



Image by evertonpestana from Pixabay

INTRODUCTION TO CLOUD ☁ FORENSICS: STORAGE & ACQUISITION

Simson L. Garfinkel*
US Census Bureau
Tuesday, December 10, 2019

DISCLAIMER:
The views in this presentation are those of the author, and not those of the US Census Bureau.

Disclaimer & Level Setting
Cloud forensics — Definitions
Amazon Web Service forensic
targets

Instance-based acquisition:

- RAM

- Instance-attached drives

- Elastic Block Service (EBS)

Service-based acquisition:

- Simple Storage Service (S3)

- Cloud Watch

Backup Slides:

- Creating an instance

- EC2 Command Line Tools

- AWS EBS

- AWS CloudTrail

- AWS EFS

- Running bulk_extractor in AWS

Outline of this briefing

Disclaimer

This presentation is based on:

- Working with Amazon Web Services (AWS) since August 2006 (S3 and EC2)
- Cloud Forensics course at George Mason University
 - http://bit.ly/Cloud_Forensics_2018
- Work at US Census Bureau in AWS GovCloud.
 - *Elastic Compute Cloud (EC2)*
 - *Simple Storage Service (S3)*
 - *Elastic Map Reduce (EMR)*
 - *Largest Cluster size: 50 r5d.24xlarge nodes = 4,800 CPU cores & 38.4 TiB RAM*

This presentation focuses on Amazon Web Services

- **Microsoft Azure** and **Google Cloud Platform** have many competitive services.
- IBM and Oracle also have significant offerings. Alibaba Cloud has great prices!
- US Census Bureau has adopted AWS GovCloud.
- I use Dreamhost for my personal stuff (limited to compute & storage).

Level Setting — Show of Hands

Who has worked with AWS EC2 instances?

Collected RAM from an EC2 instance?

Imaged EBS volumes?

Collected data from Amazon S3?


Used Amazon's internal APIs to perform undocumented functions?





Cloud Forensics: Definitions

We are all using the cloud



<https://www.pexels.com/photo/interior-of-office-building-325229/>



**Apps on
end-user
devices**



The Internet

<https://www.pexels.com/photo/facebook-like-mi-mobile-325053/>

Cloud forensics: *It's where the data are!*

"Everything is moving to the cloud."

That means digital evidence is moving to the cloud.

Digital evidence on end-user-devices is increasingly encrypted.



This works less and less.



All of the data are in the cloud.

Cloud forensics: *It's where the research is.*

The "cloud" is just a bunch of data centers.

Forensics tools for data centers are poorly developed.

- Most “cloud forensics” is really traditional incident response running in AWS.

This is an excellent opportunity!

- Virtualization
- Software Defined Networks (SDN)
- Big data analysis techniques.



Cloud computing is more than just “big, remote data centers.”

NIST Special Publication 800-145: The NIST Definition of Cloud Computing

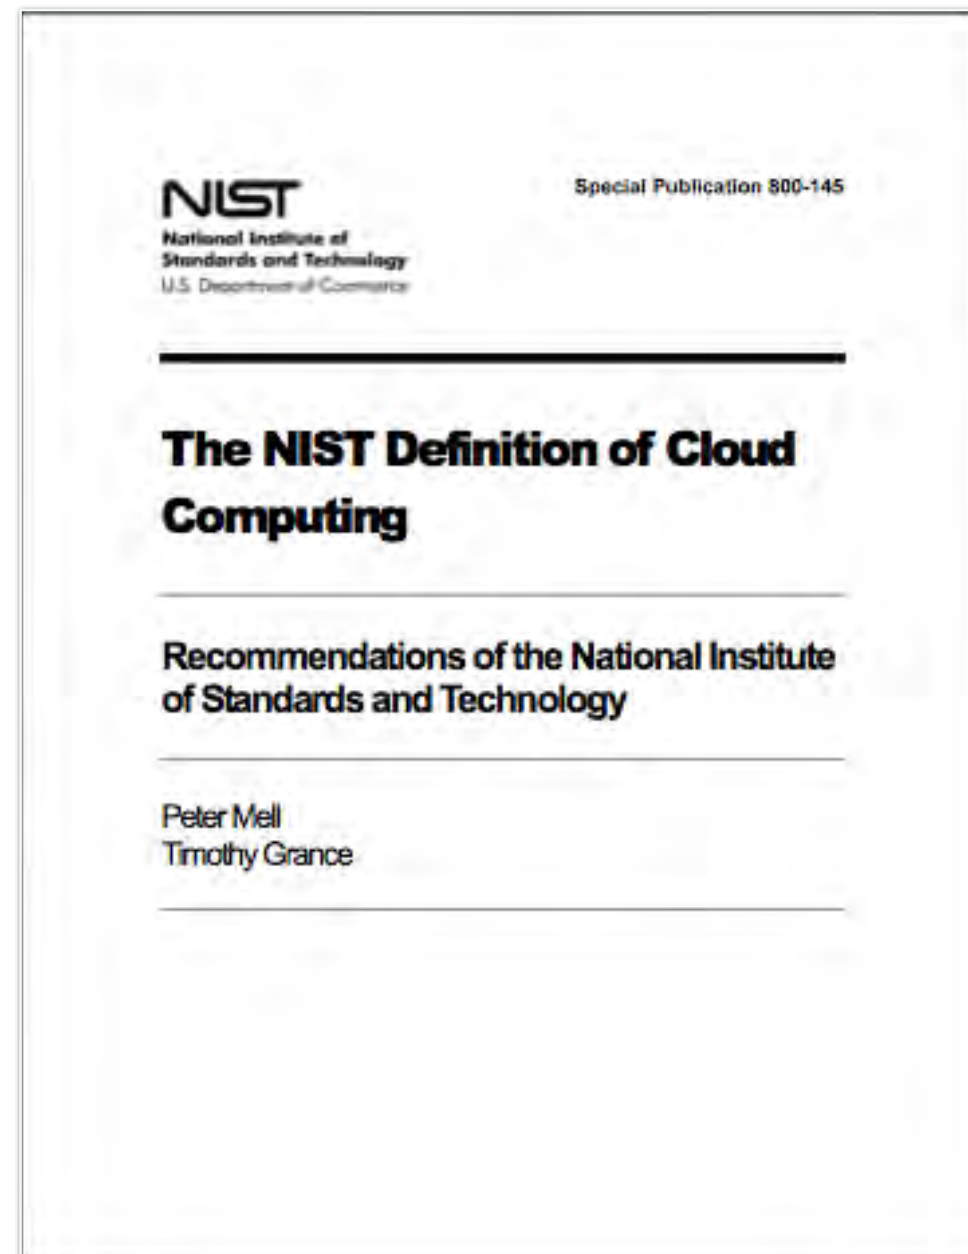
<https://csrc.nist.gov/publications/detail/sp/800-145/final>

Essential Characteristics:

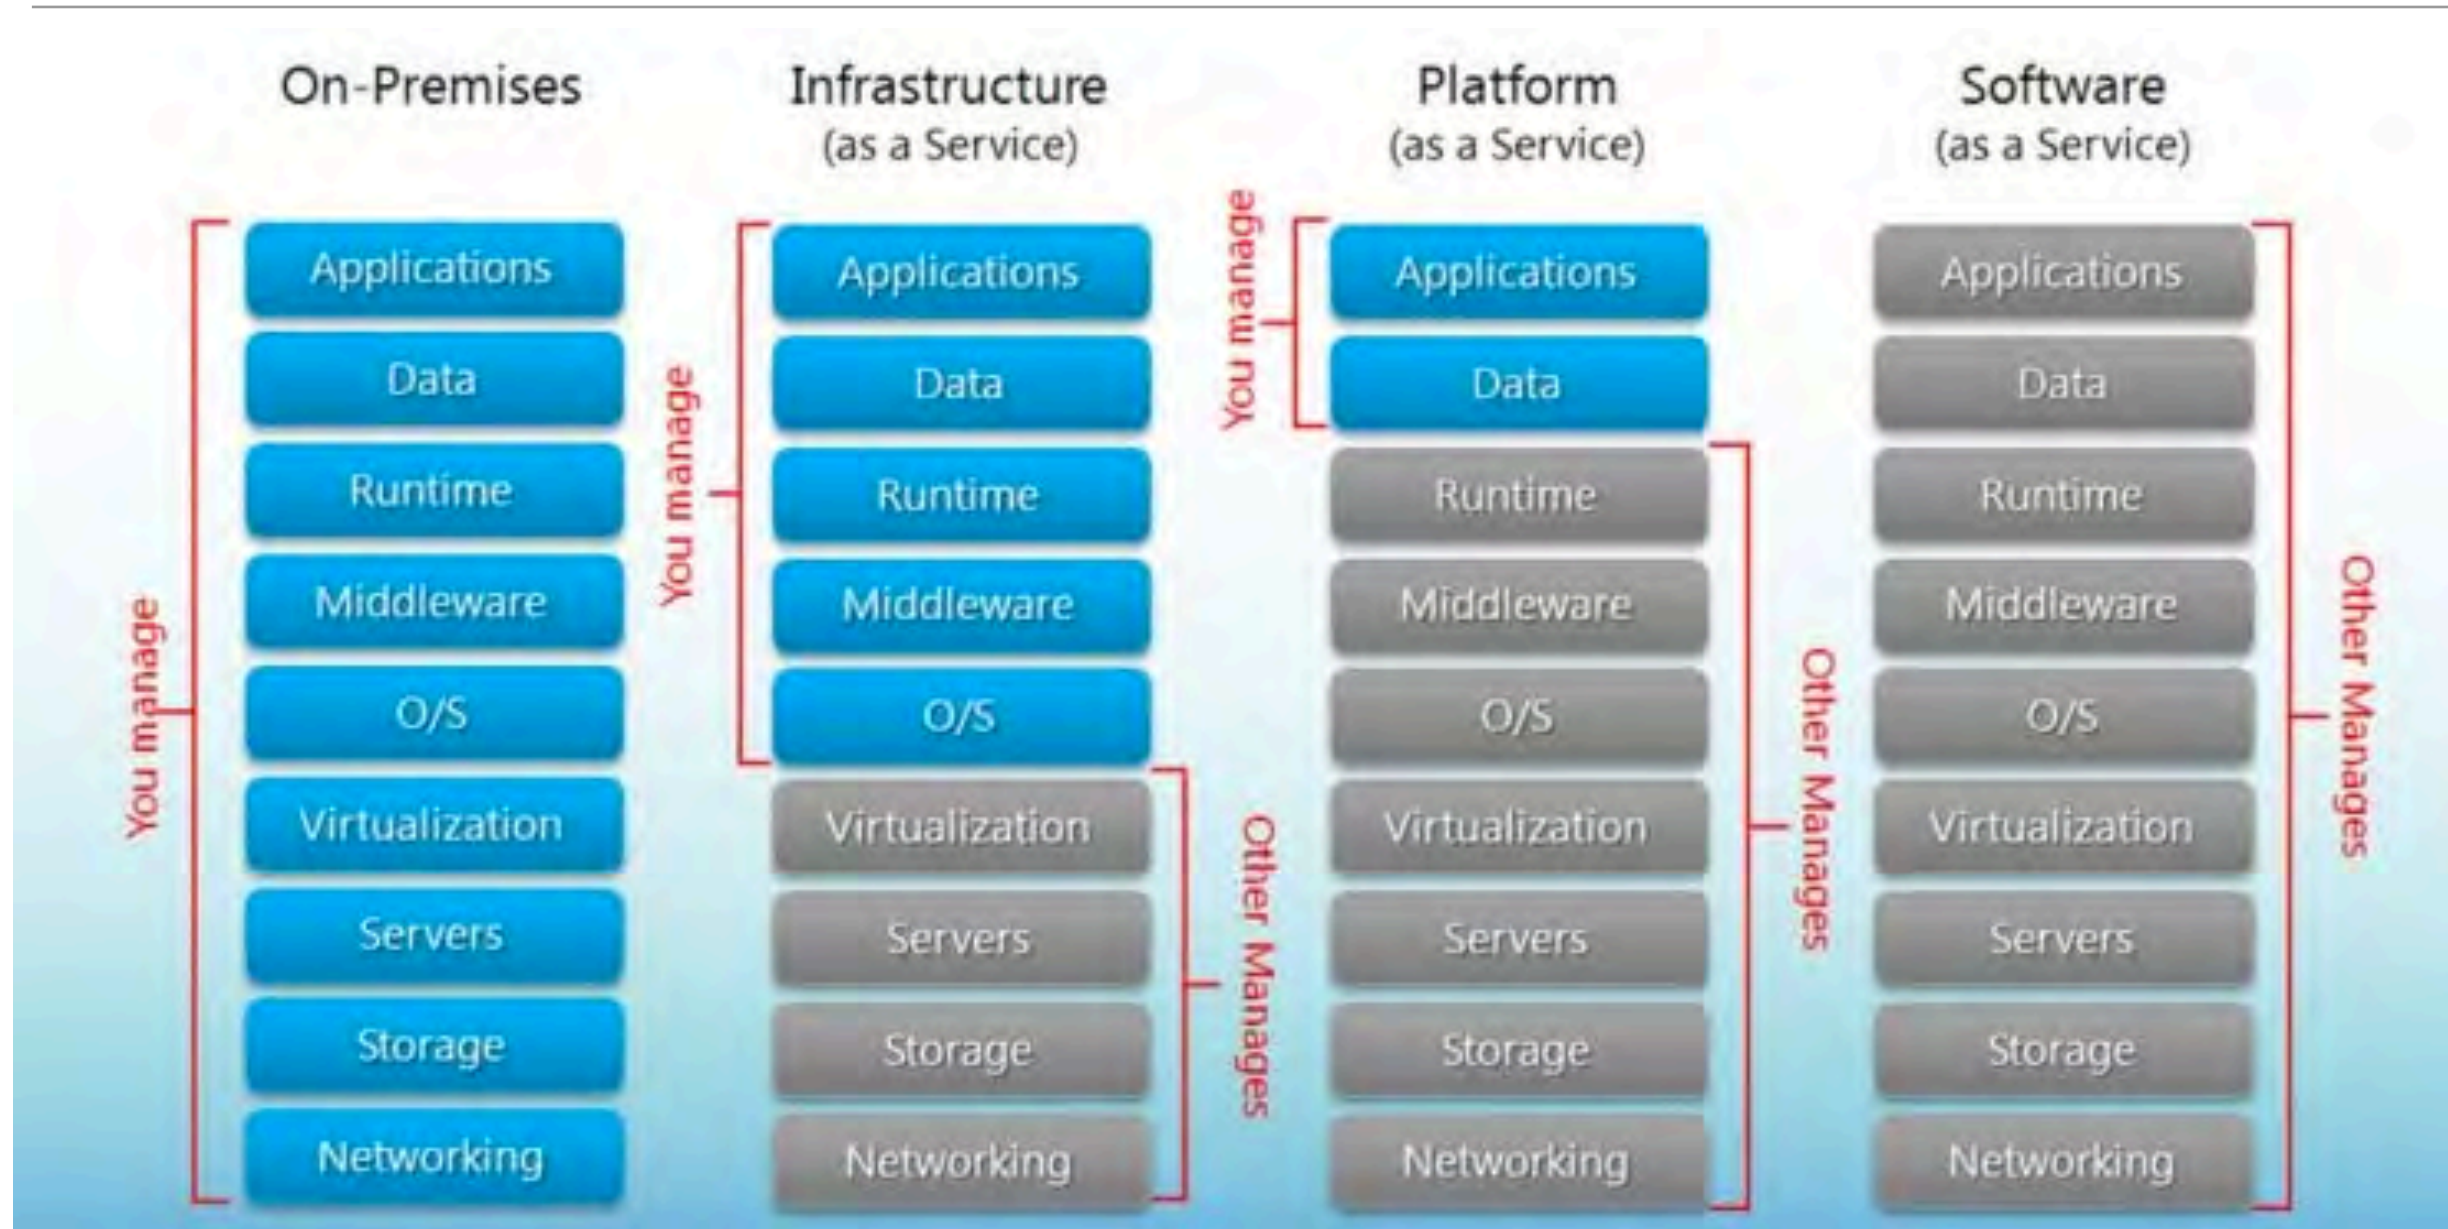
- On-demand self-service
- Broad network access
- Resource pooling
- Rapid elasticity
- Measures service

Deployment Models:

- Private cloud
- Community cloud
- Public cloud
- Hybrid cloud



Cloud hosting models



<https://www.hostingadvice.com/how-to/iaas-vs-paas-vs-saas/>

Different kinds of data are at each layer

Program Layer — code that runs in the cloud to manipulate the data

- Desktop applications (cloud-based desktops) (e.g. Amazon WorkSpaces & AppStream)
- Custom applications written in python, Scala, C++, Go, Java, etc.
- Back-ends for websites and mobile apps.

