Marathon allows to run Docker containers
- Chronos allows to run cron jobs
- Rest API to manage apps and containers



Mesos over AWS

- EC2 image for autoscaling from Mesos main worker
- AMI contains all 3 Docker images
- Script to sync amount of Marathon apps with workers number





Deployments: And automatization for all ...

- Lazy DevOps is best DevOps
- Ansible let us manage AWS in a simple (and powerful!) way
- Makefile + Dockerfile = Over 9000!
- Create and destroy environments with a single command



**CAN'T SOMEONE ELSE
JUST DO IT?**

Devil is on details

- AWS do backups, but it's worth to create your own recovery plan
- Security groups for EC2 → inbound and outbound firewall
- Remember to use protection → Http Auth for all Rest APIs
- Log centralization is a must → All hail Mesos!





- 1 Introduction
- 2 The Project
- 3 Problems and Solutions
- 4 Demo
- 5 Conclusions

Database(s): So you think 30k sites are too much

- 30k sites on a Drupal multisite means 30k databases
- MySQL → 1 folder per database, 1 file per table
- MongoDB preallocates files when a database is created
- Conclusion? Unmanageable



Divide and conquer!

- Identify each site using a unique hash
- Use that hash as prefix for tables on MySQL and MongoDB
- Group 500 sites per database, both MongoDB and MySQL
- MongoDB maximum database size is 3.952GB
- MySQL number of tables per database is 62k aprox.





The Way of the Request

- Ever seen a settings.php with steroids?
- Connect to DynamoDB
- Identify domain from the request (both drush and http)
- Calculate domain unique hash
- Ask DynamoDB for the database where the site lives
- Make sure hash is used as prefix
- Store database reference on Memcache






In soviet Russia, Nginx configures you

- Strange timeout when creating a new site from NodeJS app
- Happened when Node.JS app launched a site creation
- Fire and forget. When site was created, app is notified
- Further investigation → timeout produced on Haproxy



Another brick(s) in the wall

- Unstability on autoscaling group
- MongoDB MMS instances out of space
- PHP-fpm unstability → slowlog
- Thousand S3 buckets? God bless S3fs
- Remote & massive Drush
- Migration: 30k sites on few days
- AWS RDS speed: the bigger, the better





- 1 Introduction
- 2 The Project
- 3 Problems and Solutions
- 4 Demo**
- 5 Conclusions



BARCELONA
DRUPALCON 2015

Demo time!




- 1 Introduction
- 2 The Project
- 3 Problems and Solutions
- 4 Demo
- 5 Conclusions

Conclusions: Spoiler alert!

- Project went live on April 2015. Absolute success!
- Learned a lot about new technologies
- Cloud + Containers = Profit
- Drupal flexibility is amazing



BARCELONA
DRUPALCON 2015



Thank you!

What did you think? Please, evaluate this session

 @RabbitLair

 ezequielvazq[at]gmail[dot]com