



Cloud Hosting Practices

Lessons DuraSpace has learned

Bill Branan

Open Repositories 2013



Purpose

- Share what we've learned about applications development and hosting in a cloud environment
- Start a conversation about cloud best practices
- Enable more and better use of the cloud



Cloud History



DURACLOUD™

Established 2009



DURASPACE™

Infrastructure

Established 2010



Fedora Commons™

Demo Site

Established 2009



DSPACE

Demo Site

Established 2010



DSPACE DIRECT

Established 2012



The Task

- Your goal: Hosted app
- Constraints:
 - No local IT support
 - Up and running quickly
 - Scale with demand



Questions

- Which cloud providers to use?
- Which cloud services to use?
- How to put the pieces together?
- How to handle maintenance?
- How to ensure availability?
- How to scale?
- ...



Commercial Cloud

- Available
 - No datacenter required
- Flexible
 - No up front costs, no contracts
- Lots of options
 - AWS, Rackspace, Windows Azure, GoGrid, HP, ...

Capabilities of clouds vary greatly



Amazon Web Services

- Clear market leader
 - 6x the market share in IaaS as its next closest competitor
- Global
 - 9 regions in 6 countries
- Scale
- History of price reductions
- Pace of Growth



EC2

- Elastic Compute Cloud
 - Manage server instances
 - Amazon Machine Image (AMI)
 - Use EBS-based instances
 - Use reserved instances
 - Use spot instances when possible
 - Use elastic IPs



EBS



- Elastic Block Store
 - Block-level storage volumes
 - Think “hard drive”
 - 1 GB - 1 TB
 - Best file system storage option for EC2
 - Fast and easy to use
 - Can be moved between instances
 - Not as durable as S3
 - Make regular snapshots to S3



Storage



- S3 - Simple Storage Service
 - Highly reliable file storage
 - Files available for immediate download
 - RRS option
- Glacier
 - Cheaper than S3
 - Same durability as S3
 - Takes longer, costs more to get data out




Design

- Stateless
 - Applications
 - AMIs
- Expect failures
 - Unexpected restarts
- Plan to scale
 - Horizontal scale: More instances
 - Vertical scale: Bigger instances




Pulling it all together

- Building AMIs
- CM Automation
 - Consistency
 - Repeatability
 - Control
- Puppet
 - Define + deploy
 - Enforce



Testing and Upgrades

- Start from latest EBS snapshot
- Change local /etc/hosts file
- To deploy: Assign elastic IP




Monitoring

- Lots of options
 - monitor.us
 - Up or Down
 - AWS CloudWatch
 - Instance monitoring
 - New Relic
 - Server and application
 - Newvem
 - Costs, activity, utilization



Accounts

- Use accounts to separate environments:
 - Development, Test, Production
 - Applications
 - Users
- Consolidated billing
 - Can complicate reservations



We are learning

- Multiple availability zones
- Elastic load balancer
- Autoscaling
 - Scale up and down
 - Based on metrics (CloudWatch)
 - Even for 1 instance
- Use IAM (ID and Access Mgmt)



Thanks!

Questions?

bbranan@duraspace.org