Software Infrastructure Layer — where the code runs

- YARN, Hadoop, MapReduce, Spark, etc.
- Databases: HBase, RDS

Operating System Layer — what you log into

- Linux (Centos), Windows

Virtualization Layer — the runtime environment

- “Bare Iron” or Xen

Hardware Layer — the physical hardware on which the VMs run

- Intel or AMD systems; GPUs

Other services common in cloud computing environments

“Functions as a Service” — Serverless Computing

- Isolates business logic from the problem of running servers.
- Amazon Lambda; Google Cloud Functions; Azure Functions

Hadoop & Apache Spark — Big Data Computing

- Designed for processing data larger than the largest server.
- Amazon Elastic Map Reduce — automatically scales cluster with workload

Other big data services:

- Amazon Athena — “serverless interactive query service.”
— *Runs queries on data stored in Amazon S3.*
- Amazon Redshift — Cloud data warehouse for structured and semi-structured data.

AWS Architecture

EC2

Non-EC2 Services



AWS Forensic Targets

Amazon Web Services — Brief History



July 5, 1994 — Amazon.com was founded by Jeff Bezos

- (Originally named “Cadabra”)
- Renamed “Amazon” in 1995 with goal of being the “biggest” store in the world.
- First book ordered in 1995, *Fluid Concepts and Creative Analogies*.

By 1998, more than 100 computers processed data for every rendered page.

- Authentication; shopping cart; search results; recommendations; feedback; ...
- Amazon made organizing thousands of computers an institutional priority.

In 2006, Amazon started making its systems available as a commodity

- March: Simple Storage Service (S3) — unlimited storage
- July: Simple Queue Service (SQS) — Reliable messages up to 256KB in size.
- August:
 - Elastic Compute Cloud (EC2) — virtual machines
 - Elastic Block Store (EBS) — disks for virtual machines



Amazon EC2

CLOUD  FORENSICS

14

AWS Global Infrastructure:

22 Regions; 69 Availability Zones

Location matters:


- Speed of light: 300,000 Km/sec
- Distance to Seattle: $\approx 5,000$ Km
- Minimum time to Seattle: $= 1.6$ msec

- Distance to Reston: ≈ 30 Km
- Minimum time to Reston: $30 \text{ Km} \div 300,000 \text{ Km/sec} = 99\mu\text{sec}$




<https://aws.amazon.com/about-aws/global-infrastructure/>

AWS is divided into “regions” and “availability zones”



Code	Name
us-east-1	US East (N. Virginia)
us-east-2	US East (Ohio)
us-west-1	US West (N. California)
us-west-2	US West (Oregon)



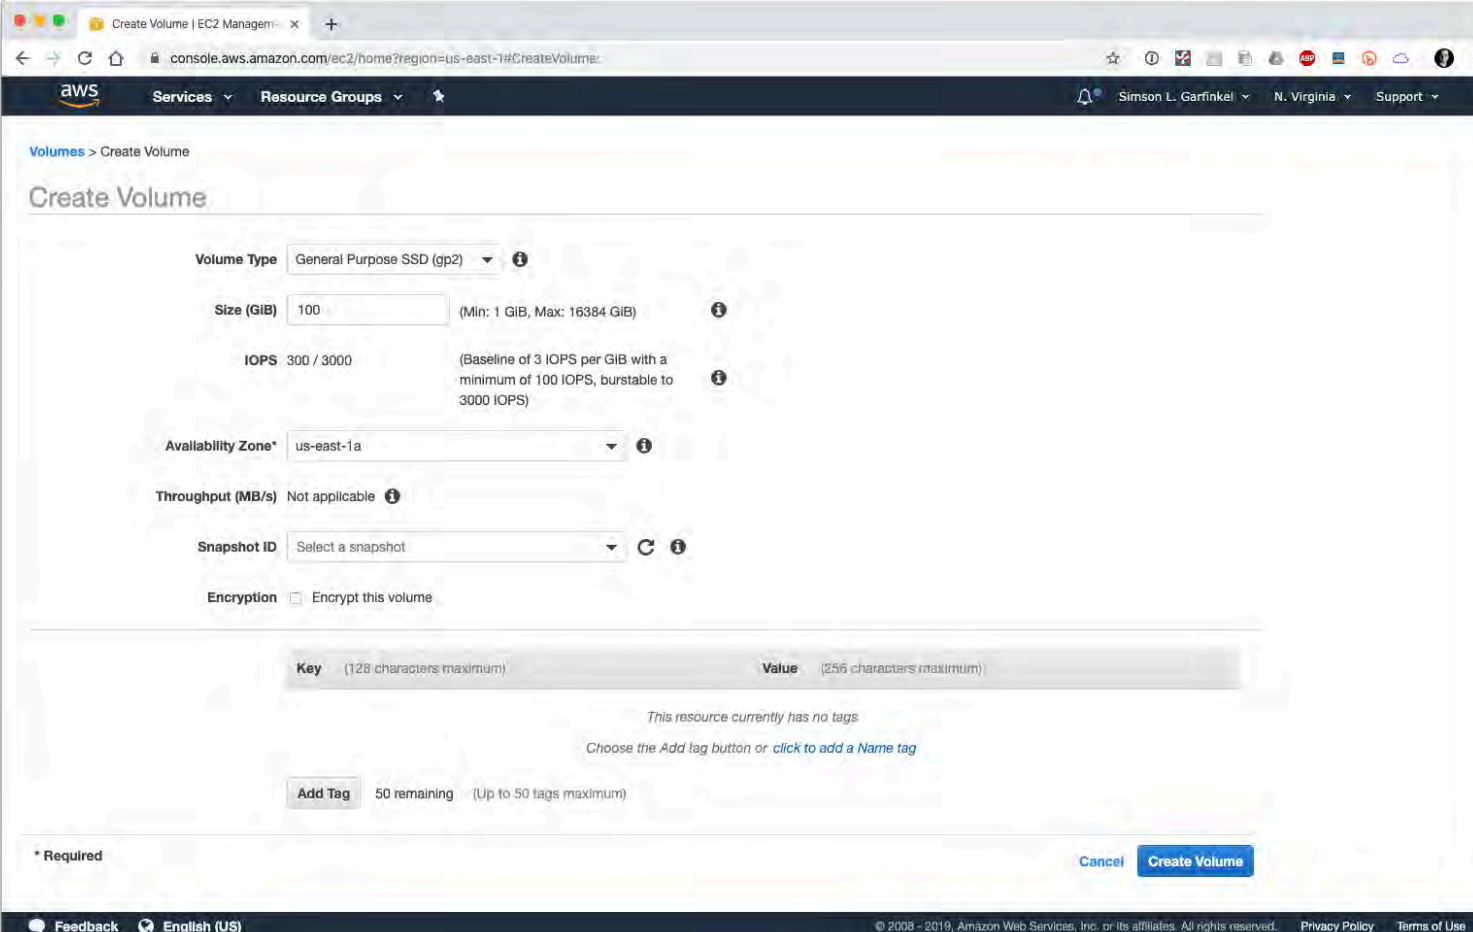
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

Manage AWS with the Graphical User Interface (GUI) or the Command Line Interface (CLI)

CLI:

```
$ aws ec2 create-volume --size 10 --region $aws_region --availability-zone us-east-1b {  
  "AvailabilityZone": "us-east-1b",  
  "Encrypted": false,  
  "VolumeType": "standard",  
  "VolumeId": "vol-46cdb6a5",  
  "State": "creating",  
  "SnapshotId": "",  
  "CreateTime": "2015-12-05T19:01:38.548Z",  
  "Size": 10  
}
```

GUI:



The screenshot shows the AWS Management Console 'Create Volume' page. The page is titled 'Create Volume' and is part of the 'Volumes' section. It features several input fields and a 'Create Volume' button at the bottom right. The fields are:

- Volume Type:** General Purpose SSD (gp2)
- Size (GiB):** 100 (Min: 1 GiB, Max: 16384 GiB)
- IOPS:** 300 / 3000 (Baseline of 3 IOPS per GiB with a minimum of 100 IOPS, burstable to 3000 IOPS)
- Availability Zone:** us-east-1a
- Throughput (MB/s):** Not applicable
- Snapshot ID:** Select a snapshot
- Encryption:** ☐ Encrypt this volume
- Tags:** A table with 'Key' and 'Value' columns. Below the table, it says 'This resource currently has no tags' and 'Choose the Add tag button or click to add a Name tag'. There is an 'Add Tag' button and a note '50 remaining (Up to 50 tags maximum)'.

At the bottom, there is a 'Create Volume' button and a 'Cancel' button. The footer of the console shows 'Feedback', 'English (US)', and copyright information for Amazon Web Services, Inc.

Amazon EC2 — Elastic Compute Cloud

Virtual machines in the cloud



EC2 is based on “Instances”

- Horizontal Scaling — Create many VMs.
- Vertical Scaling — Create small and large VMs (cores, RAM, networking)
- Geographical Diversity — Create in different locations (“availability zones”)

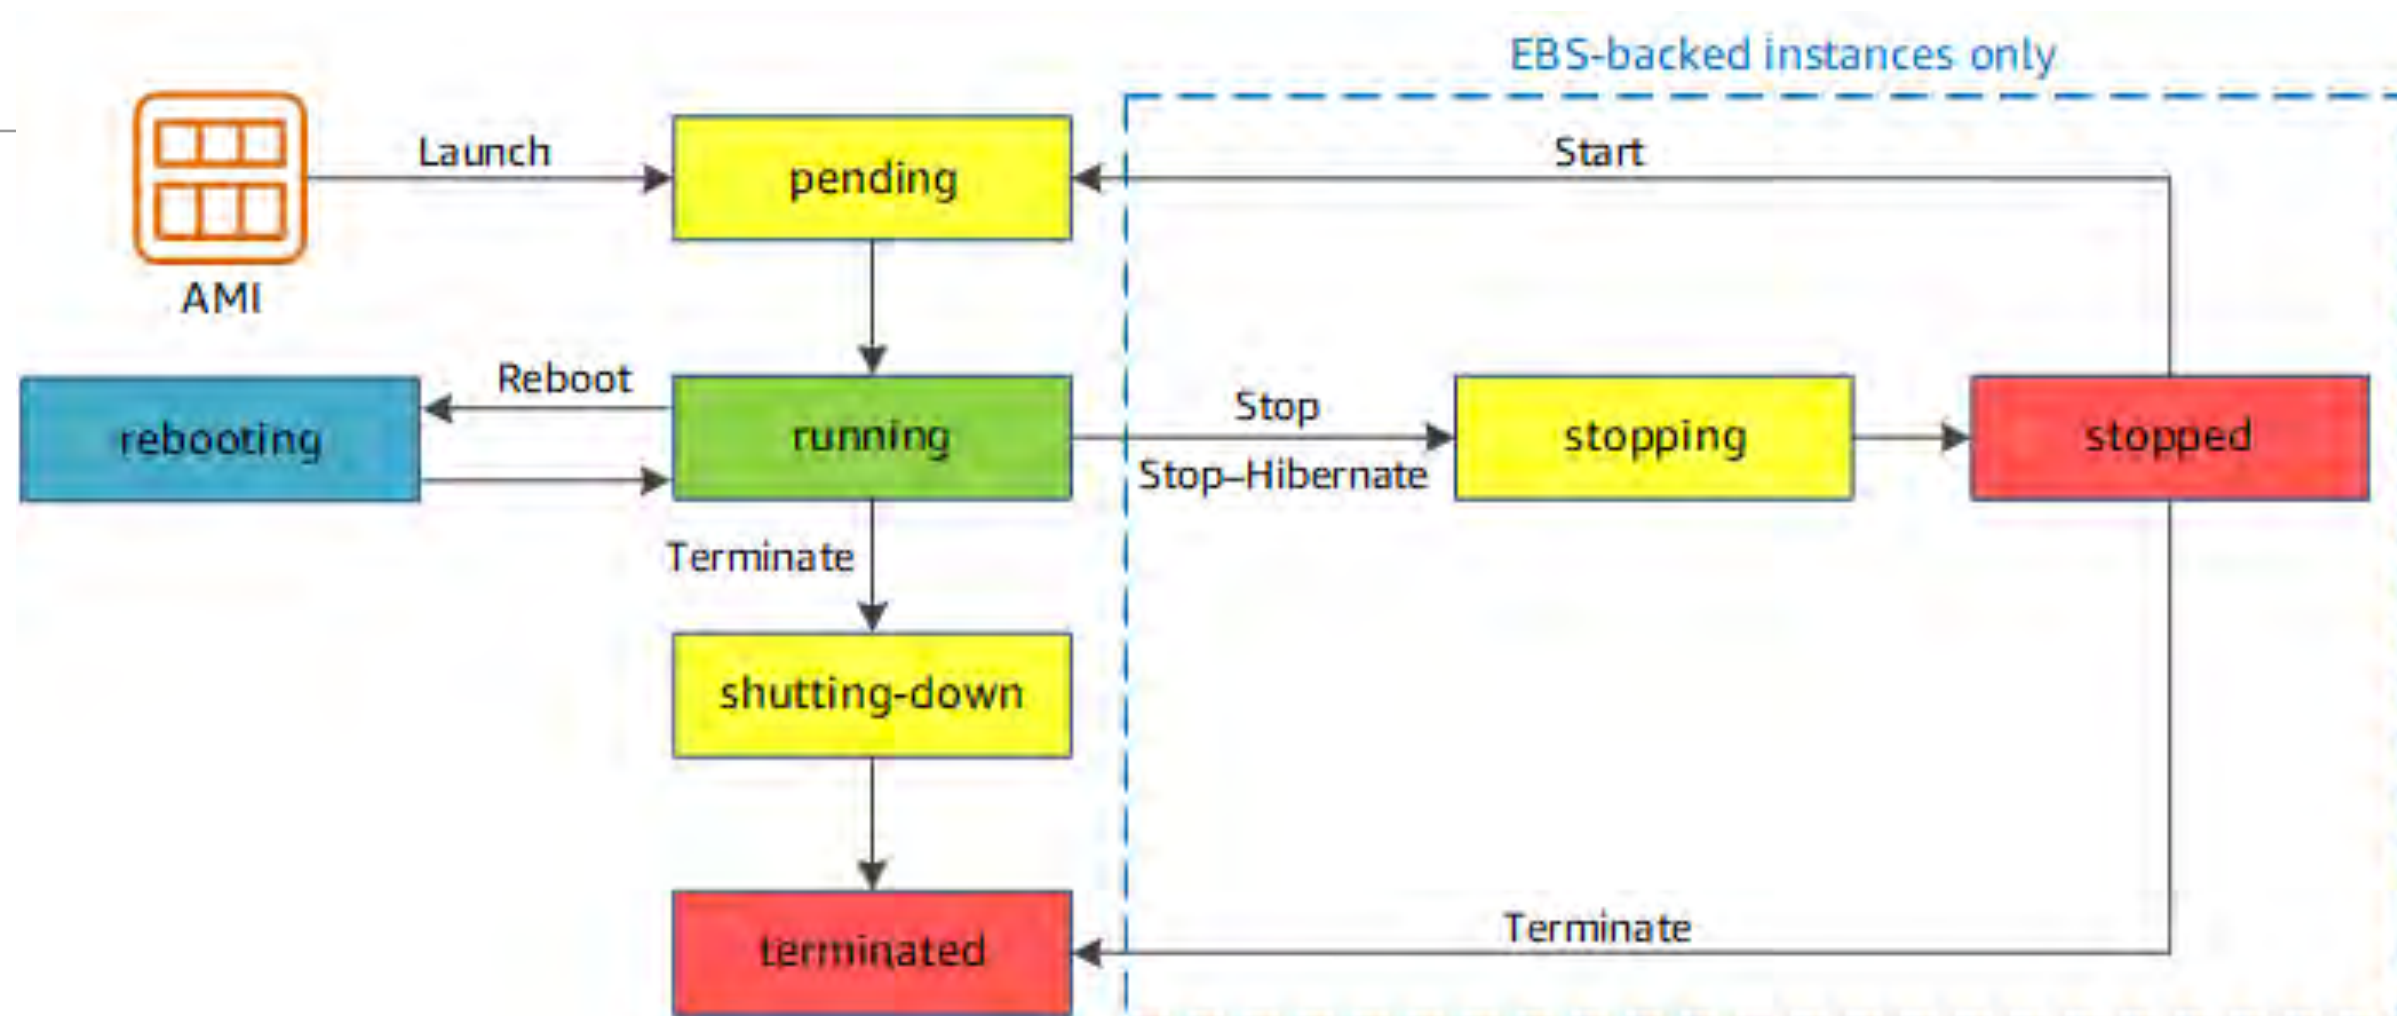
Each instance has:

- Unique Instance ID (e.g. i-04f679de246dc9c10)
- AMI — Amazon Machine Image — the initial “boot volume”
- Network interface(s) and firewall
- Instance Type (e.g. “t2.micro”) with specific CPU and RAM (1 vCPU; 1.0 GiB)
- Security Groups (what is can and cannot do)
- Key pair (used for accessing)

Instances optionally have:

- Virtual drives — Elastic Block Store; can survive shut-down.
- Attached physical drives — in the box; lost when VM terminates

EC2 Instance life cycle:



<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-lifecycle.html>

All instances boot from an AMI (you can upload your own.)

You specify if the EBS volume is kept or lost on termination.

