

# Amazon AWS in a glance

Prepared by Basil Kurian

Technical Specialist, [Dinoct.com](http://Dinoct.com)



# Primary components

- Elastic Compute Cloud (EC2)
- Simple Storage Solution (S3)
- Elastic Block Storage (EBS)
- Amazon Machine Image (AMI)
- Elastic Load Balancer (ELB)
- Elastic IP
- Amazon CloudWatch
  
- See full list at <http://aws.amazon.com/products/>



# Elastic Compute Cloud

- Central part of Amazon Web Service
- Original computing entity analogous to physical servers
- Users can create, launch, and terminate server instances as needed
- Launch instances with a variety of operating systems and hardware resources.
- Facilitate scalable deployment and pay only for the resources consumed.
- Hence the name Elastic
- Instances are launched with a choice of two types of storage for its boot device.
  - 1) Instances with EBS boot volume
  - 2) Instances without EBS boot volume

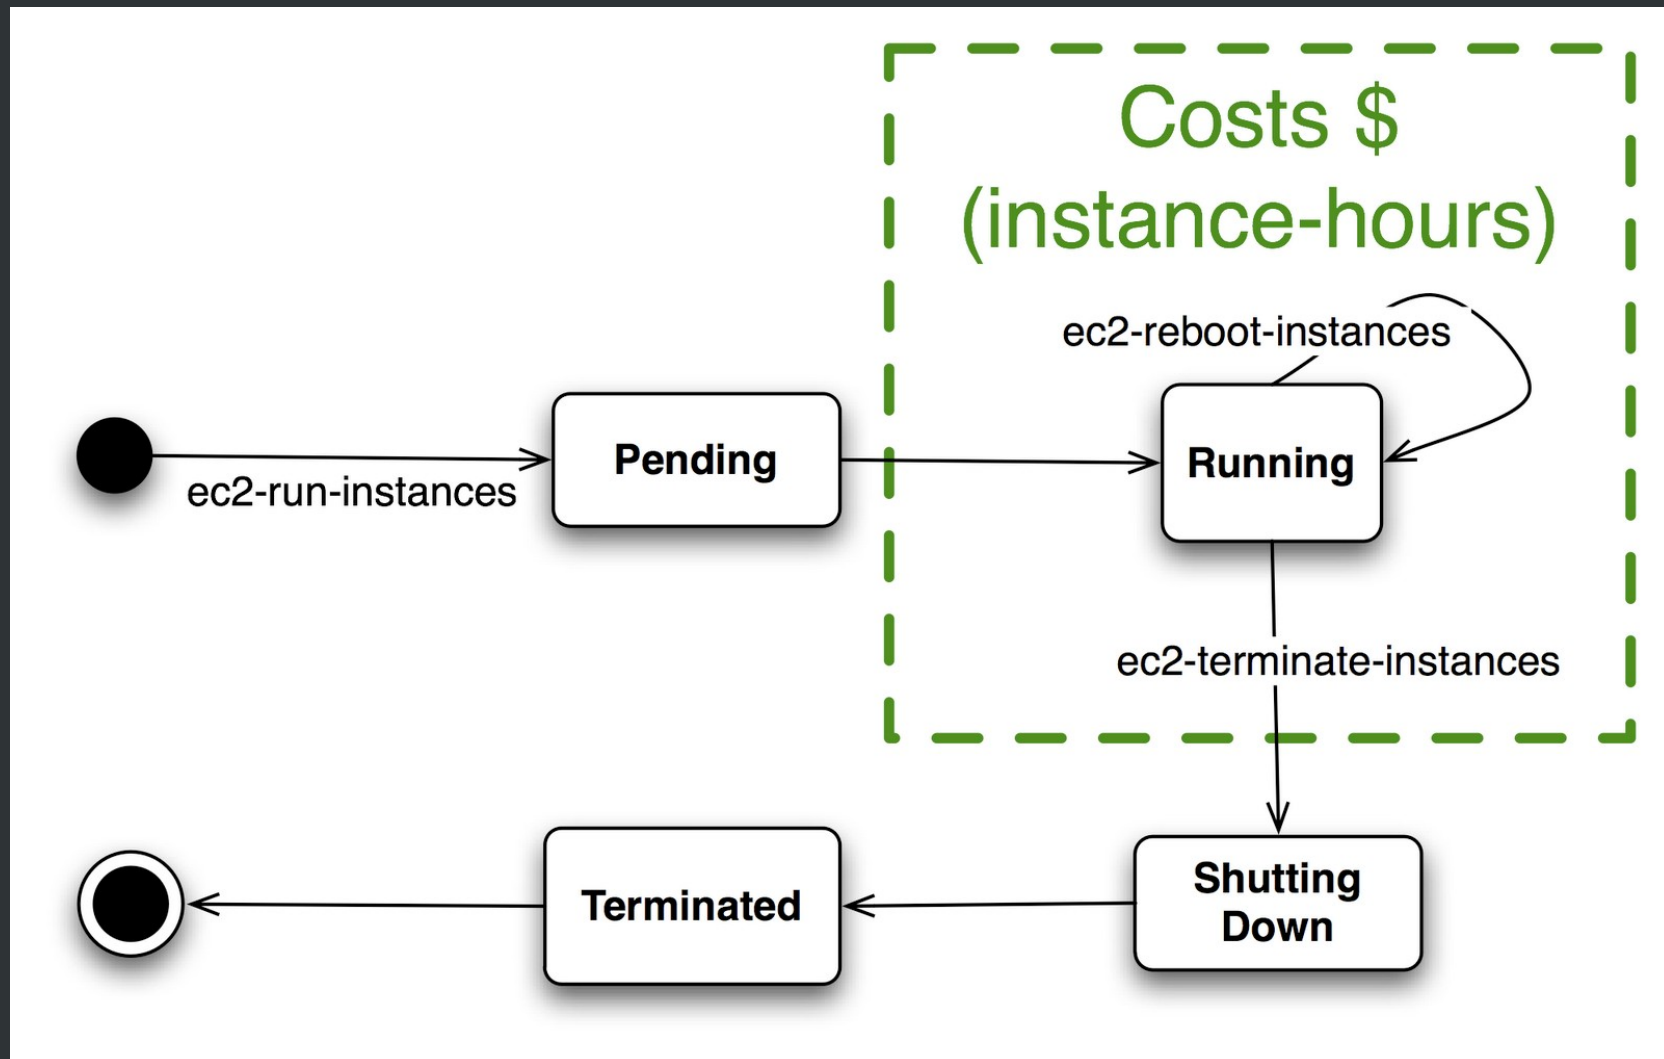


# EC2 instances without EBS boot volume

- The boot device will be on instance volume.
- Instance volume are temporary storage
- Instance volume survive rebooting an EC2 instance
- When the instance is terminated (by an API call, or due to a failure) this store is lost.
- EC2 instances without EBS can't be stopped



