

# To the Cloud

Working with Linux in AWS

# What is the cloud?

Internet based computing in which large groups of remote servers are networked so as to allow sharing of data-processing tasks, centralized data storage, and online access to computer services or resources.

The cloud is not unlimited and it's not a big mystery. Any provider with a datacenter can claim to be a cloud provider.

# The term “cloud” is abused

Every device sold is cloud enabled or optimized:

hard drives

routers

cameras

picture frames

# **AWS is known as an IaaS provider**

EC2: Elastic Compute Cloud (Xen based PV or HVM servers)

EBS: Elastic Block Storage

S3: Simple Storage Service (object store)

ELB: Elastic Load Balancer

VPC: Virtual Private Cloud

Cloudfront: Content Distribution Network

RDS: Relational Database Service (mysql, postgres)

SES: Simple Email Service (email smarthost relay)

# **Every AWS resource has an id and is taggable**

instance id: i-00197068

volume id: vol-0bb81b11

snapshot id: snap-00297207

security group: sg-103c247a

elastic ip addresses: eipalloc-4ee77921

S3 is an exception: objects are URL based

# Anatomy of an EC2 Instance

based on an AMI: Amazon Machine Image

snapshot of a bootable image

contains all files from the OS

contains metadata about the instance that can be queried via an API

storage can be ephemeral or EBS

network interface(s): elastic network interface (eni)

has one or more security groups applied

plan for and expect instances to go away when it is not convenient

instances run on real hardware that will fail

leverage a configuration management tool like Puppet, Chef, Ansible

# Launching an instance

An instance is launched in an availability zone within a region.

Global regions: us-east-1 (VA), us-west-1 (CA), us-west-2 (OR), eu-west-1 (Ireland), eu-central-1 (Frankfurt), ap-southeast-1 (Singapore), ap-northeast-1 (Tokyo), ap-southeast-2 (Sydney), and sa-east-1 (São Paulo)

Each region has two or more availability zones. Availability zones are discrete datacenters within a region. Availability zones within a region are on separate flood planes, have independent and redundant power, and are connected via low latency links.

# Load Balancing Demonstration

We will launch two instances in separate availability zones behind an elastic load balancer. We will set up a simple apache configuration and demonstrate that traffic is sent round robin to each instance. We will cause apache to fail and see the load balancer mark the instance unhealthy and continue to serve traffic to the healthy instance.



# EBS Demonstration

We will attach multiple EBS volumes to a running instance and create a software RAID volume using mdadm.

We will take a snapshot of a volume with data and create a new volume from that snapshot.

# Contact David

Email: [dcolon@dcolon.org](mailto:dcolon@dcolon.org)

Blog: <http://tech.dcolon.org>

Shoprunner is hiring

<https://www.shoprunner.com/careers/>