You pay for:

- Instances that are running*
- EBS-backed storage
- Bandwidth from EC2 → Rest of Internet

EC2 Instance control panel:

Launch InstanceConnectActions

Filter by tags and attributes or search by keyword1 to 4 of 4

	Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public D
<input type="checkbox"/>	Persistent EC2	i-5c306beb	t2.micro	us-east-1b	stopped		None	
<input type="checkbox"/>	TLSA Tester	i-eba98616	t2.micro	us-east-1b	stopped		None	
<input type="checkbox"/>	Reminance ...	i-9a48aa2c	t2.micro	us-east-1b	stopped		None	
<input type="checkbox"/>	Quicken	i-8e0b7f64	t2.micro	us-east-1b	running	2/2 checks passed	None	

Public DNS	Public IP	Key Name	Monitoring	Launch Time	Security Groups
		mucha	disabled	November 15, 2015 at 2:34:...	default
		mucha	disabled	May 8, 2015 at 5:06:32 PM ...	default
		mucha	disabled	November 25, 2015 at 5:07:...	residual-study
	52.4.178.24	windows1	disabled	April 26, 2015 at 10:40:59 A...	default

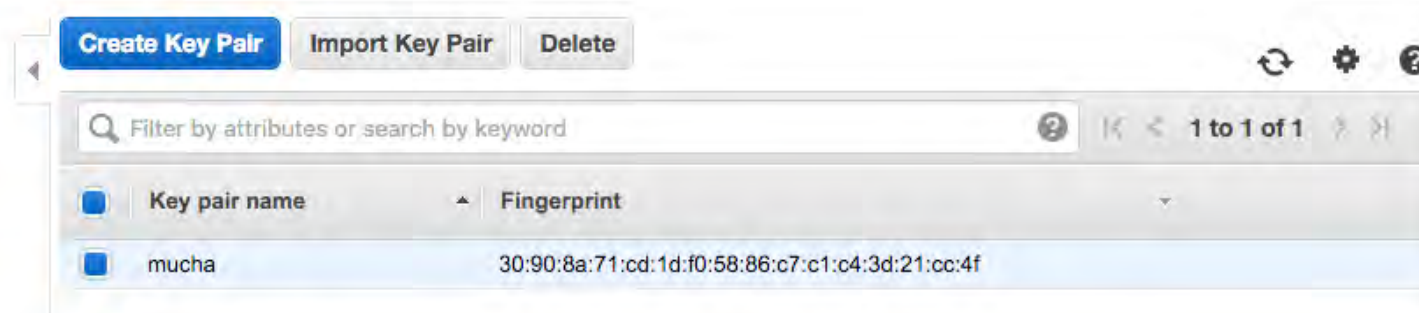
Accessing an instance: AWS key pairs

Linux instances are accessed via SSH (Secure Shell)

- AWS uses SSH “public key authentication.”
- Two ways to get your public key to Amazon:
 - You create a public/private keypair with “`ssh keygen -t rsa -f mykey.pem`” & import
 - Amazon will create the pair and you download it.
- You use the private key to authenticate.

Key pairs:

- Each key is identified by a “Fingerprint.”
- If you lose your private key, you can’t access your server.



Once an instance starts up, you can add additional users.
You can also access using network-based exploits.

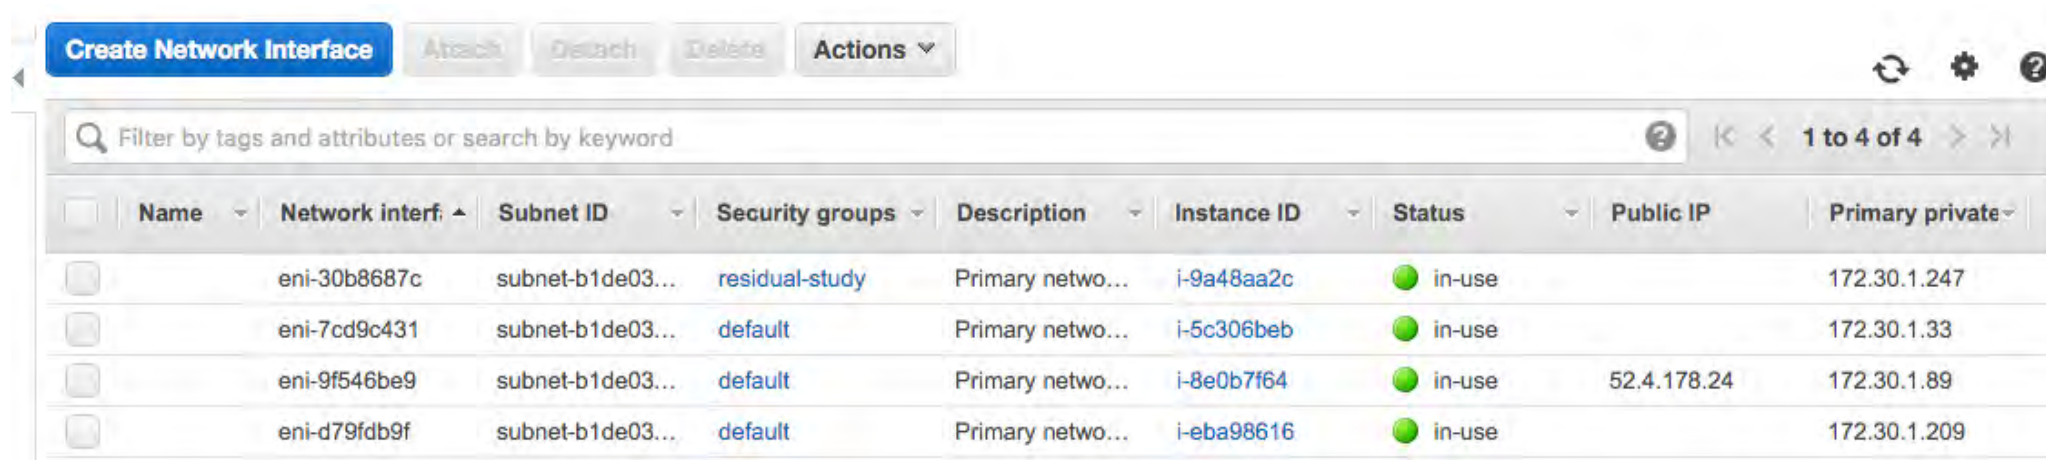
Each instance has at least 1 “virtual” interface, but possibly 2 IP addresses.

Amazon assigns a private IP address and (optionally) a public IP address.

- Private IP address is the “real” address on your private subnet.
- Amazon uses two-way NAT to provide the “public” address.
- NAT implements firewall through “security groups.”

Other options:

- You can have only private addresses. (More secure.)
- VPN to your organization.



	Name	Network interf.	Subnet ID	Security groups	Description	Instance ID	Status	Public IP	Primary private
<input type="checkbox"/>	eni-30b8687c		subnet-b1de03...	residual-study	Primary netwo...	i-9a48aa2c	in-use		172.30.1.247
<input type="checkbox"/>	eni-7cd9c431		subnet-b1de03...	default	Primary netwo...	i-5c306beb	in-use		172.30.1.33
<input type="checkbox"/>	eni-9f546be9		subnet-b1de03...	default	Primary netwo...	i-8e0b7f64	in-use	52.4.178.24	172.30.1.89
<input type="checkbox"/>	eni-d79fdb9f		subnet-b1de03...	default	Primary netwo...	i-eba98616	in-use		172.30.1.209

EC2 Instance forensic targets

Service-based acquisition:
Acquire from anywhere within AWS

Logs & ACL violations	Dynamo
S3: Simple Storage Service	Redshift
RDS: Relational Database Service	IAM
EFS: Elastic File System	Cloud Watch



Public Internet.
Many acquisition methodologies

Network Intercept

EBS: Elastic Block Service	
EC2 Instance	RAM
	Attached SSDs

Instance-based acquisition:
Acquire from the EC2 instance.



Instance-Based Acquisition

Instance-Based Acquisition: RAM

Traditional approaches for RAM acquisition in a virtualized environment:

- Suspend VM and access .vmem file
- Acquisition through the hypervisor
- Log into VM and run a RAM-dumping tool.

XenServer has the ability to dump memory too!

Acquisition on Amazon:

- Log into the VM and run a RAM-dumping tool.
- Amazon might be able to use the hypervisor to dump, but users can't!
—*You may wish to discuss this privately with Amazon under NDA...*


Instance-Based Acquisition: Disk

EC2 has two kinds of storage

Ephemeral storage / Instance Storage

- part of the instance (local drives) faster.

EBS — separate devices — slower, but can persist.



<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html>

Instance-Based Acquisition:

Instance Storage / Ephemeral Storage

Some EC2 instance types include internal storage

Instance Name	vCPUs	Memory	Local Storage	EBS-Optimized Bandwidth	Network Bandwidth
r5d.large	2	16 GiB	1 x 75 GB NVMe SSD	Up to 3.5 Gbps	Up to 10 Gbps
r5d.xlarge	4	32 GiB	1 x 150 GB NVMe SSD	Up to 3.5 Gbps	Up to 10 Gbps
r5d.2xlarge	8	64 GiB	1 x 300 GB NVMe SSD	Up to 3.5 Gbps	Up to 10 Gbps
r5d.4xlarge	16	128 GiB	2 x 300 GB NVMe SSD	3.5 Gbps	Up to 10 Gbps
r5d.12xlarge	48	384 GiB	2 x 900 GB NVMe SSD	7.0 Gbps	10 Gbps
r5d.24xlarge	96	768 GiB	4 x 900 GB NVMe SSD	14.0 Gbps	25 Gbps

The only way for you to acquire the NVMe SSD is to log into the instance.

- (Amazon might be able to acquire it using Xen hypervisor)

Instance storage cannot be used for the root file system. Most instance storage is now hardware encrypted.

Launch instance wizard | EC2 M x

console.aws.amazon.com/ec2/v2/home?region=us-east-1#LaunchInstanceWizard:

aws Services Resource Groups EC2 S3 Lambda DynamoDB WorkSpace Simson L. Garfinkel N. Virginia Support

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination	Encryption
Root	/dev/xvda	snap-020ec4f43b42c0023	8	General Purpose SSD (gp2)	100 / 3000	N/A	<input checked="" type="checkbox"/>	Not Encrypted
ephemeral0	/dev/nvme0n1	N/A	75	NVMe SSD	N/A	N/A	N/A	Hardware Encrypted

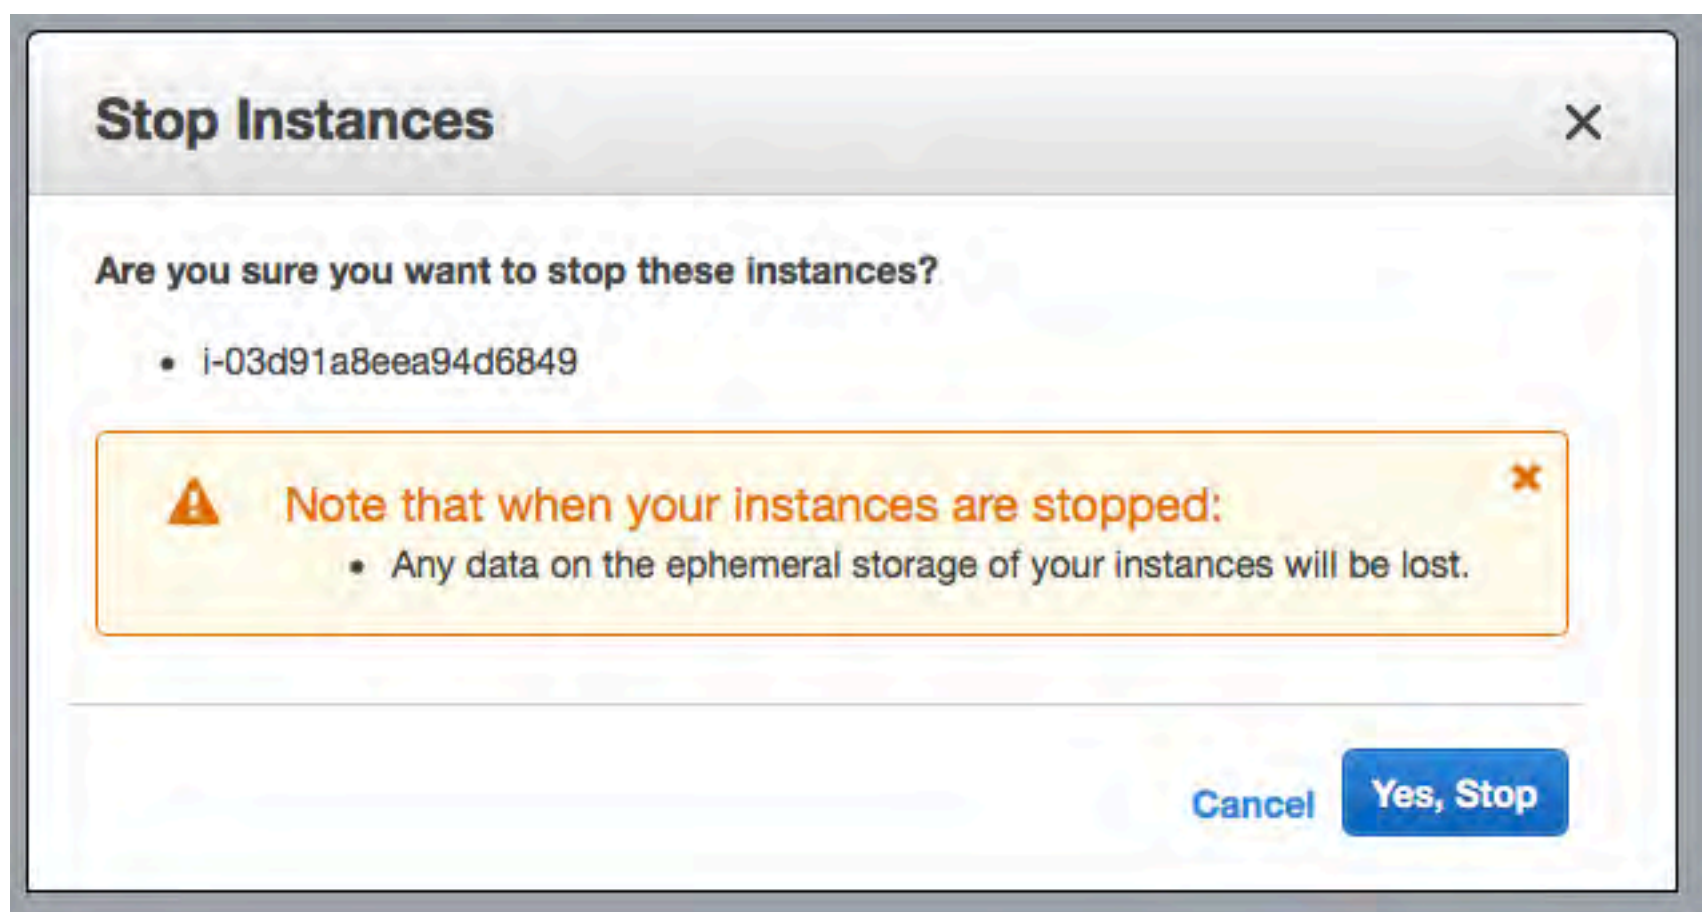
[Add New Volume](#)

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Tags](#)

Feedback English (US) © 2008 - 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use

Instance store is lost when an instance is stopped



EBS volumes can be initialized from snapshots...

EC2 Management Console

Secure | <https://console.aws.amazon.com/ec2/v2/home?region=us-east-1#LaunchInstanceWizard:>

Apps CFRS780 Discussion Board -... Schedule AWS EC2 Instance Pricing... GMU

aws Services Resource Groups

sgarfin2@gmu.edu N. Virginia Support

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination	Encrypted
Root	/dev/xvda	snap-0fae6f7252388fc12	40	General Purpose SSD (GP2)	120 / 3000	N/A	<input type="checkbox"/>	Not Encrypted
EBS	/dev/sdb	<input type="text" value="Search (case-insensit"/>	8	General Purpose SSD (GP2)	100 / 3000	N/A	<input type="checkbox"/>	<input type="checkbox"/>

[Add New Volume](#)

Free tier eligible customers can get u... usage restrictions.