# Lifecycle of an EC2 instance without EBS (with instance store)



• Courtesy : <http://shlomoswidler.com>

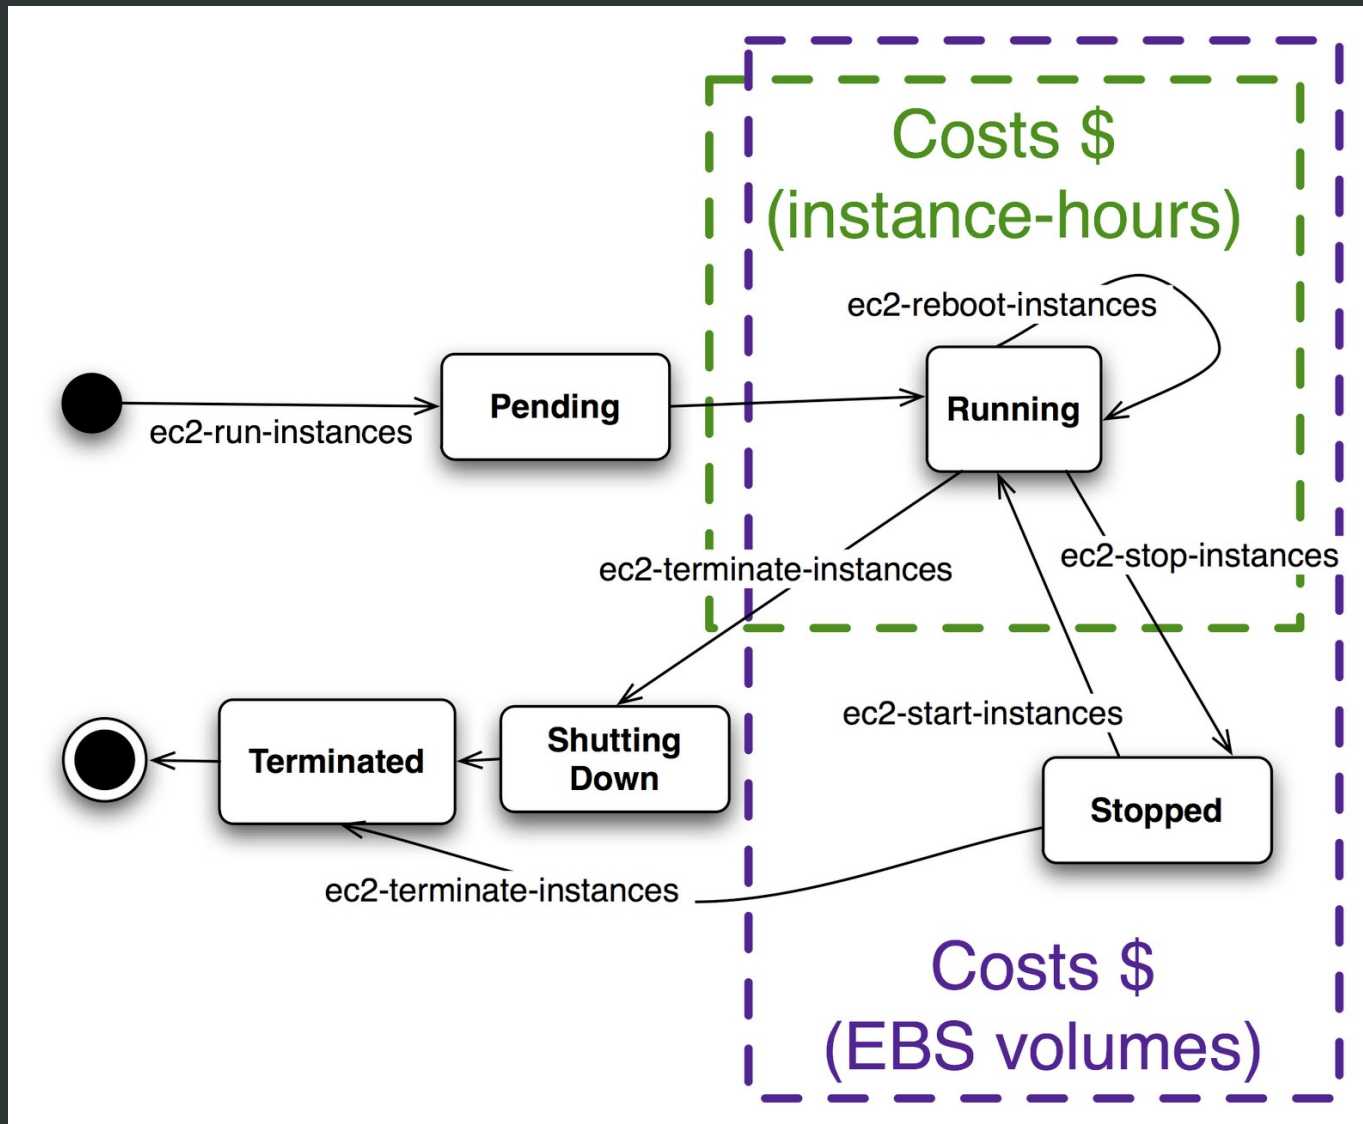


## EC2 instances with EBS boot volume

- EBS volumes are permanent and can act much like hard drives on a real server
- Appear as block devices to the operating system
- Storage volumes from 1 GB to 1 TB size can be created
- A volume can only be attached to one instance at a time, but many volumes can be attached to a single instance
- One EBS can be attached to any instance in the same Availability Zone
- Storage volume is automatically replicated within the same Availability Zone.