Description

Created by CreateImage(l-f45e0d23) for ami-1d2bd976 from vol-ae64...
snapshot of Image.rootpart
prod-crxintl-wln-30May15
amzn-ami-pv-2015.09.0.x86_64
fedimg-snap-Fedora-Cloud-Base-Rawhide-20180117.n.1.x86_64
pvlinux-redhat-5.5-x86_1.3.0.148_121231_130201
Created by CreateImage(l-07a1e2da744e6e6c5) for ami-897ff8f3 fro...
hvm/ubuntu-trusty-amd64-server-20170104
pvlinux-redhat-5.5-x86_1.3.0.148_121204_165416

Snapshot ID

snap-1d1b6157
snap-0273df398...
snap-5b7fa1c6
snap-39941fad
snap-0aeb415a5...
snap-96f541da
snap-019a8e578...
snap-0fa5bfe69...
snap-6c502b27

[Previous](#) [Review and Launch](#) [Next: Add Tags](#)

Feedback English (US) © 2008 - 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use

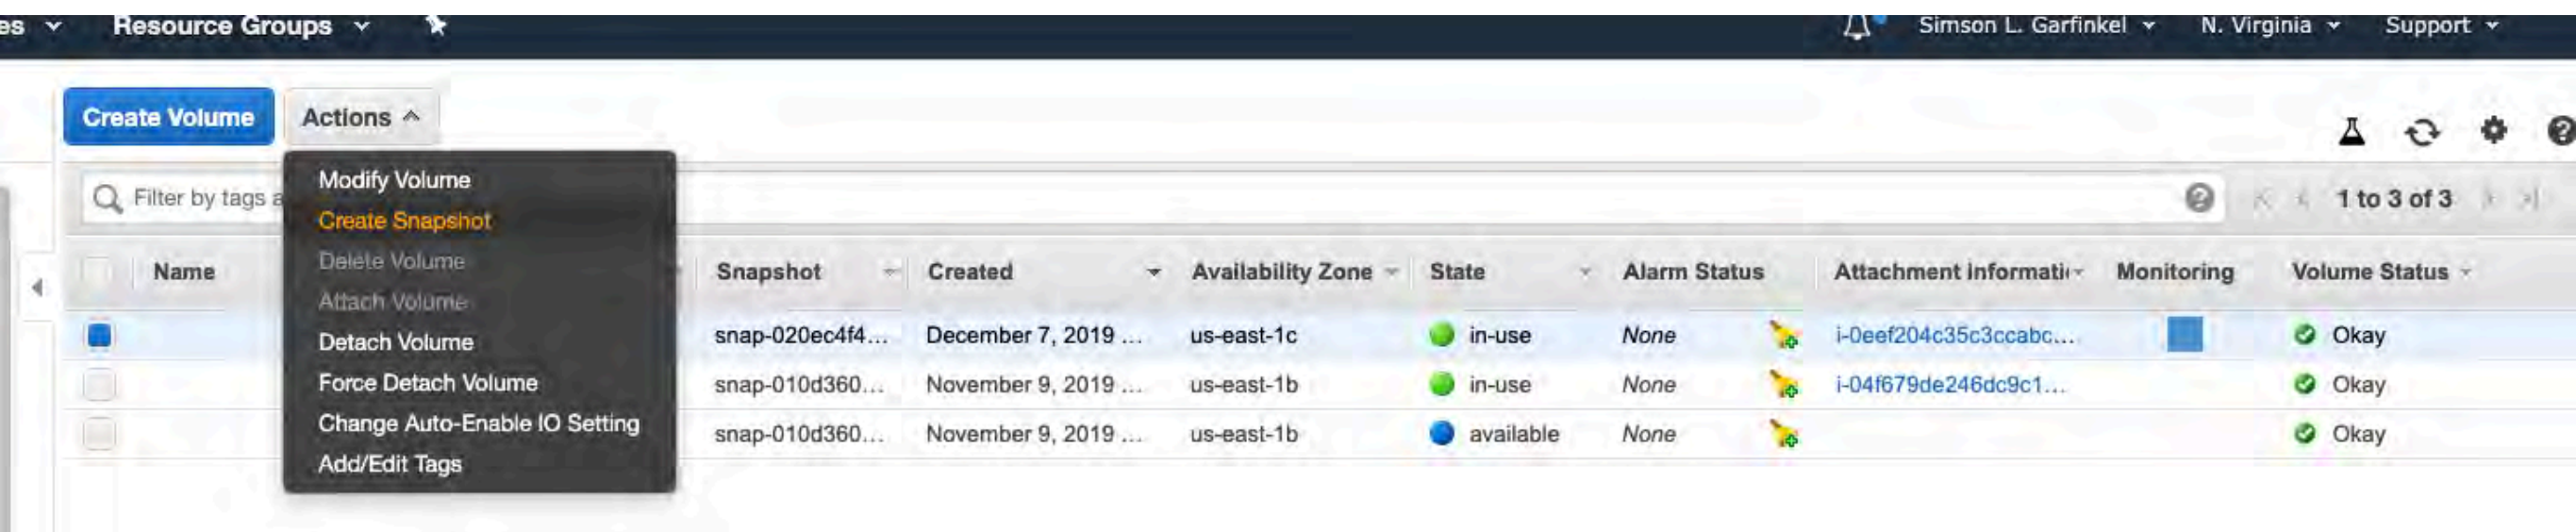
Forensic Acquisition: EBS (NETWORK ATTACHED STORAGE)

Option #1: Acquire through EC2 instance:

- Log into EC2 instance
- Run a traditional disk imaging program (e.g. dd, ewfacquire, etc.)
- Write disk image to:
 - Another device • Amazon S3 • Network socket

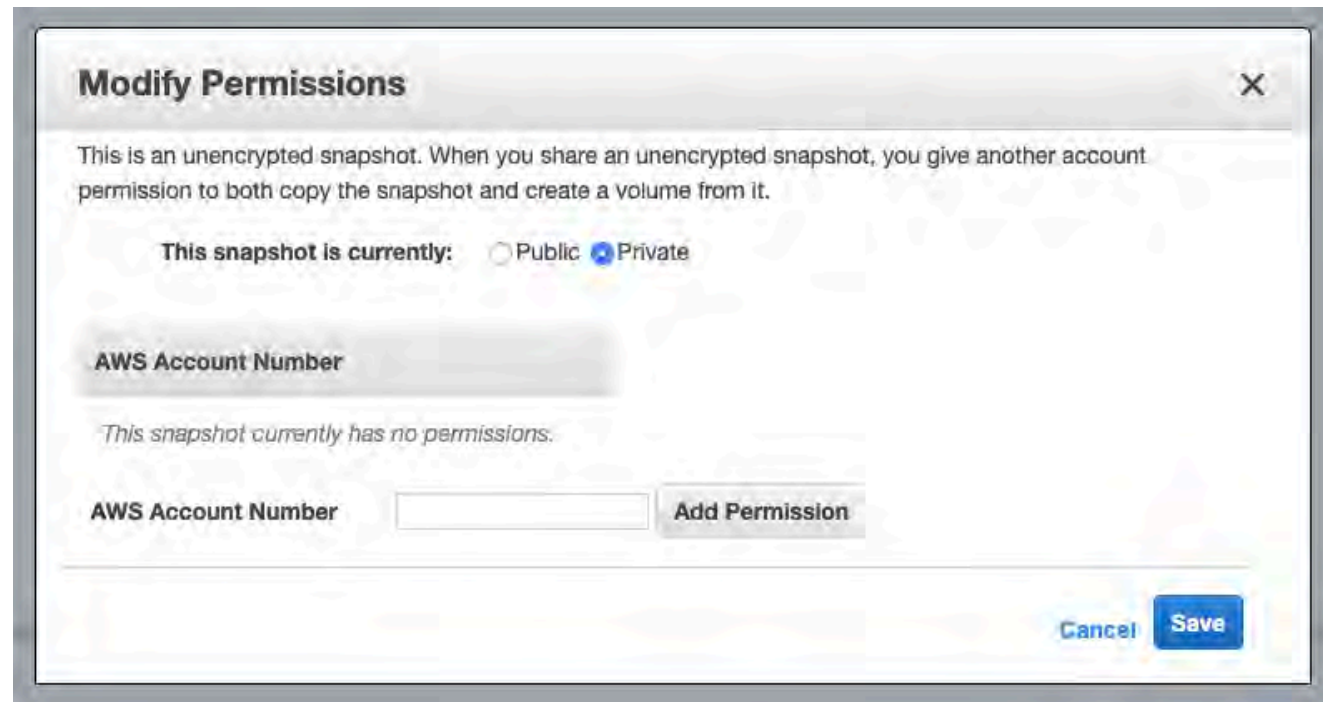
Option #2: Acquire through Amazon infrastructure:

- Use AWS GUI or CLI to snapshot the EBS volume.
- You can snapshot a running volume
- Snapshots are fast: 8GiB in a few seconds
- Restore the snapshot on a new volume on another system (CLI or GUI)



More on EBS snapshots

Snapshots can be Public or Private:



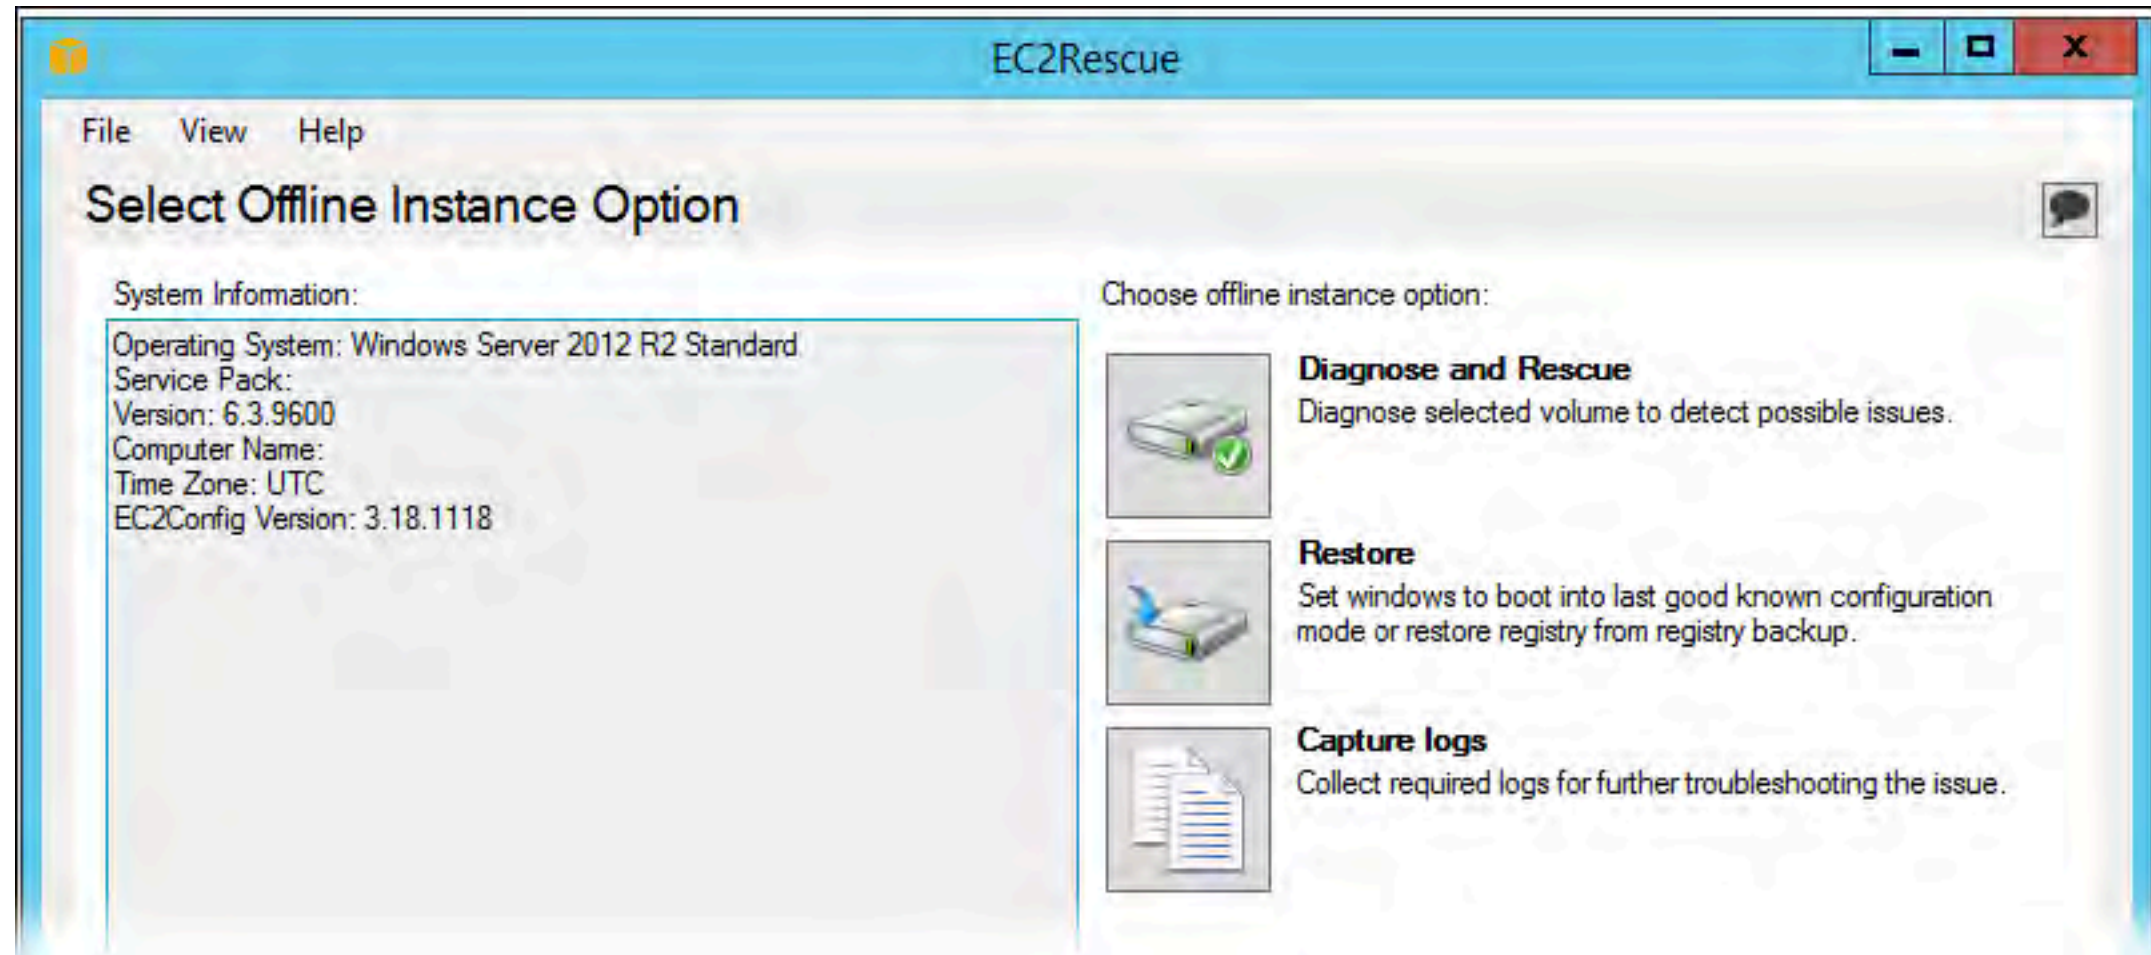
To read an EBS snapshot, restore it onto a new volume:



<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/EBSSnapshots.html>

AWS EC2 Rescue Tool

Analyzes offline instances to perform reset, restore, and log capture



<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/Windows-Server-EC2Rescue.html>

Many forensic capabilities!


Collect logs!

- <https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/ec2rw-cli.html>

FixAll — Fix things that tend to break on AWS

- System Time
- Windows Firewall
- Remote Desktop
- EC2 Config
- DHCP

ResetAccess — Reset the Administrator Password



Service-based Acquisition: S3


S3 is an object-based storage system

Every S3 bucket has:

- Name
- Owner
- Access permissions
- Region where located
- Optional event notifications
- Optional logging
- Optional static web hosting
- Optional “requester pays”
- Policy

Every S3 object has:

- Size
- URL
- Access permissions
 - e.g. *world readable*
- Optional encryption
- Optional policy




Amazon S3



AWS Regions



Buckets




Objects in the bucket

Accessing S3 data

Uses of S3:

- Storing logs
- Distributing data
- Objects for large-scale web apps (documents, JPEGs, etc.)



Advantages of S3:

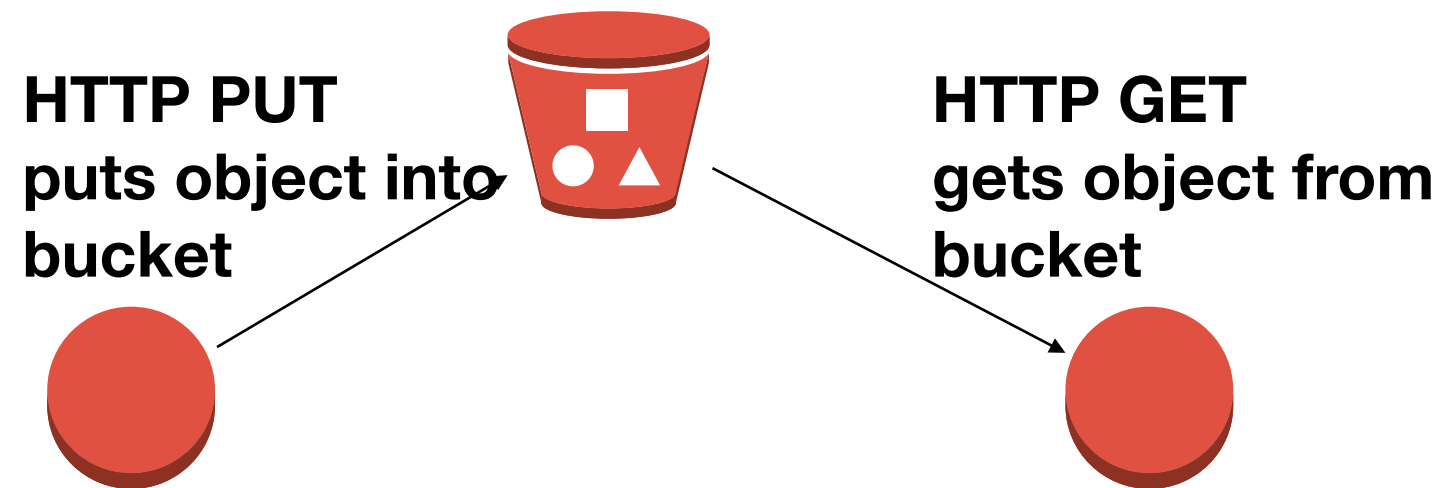
- permanence; S3 outlasts your EC2/EMR cluster*
- Pay only for what you need, rather than for virtual drives capacity.*

Disadvantage of S3

- No data locality. S3 data always moves over the network.*
- High-latency to access each object*
- Bulk data transfer must be done in parallel.*

S3 access protocol: REST

REST is built on top of HTTP.



Many ways to access data on Amazon S3

AWS GUI

- Simple object inspection can be done from a web browser

AWS CLI

```
aws s3 ls s3://bucketname/prefix/
```

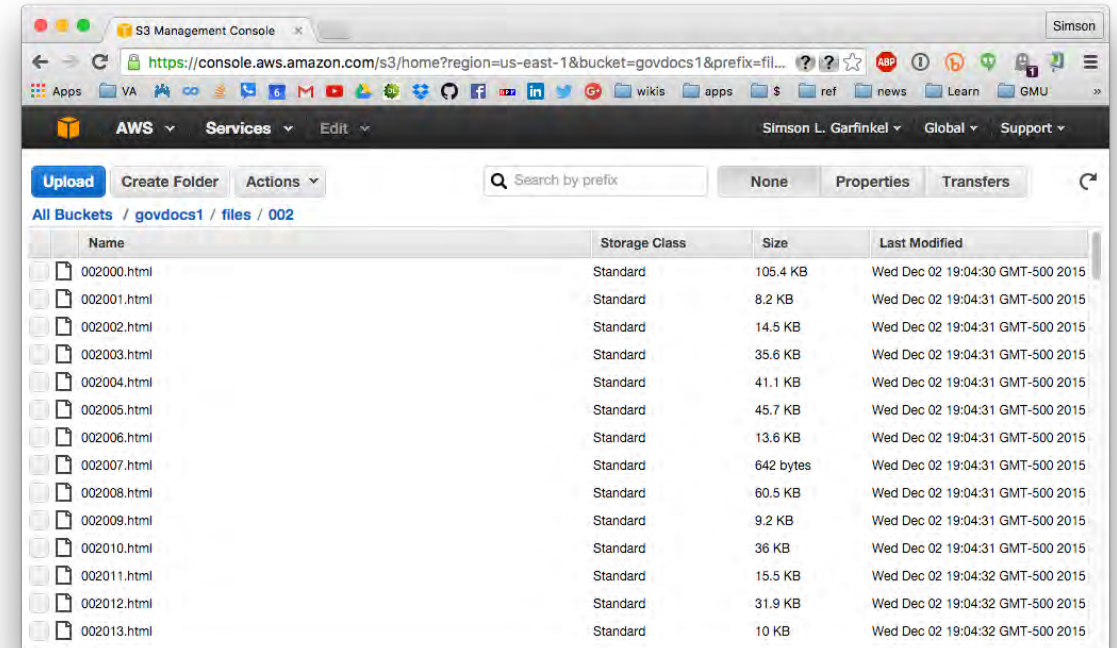
- Allows for parallelized upload and download.

HTTP / REST — Representational State Transfer

- Uses HTTP methods (with a bit of JSON)
- HTTP GET — Reads a resource without causing any side effects
- HTTP DELETE — Deletes a resources
- HTTP PUT (or POST) — Creates a new resources
- HTTP POST (or PUT) — Modify a resource's value

HTTP Hosting

- Different from REST
- Must be explicitly enabled



HTTP / SOAP — Simple Object Access Protocol

- Structure XML-based protocol
- Heavy weight; increasingly not used.

BitTorrent

- S3 can host a “tracker” and “seeds”
- Limited to objects 5GB in size

Amazon Snowball and Snowball Edge

50TB or 80TB of storage in a ruggedized container

e-ink shipping label

Snowball Edge includes:

- Amazon S3
- Amazon EC2
- Amazon Lambda



<https://aws.amazon.com/blogs/aws/aws-importexport-snowball-transfer-1-petabyte-per-week-using-amazon-owned-storage-appliances/>

Backup Slides



Creating an instance

Putting it all together...

EC2 Management Console

https://us-west-2.console.aws.amazon.com/ec2/v2/home?region=us-west-2#LaunchInstanceWizard:

AWS Services Edit

Simson Garfinkel ANLY502 Oregon Support

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Tag Instance 6. Configure Security Group 7. Review

Cancel and Exit

Step 1: Choose an Amazon Machine Image (AMI)

An AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. You can select an AMI provided by AWS, our user community, or the AWS Marketplace; or you can select one of your own AMIs.

Quick Start 1 to 22 of 22 AMIs

- Amazon Linux**
Free tier eligible
Amazon Linux AMI 2015.09.1 (HVM), SSD Volume Type - ami-f0091d91
The Amazon Linux AMI is an EBS-backed, AWS-supported image. The default image includes AWS command line tools, Python, Ruby, Perl, and Java. The repositories include Docker, PHP, MySQL, PostgreSQL, and other packages.
Root device type: ebs Virtualization type: hvm
64-bit
Select
- Red Hat**
Free tier eligible
Red Hat Enterprise Linux 7.1 (HVM), SSD Volume Type - ami-4dbf9e7d
Red Hat Enterprise Linux version 7.1 (HVM), EBS General Purpose (SSD) Volume Type
Root device type: ebs Virtualization type: hvm
64-bit
Select
- SUSE Linux**
Free tier eligible
SUSE Linux Enterprise Server 12 (HVM), SSD Volume Type - ami-d7450be7
SUSE Linux Enterprise Server 12 (HVM), EBS General Purpose (SSD) Volume Type. Public Cloud, Advanced Systems Management, Web and Scripting, and Legacy modules enabled.
Root device type: ebs Virtualization type: hvm
64-bit
Select
- Ubuntu**
Ubuntu Server 14.04 LTS (HVM), SSD Volume Type - ami-5189a661
Ubuntu Server 14.04 LTS (HVM), EBS General Purpose (SSD) Volume Type.
64-bit
Select

My AMIs
AWS Marketplace
Community AMIs
☐ Free tier only ⓘ

Feedback English

© 2008 - 2015, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use

Putting it all together...

Launch instance wizard | EC2 M x +

console.aws.amazon.com/ec2/home?region=us-east-1#LaunchInstanceWizard

aws Services Resource Groups

Simson L. Garfinkel N. Virginia Support

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 1: Choose an Amazon Machine Image (AMI)

[Cancel and Exit](#)

An AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. You can select an AMI provided by AWS, our user community, or the AWS Marketplace; or you can select one of your own AMIs.

Search for an AMI by entering a search term e.g. "Windows"






Quick Start 1 to 40 of 40 AMIs

My AMIs

AWS Marketplace

Community AMIs

☐ Free tier only

 <p>Amazon Linux Free tier eligible</p>	<p>Amazon Linux 2 AMI (HVM), SSD Volume Type - ami-00068cd7555f543d5 (64-bit x86) / ami-035240afa793cddb</p> <p>(64-bit Arm)</p> <p>Amazon Linux 2 comes with five years support. It provides Linux kernel 4.14 tuned for optimal performance on Amazon EC2, systemd 219, GCC 7.3, Glibc 2.26, Binutils 2.29.1, and the latest software packages through extras.</p> <p>Root device type: ebs Virtualization type: hvm ENA Enabled: Yes</p>	<p>Select</p> <p><input checked="" type="radio"/> 64-bit (x86) <input type="radio"/> 64-bit (Arm)</p>
 <p>Amazon Linux Free tier eligible</p>	<p>Amazon Linux AMI 2018.03.0 (HVM), SSD Volume Type - ami-00eb20669e0990cb4</p> <p>The Amazon Linux AMI is an EBS-backed, AWS-supported image. The default image includes AWS command line tools, Python, Ruby, Perl, and Java. The repositories include Docker, PHP, MySQL, PostgreSQL, and other packages.</p> <p>Root device type: ebs Virtualization type: hvm ENA Enabled: Yes</p>	<p>Select</p> <p>64-bit (x86)</p>
 <p>Red Hat Free tier eligible</p>	<p>Red Hat Enterprise Linux 8 (HVM), SSD Volume Type - ami-0c322300a1dd5dc79 (64-bit x86) / ami-03587fa4048e9eb92 (64-bit Arm)</p> <p>Red Hat Enterprise Linux version 8 (HVM), EBS General Purpose (SSD) Volume Type</p> <p>Root device type: ebs Virtualization type: hvm ENA Enabled: Yes</p>	<p>Select</p> <p><input checked="" type="radio"/> 64-bit (x86) <input type="radio"/> 64-bit (Arm)</p>
 <p>SUSE Linux Free tier eligible</p>	<p>SUSE Linux Enterprise Server 15 SP1 (HVM), SSD Volume Type - ami-0547b1fd62b28a111 (64-bit x86) / ami-008a07c569b8da5ca (64-bit Arm)</p> <p>SUSE Linux Enterprise Server 15 Service Pack 1 (HVM), EBS General Purpose (SSD) Volume Type. Public Cloud, Advanced Systems Management, Web and Scripting, and Legacy modules enabled.</p> <p>Root device type: ebs Virtualization type: hvm ENA Enabled: Yes</p>	<p>Select</p> <p><input checked="" type="radio"/> 64-bit (x86) <input type="radio"/> 64-bit (Arm)</p>
 <p>Ubuntu</p>	<p>Ubuntu Server 18.04 LTS (HVM), SSD Volume Type - ami-04b9e92b5572fa0d1 (64-bit x86) / ami-</p>	<p>Select</p>

Feedback English (US)

© 2008 - 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use

Launch instance wizard | EC2 M

console.aws.amazon.com/ec2/home?region=us-east-1#LaunchInstanceWizard:

aws

Services

Resource Groups

Simson L. Garfinkel

N. Virginia

Support

1. Choose AMI

2. Choose Instance Type

3. Configure Instance

4. Add Storage

5. Add Tags

6. Configure Security Group

7. Review

Step 2: Choose an Instance Type

Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instances are virtual servers that can run applications. They have varying combinations of CPU, memory, storage, and networking capacity, and give you the flexibility to choose the appropriate mix of resources for your applications. [Learn more](#) about instance types and how they can meet your computing needs.

Filter by: All instance types Current generation [Show/Hide Columns](#)

Currently selected: t2.micro (Variable ECUs, 1 vCPUs, 2.5 GHz, Intel Xeon Family, 1 GiB memory, EBS only)

	Family	Type	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance	IPv6 Support
<input type="checkbox"/>	General purpose	t2.nano	1	0.5	EBS only	-	Low to Moderate	Yes
<input checked="" type="checkbox"/>	General purpose	t2.micro Free tier eligible	1	1	EBS only	-	Low to Moderate	Yes
<input type="checkbox"/>	General purpose	t2.small	1	2	EBS only	-	Low to Moderate	Yes
<input type="checkbox"/>	General purpose	t2.medium	2	4	EBS only	-	Low to Moderate	Yes
<input type="checkbox"/>	General purpose	t2.large	2	8	EBS only	-	Low to Moderate	Yes
<input type="checkbox"/>	General purpose	t2.xlarge	4	16	EBS only	-	Moderate	Yes
<input type="checkbox"/>	General purpose	t2.2xlarge	8	32	EBS only	-	Moderate	Yes
<input type="checkbox"/>	General purpose	t3a.nano	2	0.5	EBS only	Yes	Up to 5 Gigabit	Yes
<input type="checkbox"/>	General purpose	t3a.micro	2	1	EBS only	Yes	Up to 5 Gigabit	Yes
<input type="checkbox"/>	General purpose	t3a.small	2	2	EBS only	Yes	Up to 5 Gigabit	Yes
<input type="checkbox"/>	General purpose	t3a.medium	2	4	EBS only	Yes	Up to 5 Gigabit	Yes
<input type="checkbox"/>	General purpose	t3a.large	2	8	EBS only	Yes	Up to 5 Gigabit	Yes

Cancel

Previous

Review and Launch

Next: Configure Instance Details

Feedback

English (US)

© 2008 - 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Privacy Policy

Terms of Use

Launch instance wizard | EC2 M x

console.aws.amazon.com/ec2/home?region=us-east-1#LaunchInstanceWizard:

aws Services Resource Groups

Simson L. Garfinkel N. Virginia Support

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 3: Configure Instance Details

Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances to take advantage of the lower pricing, assign an access management role to the instance, and more.

Number of instances ⓘ 1 [Launch into Auto Scaling Group](#) ⓘ

Purchasing option ⓘ ☐ Request Spot instances

Network ⓘ vpc-8e73cfeb [Create new VPC](#)

Subnet ⓘ subnet-8357abda | us-east-1c [Create new subnet](#)
251 IP Addresses available

Auto-assign Public IP ⓘ Use subnet setting (Enable)

Placement group ⓘ ☐ Add instance to placement group

Capacity Reservation ⓘ Open [Create new Capacity Reservation](#)

IAM role ⓘ None [Create new IAM role](#)

Shutdown behavior ⓘ Stop

Enable termination protection ⓘ ☐ Protect against accidental termination

Monitoring ⓘ ☐ Enable CloudWatch detailed monitoring
[Additional charges apply.](#)

Tenancy ⓘ Shared - Run a shared hardware instance
[Additional charges will apply for dedicated tenancy.](#)

Elastic Inference ⓘ ☐ Add an Elastic Inference accelerator
[Additional charges apply.](#)

T2/T3 Unlimited ⓘ ☐ Enable
[Additional charges may apply](#)

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Storage](#)

Launch instance wizard | EC2 M x

console.aws.amazon.com/ec2/home?region=us-east-1#LaunchInstanceWizard:

aws Services Resource Groups

Simson L. Garfinkel N. Virginia Support

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 3: Configure Instance Details

Shutdown behavior *i* Stop

Enable termination protection *i* ☐ Protect against accidental termination

Monitoring *i* ☐ Enable CloudWatch detailed monitoring
Additional charges apply.

Tenancy *i* Shared - Run a shared hardware instance
Additional charges will apply for dedicated tenancy.

Elastic Inference *i* ☐ Add an Elastic Inference accelerator
Additional charges apply.

T2/T3 Unlimited *i* ☐ Enable
Additional charges may apply

File systems *i*

▼ **Network interfaces** *i*

Device	Network Interface	Subnet	Primary IP	Secondary IP addresses	IPv6 IPs
eth0	New network interface	subnet-8357abdi	Auto-assign	Add IP	Add IP

▼ **Advanced Details**

User data *i* ☒ As text ☐ As file ☐ Input is already base64 encoded

(Optional)

Launch instance wizard | EC2 M x +

console.aws.amazon.com/ec2/home?region=us-east-1#LaunchInstanceWizard:

aws Services Resource Groups

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination	Encryption
Root	/dev/xvda	snap-020ec4f43b42c0023	8	General Purpose SSD (gp2)	100 / 3000	N/A	<input checked="" type="checkbox"/>	Not Encrypte

[Add New Volume](#)

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

Volume Type

- ✓ General Purpose SSD (gp2)
- Provisioned IOPS SSD (io1)
- Magnetic (standard)

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Tags](#)



Launch instance wizard | EC2 M x

console.aws.amazon.com/ec2/home?region=us-east-1#LaunchInstanceWizard:

aws Services Resource Groups

Simson L. Garfinkel N. Virginia Support

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 7: Review Instance Launch

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

AMI Details [Edit AMI](#)

Amazon Linux 2 AMI (HVM), SSD Volume Type - ami-00068cd7555f543d5

Free tier eligible Amazon Linux 2 comes with five years support. It provides Linux kernel 4.14 tuned for optimal performance on Amazon EC2, systemd 219, GCC 7.3, Glibc 2.26, Binutils 2.29.1, and the latest software packages through extras.

Root Device Type: ebs Virtualization

Instance Type [Edit instance type](#)

Instance Type	ECUs
t2.micro	Variable

Security Groups [Edit security groups](#)

Security group name	Description
launch-w	launch-w

Type (1)

Instance Details [Edit instance details](#)

Storage [Edit storage](#)

Tags [Edit tags](#)

Network Performance

Low to Moderate

Description (1)

[Cancel](#) [Previous](#) [Launch](#)

Select an existing key pair or create a new key pair

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about [removing existing key pairs from a public AMI](#).

Choose an existing key pair

Select a key pair

id_rsa

☐ I acknowledge that I have access to the selected private key file (id_rsa.pem), and that without this file, I won't be able to log into my instance.

[Cancel](#) [Launch instances](#)

Feedback English (US)

©2008–2019, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use


Launch instance wizard | EC2 M x


console.aws.amazon.com/ec2/home?region=us-east-1#LaunchInstanceWizard:

aws Services Resource Groups

Simson L. Garfinkel N. Virginia Support

Launch Status

 **Your instances are now launching**
The following instance launches have been initiated: [i-0eef204c35c3ccabc](#) [View launch log](#)

 **Get notified of estimated charges**
[Create billing alerts](#) to get an email notification when estimated charges on your AWS bill exceed an amount you define (for example, if you exceed the free usage tier).