# Lifecycle of an EC2 instance with EBS



# EBS snapshots

- Creating a snapshot causes a copy of the data in the volume to be written to S3
- Successive snapshots are incremental
- We can restore a snapshot by creating a new volume from it
- Snapshots can't be manipulated externally (we can't move it)
- Only 500 snapshot allowed per account. Not per machine
- On deleting a snapshot, the blocks that are only used by that snapshot are deleted.



# EC2 Availability Zones

- Amazon EC2 has several regions and multiple availability zones in each regions
- A region can be thought of as a specific area of the world.
- An availability zone can be thought of roughly as a data center
- No single failure scenario will affect two availability zones
- The current regions are:
  - 1) ap-southeast-1
  - 2) eu-west-1
  - 3) us-east-1
  - 4) us-west-1



## EC2 Availability Zones (continued)

- Availability zones in us-east-1 are

us-east-1a      us-east-1b

us-east-1c      us-east-1d

- Similarly there are several zones in other regions also.
- Note : EC2 availability zone names in different accounts do not match to the same underlying physical infrastructure



# Simple Storage Solution (S3)

- Simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web.
- Write, read, and delete objects containing from 1 byte to 5 terabytes
- Designed to provide 99.999999999% durability and 99.99% availability of objects over a given year.
- A bucket can be stored in one of several regions.
- Buckets and objects can be created, listed, and retrieved using either a REST-style HTTP interface or a SOAP interface.



# Amazon Machine Image (AMI)

- Special type of virtual appliance which is used to start an EC2 instance
- It serves as the basic unit of deployment for services delivered using EC2
- Main component is a read-only filesystem image which includes an operating system and any additional software required to deliver a service or a portion of it.
- The AMI filesystem is compressed, encrypted, signed, split into a series of 10MB chunks and uploaded into Amazon S3 for storage.
- An XML file stores information about the AMI, including name, version, architecture, default kernel id, decryption key and digests for all of the filesystem chunks.
- An AMI does not include a kernel image, only a pointer to the default kernel id, which can be chosen from an approved list of safe kernels maintained by Amazon and its partners (For example Canonical or Redhat)



# Amazon Elastic Load Balancing (ELB)

- Distributes incoming application traffic across multiple Amazon EC2 instances
- Enables us to achieve even greater fault tolerance in our applications
- Detects unhealthy instances within a pool and automatically reroutes traffic to healthy instances until the unhealthy instances have been restored.
- Can distribute incoming traffic across your Amazon EC2 instances in a single Availability Zone or multiple Availability Zones
- Supports the ability to stick user sessions to specific EC2 instances
- Supports SSL termination at the Load Balancer
- With the Auto Scaling feature, ELB automatically adds the requisite amount of EC2 instances to our Auto Scaling Group when the response time increases.
- IP address can't be assigned to an ELB.



# Amazon Elastic IP

- Similar to static IP address in traditional data centers.
- User can map an Elastic IP Address to any EC2 instance using API
- Elastic IP Address belongs to the account and not to a virtual machine instance
- We can't assign Elastic IP to other entities like ELB
- By default, each Amazon EC2 customer has a limit of 5 Elastic IP addresses.



# Amazon CloudWatch

- A web service that provides monitoring for AWS cloud resources.
- Provides customers with visibility into resource utilization, operational performance, and overall demand patterns including metrics such as CPU utilization, disk reads and writes, and network traffic.
- Collected data can be accessed using the AWS Management Console, web service APIs or Command Line Tools.



Thanks

Basil Kurian  
basil@basilkurian.co.cc