How to connect to your instances

Your instances are launching, and it may take a few minutes until they are in the **running** state, when they will be ready for you to use. Usage hours on your new instances will start immediately and continue to accrue until you stop or terminate your instances.

Click **View Instances** to monitor your instances' status. Once your instances are in the **running** state, you can **connect** to them from the Instances screen. [Find out](#) how to connect to your instances.

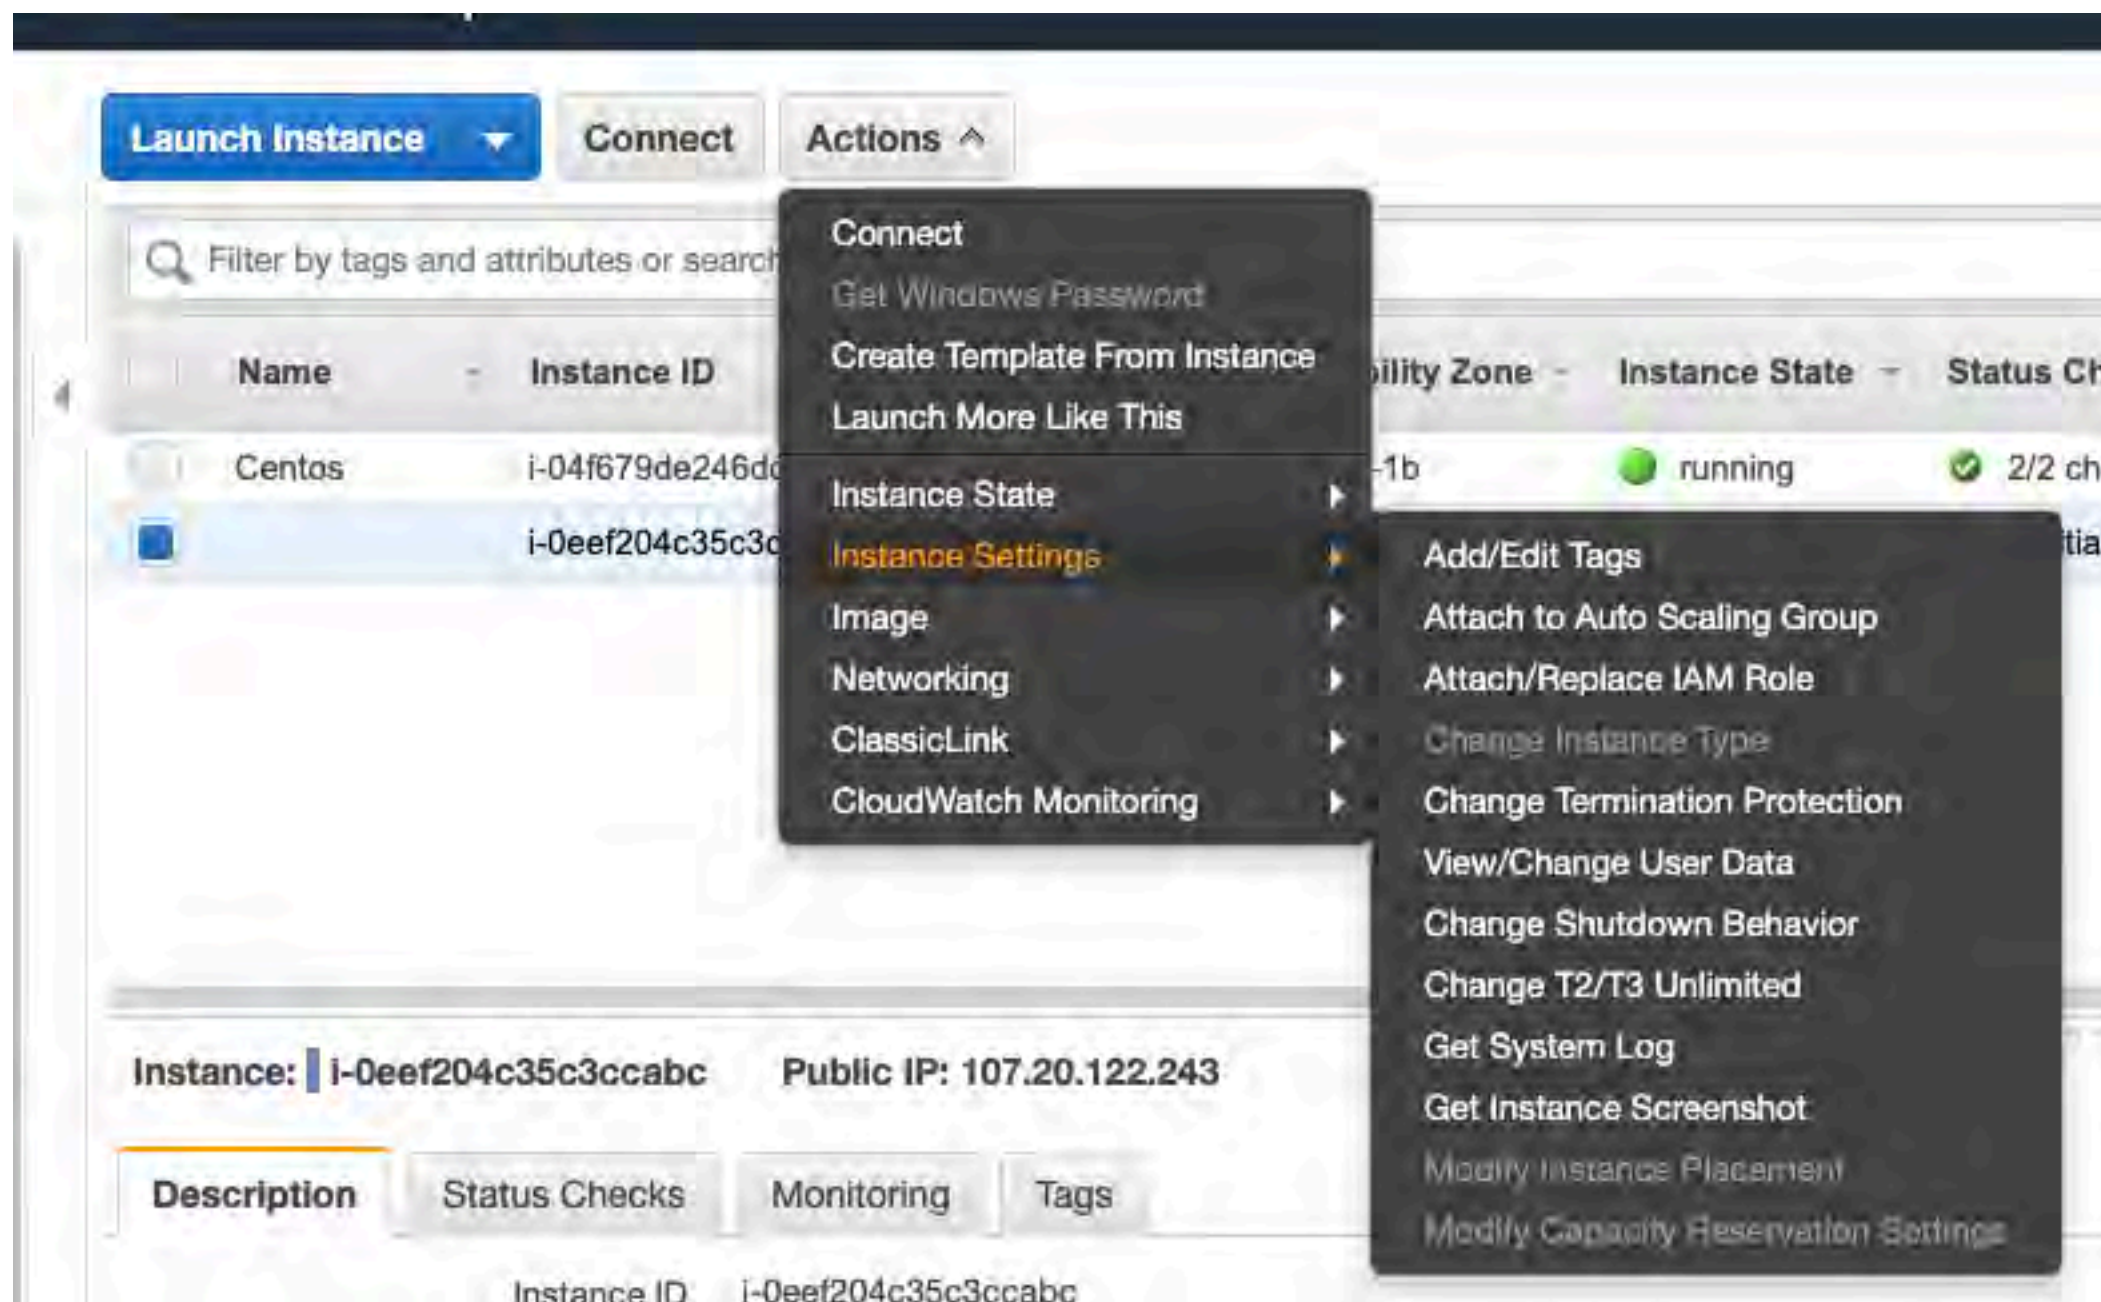
▼ Here are some helpful resources to get you started

- [How to connect to your Linux instance](#)
- [Learn about AWS Free Usage Tier](#)
- [Amazon EC2: User Guide](#)
- [Amazon EC2: Discussion Forum](#)

While your instances are launching you can also

- [Create status check alarms](#) to be notified when these instances fail status checks. (Additional charges may apply)
- [Create and attach additional EBS volumes](#) (Additional charges may apply)
- [Manage security groups](#)

[View Instances](#)



System.

Starting Reload Configuration from the Real Root...

```
[ [32m OK [0m] Started Reload Configuration from the Real Root.
[ [32m OK [0m] Reached target Initrd File Systems.
[ [32m OK [0m] Reached target Initrd Default Target.
```

Starting dracut pre-pivot and cleanup hook...

```
[ [32m OK [0m] Started dracut pre-pivot and cleanup hook.
```


Starting Cleaning Up and Shutting Down Daemons...

```
[ [32m OK [0m] Stopped target Timers.
[ [32m OK [0m] Stopped Cleaning Up and Shutting Down Daemons.
[ [32m OK [0m] Stopped dracut pre-pivot and cleanup hook.
[ [32m OK [0m] Stopped target Remote File Systems.
[ [32m OK [0m] Stopped target Remote File Systems (Pre).
[ [32m OK [0m] Stopped dracut initqueue hook.
[ [32m OK [0m] Stopped target Initrd Default Target.
[ [32m OK [0m] Stopped target Basic System.
[ [32m OK [0m] Stopped target Sockets.
[ [32m OK [0m] Stopped target System Initialization.
[ [32m OK [0m] Stopped Apply Kernel Variables.
[ [32m OK [0m] Stopped udev Coldplug all Devices.
[ [32m OK [0m] Stopped dracut pre-trigger hook.
[ [32m OK [0m] Stopped target Swap.
[ [32m OK [0m] Stopped target Local File Systems.
[ [32m OK [0m] Stopped target Slices.
[ [32m OK [0m] Stopped target Paths.
[ [32m OK [0m] Stopped Dispatch Password Requests to Console Directory Watch.
[ [32m OK [0m] Stopped udev Kernel Device Manager.
[ [32m OK [0m] Stopped Create Static Device Nodes in /dev.
[ [32m OK [0m] Stopped Create list of required sta...ce nodes for the current ker
[ [32m OK [0m] Stopped dracut pre-udev hook.
[ [32m OK [0m] Stopped dracut cmdline hook.
```

Close

Get instance screenshot

Below is a screenshot of i-0eef204c35c3ccabc at 2019-12-07T14:46:08.141-05:00.

 Refresh




```
Amazon Linux 2  
Kernel 4.14.152-127.182.amzn2.x86_64 on an x86_64  
ip-172-30-4-236 login: _
```

Close

Instance is running...

Launch Instance Connect Actions ▾

Filter by tags and attributes or search by keyword

<input type="checkbox"/>	Name ▾	Instance ID ▴	Instance Type ▾	Availability Zone ▾	Instance State ▾	Status Checks ▾	Alarm Status
<input type="checkbox"/>		i-3b4c05ff	t2.micro	us-west-2b	 running	 2/2 checks ...	None 

1 to 1 of 1

Public DNS ▾	Public IP ▾	Key Name ▾	Monitoring ▾	Launch Time ▾	Security Groups ▾
ec2-52-33-99-98.us-we...	52.33.99.98	anly502	<input type="checkbox"/> disabled	November 29, 2015 at 4:04:...	launch-wizard-1

Connect...

```
simsong@nimi ~ % ssh ec2-user@107.20.122.243
```

```
The authenticity of host '107.20.122.243 (107.20.122.243)' can't be established.  
ECDSA key fingerprint is SHA256:MKlTMdgi3FvK9rCSe++Q0Bt+/MQfqicf63pkVsD9YDk.
```

```
Are you sure you want to continue connecting (yes/no)? yes
```

```
Warning: Permanently added '107.20.122.243' (ECDSA) to the list of known hosts.
```

```
  ____|  ____|  )  
  _|  (  ____|  /  Amazon Linux 2 AMI  
  ____| \ ____|  |
```

```
https://aws.amazon.com/amazon-linux-2/
```

```
5 package(s) needed for security, out of 13 available
```

```
Run "sudo yum update" to apply all updates.
```

```
[ec2-user@ip-172-30-4-236 ~]$
```


We have a running instance!

```
[ec2-user@ip-172-30-4-236 ~]$ df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        475M   0  475M   0% /dev
tmpfs           492M   0  492M   0% /dev/shm
tmpfs           492M 400K  492M   1% /run
tmpfs           492M   0  492M   0% /sys/fs/cgroup
/dev/xvda1      8.0G 1.3G  6.8G  16% /
tmpfs           99M   0   99M   0% /run/user/1000
[ec2-user@ip-172-30-4-236 ~]$ top
top - 19:50:44 up 6 min,  1 user,  load average: 0.00, 0.05, 0.03
Tasks:  83 total,   1 running, 46 sleeping,   0 stopped,   0 zombie
%Cpu(s):  0.0 us,   0.0 sy,   0.0 ni,100.0 id,   0.0 wa,   0.0 hi,   0.0 si,   0.0 st
KiB Mem : 1007276 total,   610048 free,    60176 used,   337052 buff/cache
KiB Swap:   0 total,         0 free,         0 used.  807136 avail Mem
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
1	root	20	0	125512	5500	4064	S	0.0	0.5	0:01.59	systemd
2	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kthreadd
4	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	kworker/0:0H
5	root	20	0	0	0	0	I	0.0	0.0	0:00.02	kworker/u30:0
6	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	mm_percpu_wq
7	root	20	0	0	0	0	S	0.0	0.0	0:00.04	ksoftirqd/0
8	root	20	0	0	0	0	I	0.0	0.0	0:00.15	rcu_sched
9	root	20	0	0	0	0	I	0.0	0.0	0:00.00	rcu_bh
10	root	rt	0	0	0	0	S	0.0	0.0	0:00.00	migration/0
11	root	rt	0	0	0	0	S	0.0	0.0	0:00.00	watchdog/0
12	root	20	0	0	0	0	S	0.0	0.0	0:00.00	cpuhp/0
14	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kdevtmpfs
15	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	netns
21	root	20	0	0	0	0	S	0.0	0.0	0:00.00	xenbus



EC2 Command Line Tools

Amazon provides command line tools

Can be run from *any* Linux, Mac or Windows computer.

- Faster interaction than web interface.
- Can be scripted.

AWS Command Line Interface

- Run through “aws” command
- Flexible output — JSON, text, tables
- List EC2 instance: `$ aws ec2 describe-instances`

<https://aws.amazon.com/cli/>

<http://docs.aws.amazon.com/cli/latest/userguide/cli-chap-welcome.html>

Elastic Comput Cloud CLI

- Run through 176 different `ec2-*` commands
- List EC2 instances: `$ ec2-describe-instances`

<http://docs.aws.amazon.com/AWSEC2/latest/CommandLineReference/ApiReference-cmd-DescribeVolumes.html>

Credentials:

- Credentials kept in `$HOME/.aws/` directory
- Credentials kept in `AWS_USERNAME`, `AWS_ACCESS_KEY`, `AWS_SECRET_KEY` environment variables.

Both are pre-installed on Amazon’s AMIs. Use the AWS CLI if possible.

Set up your environment variables and test:

AWS CLI command:

```
$ aws ec2 describe-regions
REGIONS      ec2.eu-north-1.amazonaws.com eu-north-1
REGIONS      ec2.ap-south-1.amazonaws.com ap-south-1
REGIONS      ec2.eu-west-3.amazonaws.com  eu-west-3
REGIONS      ec2.eu-west-2.amazonaws.com  eu-west-2
REGIONS      ec2.eu-west-1.amazonaws.com  eu-west-1
REGIONS      ec2.ap-northeast-2.amazonaws.com  ap-northeast-2
REGIONS      ec2.ap-northeast-1.amazonaws.com  ap-northeast-1
REGIONS      ec2.sa-east-1.amazonaws.com  sa-east-1
REGIONS      ec2.ca-central-1.amazonaws.com  ca-central-1
REGIONS      ec2.ap-southeast-1.amazonaws.com  ap-southeast-1
REGIONS      ec2.ap-southeast-2.amazonaws.com  ap-southeast-2
REGIONS      ec2.eu-central-1.amazonaws.com  eu-central-1
REGIONS      ec2.us-east-1.amazonaws.com  us-east-1
REGIONS      ec2.us-east-2.amazonaws.com  us-east-2
REGIONS      ec2.us-west-1.amazonaws.com  us-west-1
REGIONS      ec2.us-west-2.amazonaws.com  us-west-2
```


EC2 has a command-line interface

Show running instances:

```
[ec2-user@ip-172-30-4-236 ~]$ aws ec2 describe-instances
RESERVATIONS      376778049323      086189789714      r-08090e555597ca8cd
INSTANCES         0                x86_64            157333535705912490  False      True      xen
ami-02eac2c0129f6376b      i-04f679de246dc9c10      t2.micro      id_rsa
2019-11-15T19:18:05.000Z      ip-172-30-1-55.ec2.internal      172.30.1.55
18.212.220.250      /dev/sda1      ebs      True      subnet-blde03c6      hvm
vpc-8e73cfeb
BLOCKDEVICEMAPPINGS      /dev/sda1
EBS      2019-11-09T21:36:00.000Z      False      attached      vol-0aab795976166105c
CAPACITYRESERVATIONSPECIFICATION      open
CPUOPTIONS      1      1
HIBERNATIONOPTIONS      False
MONITORING      disabled
NETWORKINTERFACES      Primary network interface      0a:e8:03:05:29:67
eni-094d36080627e1676      376778049323      172.30.1.55      True      in-use      subnet-
blde03c6      vpc-8e73cfeb
ASSOCIATION      amazon      18.212.220.250
ATTACHMENT      2019-11-09T21:35:59.000Z      eni-attach-092f65b6f1f7f26be      True
0      attached
GROUPS      sg-06ec94cd40903890a      CentOS 7 -x86_64- - with Updates HVM-1901_01-
AutogenByAWSMP-1
PRIVATEIPADDRESSES      True      172.30.1.55
ASSOCIATION      amazon      18.212.220.250
PLACEMENT      us-east-1b      default
PRODUCTCODES      aw0evgkw8e5c1q413zgy5pjce      marketplace
SECURITYGROUPS      sg-06ec94cd40903890a      CentOS 7 -x86_64- - with Updates HVM-1901_01-
AutogenByAWSMP-1
STATE 16      running
TAGS Name      Centos
```

Use “help” to get help

```
$ aws ec2 describe-instances help
```

NAME

```
describe-instances -
```

DESCRIPTION

Describes one or more of your instances.

If you specify one or more instance IDs, Amazon EC2 returns information for those instances. If you do not specify instance IDs, Amazon EC2 returns information for all relevant instances. If you specify an instance ID that is not valid, an error is returned. If you specify an instance that you do not own, it is not included in the returned results.

Recently terminated instances might appear in the returned results. This interval is usually less than one hour.

`describe-instances` is a paginated operation. Multiple API calls may be issued in order to retrieve the entire data set of results. You can disable pagination by providing the `--no-paginate` argument. When using `--output text` and the `--query` argument on a paginated response, the `--query` argument must extract data from the results of the following query expressions: `Reservations`

\$ ec2-describe-instance-status — see what's running

```
$ aws ec2 describe-instance-status --output=text
INSTANCESTATUSES  us-east-1b  i-5c306beb
INSTANCESTATE     16          running
INSTANCESTATUS    ok
DETAILS           reachability passed
SYSTEMSTATUS      ok
DETAILS           reachability passed
INSTANCESTATUSES  us-east-1b  i-9a48aa2c
INSTANCESTATE     16          running
INSTANCESTATUS    ok
DETAILS           reachability passed
SYSTEMSTATUS      ok
DETAILS           reachability passed
$
```

Change output format:

```
[ec2-user@ip-172-30-4-236 ~]$ aws ec2 describe-instances --output table
```


DescribeInstances	
Reservations	
OwnerId	376778049323
RequesterId	086189789714
ReservationId	r-08090e555597ca8cd
Instances	
AmiLaunchIndex	0
Architecture	x86_64
ClientToken	157333535705912490
EbsOptimized	False
EnaSupport	True
Hypervisor	xen
ImageId	ami-02eac2c0129f6376b
InstanceId	i-04f679de246dc9c10
InstanceType	t2.micro
KeyName	id_rsa
LaunchTime	2019-11-15T19:18:05.000Z
PrivateDnsName	ip-172-30-1-55.ec2.internal
PrivateIpAddress	172.30.1.55

JSON output is more useful for scripting

```
$ aws ec2 describe-instance-status --output=json
  "Reservations": [
    {
      "Instances": [
        {
          "Monitoring": {
            "State": "disabled"
          },
          "PublicDnsName": "",
          "State": {
            "Code": 16,
            "Name": "running"
          },
          "EbsOptimized": false,
          "LaunchTime": "2019-11-15T19:18:05.000Z",
          "PublicIpAddress": "18.212.220.250",
          "PrivateIpAddress": "172.30.1.55",
          "ProductCodes": [
            {
              "ProductCodeId": "aw0evgkw8e5c1q413zgy5pjce",
              "ProductCodeType": "marketplace"
            }
          ],
          "VpcId": "vpc-8e73cfef",
          "CpuOptions": {
            "CoreCount": 1,
            "ThreadsPerCore": 1
          },
          "StateTransitionReason": "",
          "InstanceId": "i-04f679de246dc9c10",
          "EnaSupport": true,
          "ImageId": "ami-02eac2c0129f6376b",
          "PrivateDnsName": "ip-172-30-1-55.ec2.internal",
          "KeyName": "id_rsa",
          "SecurityGroups": [
```


Get console output!

```
$ aws ec2 get-console-output --instance-id i-0042cc0b3e4175345 --output text
i-0042cc0b3e4175345 [ 0.000000] Linux version 4.9.76-3.78.amzn1.x86_64 (mockbuild@gobi-
build-60009) (gcc version 7.2.1 20170915 (Red Hat 7.2.1-2) (GCC) ) #1 SMP Fri Jan 12 19:51:35 UTC 2018
[ 0.000000] Command line: root=LABEL=/ console=tty1 console=ttyS0 selinux=0
nvme_core.io_timeout=4294967295
[ 0.000000] x86/fpu: Supporting XSAVE feature 0x001: 'x87 floating point registers'
[ 0.000000] x86/fpu: Supporting XSAVE feature 0x002: 'SSE registers'
[ 0.000000] x86/fpu: Supporting XSAVE feature 0x004: 'AVX registers'
[ 0.000000] x86/fpu: xstate_offset[2]: 576, xstate_sizes[2]: 256
[ 0.000000] x86/fpu: Enabled xstate features 0x7, context size is 832 bytes, using 'standard' format.
[ 0.000000] x86/fpu: Using 'eager' FPU context switches.
[ 0.000000] e820: BIOS-provided physical RAM map:
[ 0.000000] BIOS-e820: [mem 0x0000000000000000-0x0000000000009dfff] usable
[ 0.000000] BIOS-e820: [mem 0x0000000000009e000-0x0000000000009ffff] reserved
[ 0.000000] BIOS-e820: [mem 0x000000000000e0000-0x000000000000fffff] reserved
[ 0.000000] BIOS-e820: [mem 0x00000000000100000-0x000000000003fffff] usable
[ 0.000000] BIOS-e820: [mem 0x00000000000fc00000-0x00000000000fffff] reserved
[ 0.000000] NX (Execute Disable) protection: active
[ 0.000000] SMBIOS 2.7 present.
[ 0.000000] Hypervisor detected: Xen
[ 0.000000] Xen version 4.2.
[ 0.000000] Netfront and the Xen platform PCI driver have been compiled for this kernel: unplug
emulated NICs.
[ 0.000000] Blkfront and the Xen platform PCI driver have been compiled for this kernel: unplug
emulated disks.
[ 0.000000] You might have to change the root device
[ 0.000000] from /dev/hd[a-d] to /dev/xvd[a-d]
[ 0.000000] in your root= kernel command line option
[ 0.000000] e820: last_pfn = 0x40000 max_arch_pfn = 0x400000000
[ 0.000000] x86/PAT: Configuration [0-7]: WB WC UC- UC WB WC UC- WT
[ 0.000000] found SMP MP-table at [mem 0x000fbc20-0x000fbc2f] mapped at [ffff88000
[ 0.000000] RAMDISK: [mem 0x371e1000-0x37feffff]
```



Per-instance metadata: Letting the instance know what it is

HTTP API:

```
$ curl http://169.254.169.254/latest/meta-data/instance-id
i-5c306beb$

$ aws_instance=$(wget -q -O- http://169.254.169.254/latest/meta-data/instance-id)
$ aws_region=$(wget -q -O- http://169.254.169.254/latest/meta-data/hostname)
$ echo $aws_instance $aws_region
i-5c306beb ip-172-30-1-33.ec2.internal
$
```

ec2-metadata:

```
$ ec2-metadata -i
instance-id: i-5c306beb
$ ec2-metadata -i | awk '{print $2;}'
i-5c306beb
```

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html>

Instance devices show up in the console

2 block devices:

Instance: i-08f7f8534c056d39b		Public DNS: ec2-54-165-59-97.compute-1.amazonaws.com	
Description	Status Checks	Monitoring	Tags
Instance ID	i-08f7f8534c056d39b		
Instance state	running		
Instance type	c3.large		
Elastic IPs			
Availability zone	us-east-1a		
Security groups	launch-wizard-4 . view inbound rules		
Scheduled events	No scheduled events		
AMI ID	amzn-ami-hvm-2017.09.1.20180115-x86_64-gp2 (ami-97785bed)		
Platform	-		
IAM role	-		
Key pair name	CFRS780		
EBS-optimized	False		
Root device type	ebs		
Root device	/dev/xvda		
Block devices	/dev/xvda /dev/sdc		

Here's what it looks like

```
[ec2-user@ip-172-31-30-242 ~]$ df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        1.9G   72K   1.9G   1% /dev
tmpfs           1.9G    0   1.9G   0% /dev/shm
/dev/xvda1       7.8G  1.1G   6.7G  14% /
/dev/xvdb        15G   39M   15G   1% /media/ephemeral0
[ec2-user@ip-172-31-30-242 ~]$
```


Redo the benchmark

```
$ sudo time dd if=/dev/zero of=/bigfile bs=65536 count=16384 conv=fdatasync
16384+0 records in
16384+0 records out
1073741824 bytes (1.1 GB) copied, 16.1457 s, 66.5 MB/s
0.00user 0.72system 0:16.17elapsed 4%CPU (0avgtext+0avgdata 2184maxresident)k
8inputs+2097152outputs (0major+100minor)pagefaults 0swaps
$
```

```
$ sudo time dd if=/dev/zero of=/media/ephemeral0/bigfile bs=65536 count=16384
conv=fdatasync
16384+0 records in
16384+0 records out
1073741824 bytes (1.1 GB) copied, 11.6586 s, 92.1 MB/s
0.02user 1.06system 0:11.68elapsed 9%CPU (0avgtext+0avgdata 2192maxresident)k
64inputs+2097152outputs (0major+100minor)pagefaults 0swaps
$
```

Working with EBS volumes

EBS volumes are virtual disks. Each one has:

- Volume ID: vol-0490630760213a246
- Size
- Volume Type
- IOPS
- Snapshot it was created from
- Created Time
- Availability Zone
- State
- Alarm Status
- Attachment Information — the EC2 instance it's attached to
- Volume Status
- Encryption Status

Actions for EBS Volumes

Modify Volume —

- Migrate to a different storage
- Make it bigger or smaller.
- Coordinate with OS!

The screenshot displays the AWS Management Console interface for EBS volumes. At the top, there is a 'Create Volume' button and an 'Actions' dropdown menu. The 'Actions' menu is open, showing options: 'Modify Volume', 'Create Snapshot', 'Delete Volume', 'Attach Volume', 'Detach Volume', 'Force Detach Volume', 'Change Auto-Enable IO Setting', and 'Add/Edit Tags'. Below the menu is a table of EBS volumes. The table has columns for 'Name', 'Size', and 'Volume Type'. One volume is selected, and the 'Modify Volume' dialog is open.

Name	Size	Volume Type
	8 GiB	gp2
	8 GiB	gp2
	8 GiB	gp2
	40 GiB	gp2
vol-028aedee7d1484adc	8 GiB	gp2
		gp2

Modify Volume

Volume ID: vol-091efb4f63d7ab302

Volume Type: General Purpose SSD (GP2) ⓘ

Size: 8 (Min: 1 GiB, Max: 16384 GiB) ⓘ

Iops: 100 / 3000 (Baseline of 3 IOPS per GiB with a minimum of 100 IOPS, burstable to 3000 IOPS) ⓘ

Cancel Modify

Actions and terminology

From a live VM, you can:

- Snapshot — A copy of just the blocks
- Image — A bootable AMI (Amazon Machine Image)

Snapshots can be turned into an AMI by:

- Registering it with metadata
- Specifying the correct kernel

AMI includes:

- Disk image
- Metadata — architecture, kernel, AMI name, description, device mappings

EBS volume:

- Can start as a blank volume or as a copy of an image.
- Can be mounted on a device.

All of these commands, and more, can be run from the CLI

\$ aws ec2 help

336 commands!

accept-reserved-instances-exchange-quote
accept-vpc-peering-connection
allocate-address
allocate-hosts
assign-ipv6-addresses
assign-private-ip-addresses
associate-address
associate-dhcp-options
associate-iam-instance-profile
associate-route-table
associate-subnet-cidr-block
associate-vpc-cidr-block
attach-classic-link-vpc
attach-internet-gateway
attach-network-interface
attach-volume
attach-vpn-gateway
authorize-security-group-egress
authorize-security-group-ingress
bundle-instance
cancel-bundle-task
cancel-conversion-task
cancel-export-task
cancel-import-task
cancel-reserved-instances-listing
cancel-spot-fleet-requests
cancel-spot-instance-requests
confirm-product-instance
copy-image
copy-snapshot
create-customer-gateway
create-default-vpc
create-dhcp-options
create-egress-only-internet-gateway
create-flow-logs
create-fpga-image
create-image
create-instance-export-task
create-internet-gateway
create-key-pair
create-nat-gateway
create-network-acl
create-network-acl-entry
create-network-interface
create-network-interface-permission
create-placement-group
create-reserved-instances-listing
create-route
create-route-table
create-security-group

create-snapshot
create-spot-datafeed-subscription
create-subnet
create-tags
create-volume
create-vpc
create-vpc-endpoint
create-vpc-peering-connection
create-vpn-connection
create-vpn-connection-route
create-vpn-gateway
delete-customer-gateway
delete-dhcp-options
delete-egress-only-internet-gateway
delete-flow-logs
delete-internet-gateway
delete-key-pair
delete-nat-gateway
delete-network-acl
delete-network-acl-entry
delete-network-interface
delete-network-interface-permission
delete-placement-group
delete-route
delete-route-table
delete-security-group
delete-snapshot
delete-spot-datafeed-subscription
delete-subnet
delete-tags
delete-volume
delete-vpc
delete-vpc-endpoints
delete-vpc-peering-connection
delete-vpn-connection
delete-vpn-connection-route
delete-vpn-gateway
deregister-image
describe-account-attributes
describe-addresses
describe-availability-zones
describe-bundle-tasks
describe-classic-link-instances
describe-conversion-tasks
describe-customer-gateways
describe-dhcp-options
describe-egress-only-internet-gateways
describe-elastic-gpus
describe-export-tasks
describe-flow-logs

describe-fpga-images
describe-host-reservation-offerings
describe-host-reservations
describe-hosts
describe-iam-instance-profile-associations
describe-id-format
describe-identity-id-format
describe-image-attribute
describe-images
describe-import-image-tasks
describe-import-snapshot-tasks
describe-instance-attribute
describe-instance-status
describe-instances
describe-internet-gateways
describe-key-pairs
describe-moving-addresses
describe-nat-gateways
describe-network-acls
describe-network-interface-attribute
describe-network-interface-permissions
describe-network-interfaces
describe-placement-groups
describe-prefix-lists
describe-regions
describe-reserved-instances
describe-reserved-instances-listings
describe-reserved-instances-modifications
describe-reserved-instances-offerings
describe-route-tables
describe-scheduled-instance-availability
describe-scheduled-instances
describe-security-group-references
describe-security-groups
describe-snapshot-attribute
describe-snapshots
describe-spot-datafeed-subscription
describe-spot-fleet-instances
describe-spot-fleet-request-history
describe-spot-fleet-requests
describe-spot-instance-requests
describe-spot-price-history
describe-stale-security-groups
describe-subnets
describe-tags
describe-volume-attribute
describe-volume-status

describe-volumes
describe-volumes-modifications
describe-vpc-attribute
describe-vpc-classic-link
describe-vpc-classic-link-dns-support
describe-vpc-endpoint-services
describe-vpc-endpoints
describe-vpc-peering-connections
describe-vpcs
describe-vpn-connections
describe-vpn-gateways
detach-classic-link-vpc
detach-internet-gateway
detach-network-interface
detach-volume
detach-vpn-gateway
disable-vgw-route-propagation
disable-vpc-classic-link
disable-vpc-classic-link-dns-support
disassociate-address
disassociate-iam-instance-profile
disassociate-route-table
disassociate-subnet-cidr-block
disassociate-vpc-cidr-block
enable-vgw-route-propagation
enable-volume-io
enable-vpc-classic-link
enable-vpc-classic-link-dns-support
get-console-output
get-console-screenshot
get-host-reservation-purchase-preview
get-password-data
get-reserved-instances-exchange-quote
help
import-image
import-key-pair
import-snapshot
modify-hosts
modify-id-format
modify-identity-id-format
modify-image-attribute
modify-instance-attribute
modify-instance-placement
modify-network-interface-attribute
modify-reserved-instances
modify-snapshot-attribute
modify-spot-fleet-request
modify-subnet-attribute
modify-volume

modify-volume-attribute
modify-vpc-attribute
modify-vpc-endpoint
modify-vpc-peering-connection-options
monitor-instances
move-address-to-vpc
purchase-host-reservation
purchase-reserved-instances-offering
purchase-scheduled-instances
reboot-instances
register-image
reject-vpc-peering-connection
release-address
release-hosts
replace-iam-instance-profile-association
replace-network-acl-association
replace-network-acl-entry
replace-route
replace-route-table-association
report-instance-status
request-spot-fleet
request-spot-instances
reset-image-attribute
reset-instance-attribute
reset-network-interface-attribute
reset-snapshot-attribute
restore-address-to-classic
revoke-security-group-egress
revoke-security-group-ingress
run-instances
run-scheduled-instances
start-instances
stop-instances
terminate-instances
unassign-ipv6-addresses
unassign-private-ip-addresses
unmonitor-instances
wait



AWS EBS


EBS: Virtual Disk Volumes

EBS volumes:

- Created automatically when EC2 instance starts up.
- Snapshots on the fly.

Options:

- Magnetic or SSD
- Destroy or persist on instance termination
- Not Encrypted / Encrypted
- Provisioned IOPS



Uses:

- Boot drives
- Read-only drives to share static databases. (Make 1 TB drive and mount)
- Database drives for MySQL, etc.. (But you should use Amazon's managed service.)

EBS offers three classes of service.

Characteristic	General Purpose (SSD)	Provisioned IOPS (SSD)	Magnetic
Use cases	<ul style="list-style-type: none">• System boot volumes• Virtual desktops• Small to medium sized databases• Development and test environments	<ul style="list-style-type: none">• Critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume• Large database workloads, such	<ul style="list-style-type: none">• Cold workloads where data is infrequently accessed• Scenarios where the lowest storage cost is important
Volume size	1 GiB – 16 TiB	4 GiB – 16 TiB	1 GiB – 1 TiB
Maximum throughput	160 MiB/s	320 MiB/s	40–90 MiB/s
IOPS performance	Baseline performance of 3 IOPS/GiB (up to 10,000 IOPS) with the ability to burst to 3,000	Consistently performs at provisioned level, up to 20,000 IOPS maximum	Averages 100 IOPS, with the ability to burst to hundreds of IOPS
API and CLI volume	gp2	io1	standard

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>

Pricing — you are probably best off with SSD General Purpose.

Amazon EBS Pricing

With Amazon EBS, you only pay for what you use. The pricing for Amazon EBS volumes is listed below.

Region:

Amazon EBS General Purpose (SSD) volumes

- \$0.10 per GB-month of provisioned storage

Amazon EBS Provisioned IOPS (SSD) volumes

- \$0.125 per GB-month of provisioned storage
- \$0.065 per provisioned IOPS-month

Amazon EBS Magnetic volumes

- \$0.05 per GB-month of provisioned storage
- \$0.05 per 1 million I/O requests

Amazon EBS Snapshots to Amazon S3

- \$0.095 per GB-month of data stored

EBS volumes can be created and used for: extra storage, sharing data

Each EBS volume has:

- Size e.g. 40GB
- Name e.g. vol-65202e2d
- Region / AvailabilityZone e.g. us-east-1 / us-east-1b
- Attributes e.g. CreateTime, Encrypted, Iops,

Volumes can be mounted:

- read/write on a single instance
- read-only on multiple instances

Create and share an instance:

```
$ aws ec2 create-volume --size 10 --availability-zone us-east-1a
```

You must specify a region. You can also configure your region by running "aws configure".

```
$ aws ec2 create-volume --size 10 --region us-east-1 --availability-zone us-east-1b
```

```
{
  "AvailabilityZone": "us-east-1b",
  "Encrypted": false,
  "VolumeType": "standard",
  "VolumeId": "vol-95cab176",
  "State": "creating",
  "SnapshotId": "",
  "CreateTime": "2015-12-05T18:55:28.052Z",
  "Size": 10
}
```

```
$
```

Attach the EBS volume to your VM

(Be sure EBS is in same region & availability zone)

First get a volume...

```
$ aws_zone=$(curl -s http://169.254.169.254/latest/meta-data/placement/availability-zone)
$ aws_instance=$(curl -s http://169.254.169.254/latest/meta-data/instance-id)
$ aws_region=$(curl -s http://169.254.169.254/latest/dynamic/instance-identity/document | grep
region|awk -F\" '{print $4}')
$ aws ec2 create-volume --size 10 --region $aws_region --availability-zone $aws_zone
{
  "AvailabilityZone": "us-east-1b",
  "Encrypted": false,
  "VolumeType": "standard",
  "VolumeId": "vol-46cdb6a5",
  "State": "creating",
  "SnapshotId": "",
  "CreateTime": "2015-12-05T19:01:38.548Z",
  "Size": 10
}
$ aws ec2 attach-volume --volume-id=vol-46cdb6a5 --instance-id=$aws_instance \
--device=/dev/sdb --region=$aws_region
{
  "AttachTime": "2015-12-05T19:02:11.541Z",
  "InstanceId": "i-5c306beb",
  "VolumeId": "vol-46cdb6a5",
  "State": "attaching",
  "Device": "/dev/sdb"
}
$
```

Now we need to make a file system...

Create a file system on the volume

```
$ sudo mkfs -t ext4 /dev/sdb
mke2fs 1.42.12 (29-Aug-2014)
Creating filesystem with 2621440 4k blocks and 655360 inodes
Filesystem UUID: 681c57f0-1461-4dae-b956-032656ba82a9
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632
```

```
Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

```
$ sudo mount /dev/sdb /mnt/extra/
[ip-172-30-1-33 ~ 19:04:44]$ df
Filesystem      1K-blocks      Used Available Use% Mounted on
/dev/xvda1      41151788 6506728  34544812  16% /
devtmpfs         500712         60    500652    1% /dev
tmpfs            509724          0    509724    0% /dev/shm
/dev/xvdb       10190136    23028   9626436    1% /mnt/extra
```

```
$ lsblk
NAME      MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
xvda      202:0    0   40G  0 disk
└─xvda1   202:1    0   40G  0 part /
xvdb      202:16   0   10G  0 disk /mnt/extra
$
```

AWS CloudTrail



U.S. Department of Commerce
Economics and Statistics Administration
U.S. CENSUS BUREAU
[census.gov](https://www.census.gov)

S3 Management Console

CloudTrail Management Console

GMU

Secure

https://console.aws.amazon.com/cloudtrail/home?region=us-east-1#/configuration

Apps

CFRS780

Discussion Board

Schedule

AWS

EC2 Management C...

EC2 Instance Pricing...

\$\$

GMU

aws

Services

Resource Groups

EC2

EMR

sgarfin2@gmu.edu

N. Virginia

Support

CloudTrail

Dashboard

Event history

Trails

Trails

Deliver logs to an Amazon S3 bucket. CloudTrail events can be processed by one trail for free. There is a charge for processing events with additional trails. For more information, see [AWS CloudTrail Pricing](#).

Create trail

Name	Region	S3 bucket	Log file prefix	CloudWatch Logs Log group	Status
No trails have been added					

Learn more

[Pricing](#)
[Documentation](#)
[Forums](#)
[FAQs](#)

Feedback

English (US)

© 2008 - 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Privacy Policy

Terms of Use

CloudTrail log validation

The screenshot shows the AWS CloudTrail console configuration page for a new trail. The browser tabs at the top include 'S3 Management Console' and 'CloudTrail Management Console'. The address bar shows the URL: <https://console.aws.amazon.com/cloudtrail/home?region=us-east-1#/configuration/new>. The AWS navigation bar at the top includes the AWS logo, 'Services', 'Resource Groups', and icons for EC2, EMR, and other services. The user's email 'sgarfin2@gmu.edu' and the region 'N. Virginia' are also visible.

On the left sidebar, the 'CloudTrail' section is expanded, showing 'Dashboard', 'Event history', and 'Trails'. The 'Trails' section is currently selected.

The main configuration area includes the following options:

- Create a new S3 bucket:** ☒ Yes ☐ No
- S3 bucket*:**
- Log file prefix:**
Location: /AWSLogs/309467262965/CloudTrail/us-east-1
- Encrypt log files:** ☐ Yes ☒ No
- Enable log file validation:** ☒ Yes ☐ No
- Send SNS notification for every log file delivery:** ☐ Yes ☒ No

A tooltip is displayed over the 'Enable log file validation' option, containing the text: "To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file validation. [Learn more.](#)"

At the bottom of the configuration area, there is a note: "* Required field". To the right, a partial note reads: "Additional charges may ap". A blue 'Create' button is visible at the bottom right.

Below the configuration area, there is a 'Learn more' link and a 'Pricing' link.

S3 Management Console
CloudTrail Management Console
Validating CloudTrail Log File Integrity
GMU

Secure
https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html?icmpid=docs_cloudtrail...

Apps
CFRS780
Discussion Board
Schedule
AWS
EC2 Management C...
EC2 Instance Pricing...
\$\$
GMU

Menu
aws
English
Sign In to the Console

AWS CloudTrail

User Guide (Version 1.0)

Documentation - This Guide

Search

- + What Is AWS CloudTrail?
- + Getting Started with CloudTrail
- Working with CloudTrail Log Files
 - Create Multiple Trails
 - Logging Data and Management Events for Trails
 - Receiving CloudTrail Log Files from Multiple Regions
- + Monitoring CloudTrail Log Files with Amazon CloudWatch Logs
- + Receiving CloudTrail Log Files from Multiple Accounts
- + Sharing CloudTrail Log Files Between AWS Accounts
- + Encrypting CloudTrail Log Files with AWS KMS-Managed Keys (SSE-KMS)

[AWS Documentation](#) » [AWS CloudTrail](#) » [User Guide](#) » [Working with CloudTrail Log Files](#) » Validating CloudTrail Log File Integrity

Validating CloudTrail Log File Integrity

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the AWS CLI to validate the files in the location where CloudTrail delivered them.

Why Use It?

Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.

How It Works

When you enable log file integrity validation, CloudTrail creates a hash for every log file that it delivers. Every hour, CloudTrail also creates and delivers a file that references the log files

Terms of Use | © 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

S3 Management Console

CloudTrail Management

Amazon SNS Topic Poll

CloudTrail Trail Naming

AWS CloudTrail | Pricing

GMU

Secure

https://console.aws.amazon.com/cloudtrail/home?region=us-east-1#/configuration

Apps

CFRS780

Discussion Board

Schedule

AWS

EC2 Management C...

EC2 Instance Pricing...

\$\$

GMU

aws

Services

Resource Groups

EC2

EMR

sgarfin2@gmu.edu

N. Virginia

Support

CloudTrail

Dashboard

Event history

Trails

Trails

Deliver logs to an Amazon S3 bucket. CloudTrail events can be processed by one trail for free. There is a charge for processing events with additional trails. For more information, see [AWS CloudTrail Pricing](#).

Create trail

Name	Region	S3 bucket	Log file prefix	CloudWatch Logs Log group	Status
CRFS780_Trails	All	crfs780-trails			✓

Learn more

[Pricing](#)
[Documentation](#)
[Forums](#)
[FAQs](#)

Feedback

English (US)

© 2008 - 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Privacy Policy

Terms of Use

Cloud Trails Works!

Note spelling error!

CloudTrail — the raw files

CloudTrail-Digest — metadata about the raw files

```
[hadoop@ip-172-31-49-78 ~]$ aws s3 ls crfs780-trails/AWSLogs/309467262965/
PRE CloudTrail-Digest/
PRE CloudTrail/
2018-02-19 00:56:16      0
[hadoop@ip-172-31-49-78 ~]$ aws s3 ls crfs780-trails/AWSLogs/309467262965/CloudTrail/
PRE ap-northeast-1/
PRE ap-northeast-2/
PRE ap-northeast-3/
PRE ap-south-1/
PRE ap-southeast-1/
PRE ap-southeast-2/
PRE ca-central-1/
PRE eu-central-1/
PRE eu-west-1/
PRE eu-west-2/
PRE eu-west-3/
PRE sa-east-1/
PRE us-east-1/
PRE us-east-2/
PRE us-west-1/
PRE us-west-2/
```

Cloud Trails Works!

CloudTrail — the raw files

CloudTrail-Digest — metadata about the raw files

aws s3 ls crfs780-trails/AWSLogs/309467262965/CloudTrail-Digest/us-east-1/2018/02/19/

```
[hadoop@ip-172-31-49-78 ~]$ aws s3 ls crfs780-trails/AWSLogs/309467262965/CloudTrail-Digest/us-east-1/2018/02/19/
2018-02-19 01:47:32      342 309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T005615Z.json.gz
2018-02-19 02:47:38     4037 309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T015615Z.json.gz
2018-02-19 03:47:23     4192 309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T025615Z.json.gz
2018-02-19 04:47:26     4124 309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T035615Z.json.gz
2018-02-19 05:47:34     4114 309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T045615Z.json.gz
2018-02-19 06:47:30     4128 309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T055615Z.json.gz
2018-02-19 07:47:27     4186 309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T065615Z.json.gz
2018-02-19 08:47:31     4122 309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T075615Z.json.gz
2018-02-19 09:47:36     4202 309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T085615Z.json.gz
...
```

[hadoop@ip-172-31-49-78 ~]\$ aws s3 cp s3://crfs780-trails/AWSLogs/309467262965/CloudTrail-Digest/us-east-1/2018/02/19/309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T015615Z.json.gz trails.json.gz

download: s3://crfs780-trails/AWSLogs/309467262965/CloudTrail-Digest/us-east-1/2018/02/19/309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T015615Z.json.gz to ./trails.json.gz


```
[hadoop@ip-172-31-49-78 ~]$ ls -l trails.json.gz
-rw-rw-r-- 1 hadoop hadoop 4037 Feb 19 02:47 trails.json.gz
[hadoop@ip-172-31-49-78 ~]$ gunzip trails.json.gz
[hadoop@ip-172-31-49-78 ~]$ ls -l trails*
-rw-rw-r-- 1 hadoop hadoop 17833 Feb 19 02:47 trails.json
[hadoop@ip-172-31-49-78 ~]$ more trails.json
```

...

```
{"awsAccountId":"309467262965","digestStartTime":"2018-02-19T00:56:15Z","digestEndTime":"2018-02-19T01:56:15Z","digestS3Bucket":"crfs780-trails","digestS3Object":"AWSLogs/309467262965/CloudTrail-Digest/us-east-1/2018/02/19/309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T015615Z.json.gz","digestPublicKeyFingerprint":"3fc64187ed954b65bd053279acc75d70","digestSignatureAlgorithm":"SHA256withRSA","newestEventTime":"2018-02-19T01:52:22Z","oldestEventTime":"2018-02-19T00:44:47Z","previousDigestS3Bucket":"crfs780-trails","previousDigestS3Object":"AWSLogs/309467262965/CloudTrail-Digest/us-east-1/2018/02/19/309467262965_CloudTrail-Digest_us-east-1_CRFS780_Trails_us-east-1_20180219T005615Z.json.gz","previousDigestHashValue":"cc6026c8656ab9d6f295bfaed9d2760b969487a8b8148704f27455e344ed27f4","previousDigestHashAlgorithm":"SHA-256","previousDigestSignature":"76f4d42d0b12bd9f4fad874c95adaaf6ba05f66c6adc8f9ed7c1d22fa827b8c8de516b112bca0bad76d9833bcd3e040be64c321a25bb765505ddecafb0b132db53edc4e122eed34c4932fb4441b86589dec382d2356157d457e9d3a5114d53761f629bb9e4487e764d3121dfcdb1bb79d274b497161b9a44143560704249d75dd57150b42e4d757602c92637aa704f822c899acf4f47f9e1166ae0692590087eb2ea3aaec97e6e937efde434e36739f37a2eac372e8b83d37f354cb7953d54d85780ea5aabd430c6d4f132edf83e9525644fe595f77ca6987f491da7524a8f5a66779bab7e13c9cf5494e2d064105c592cf2aa332dcabc734ba5b8a8c49a0fd4","logFiles":[{"s3Bucket":"crfs780-trails","s3Object":"AWSLogs/309467262965/CloudTrail/us-east-1/2018/02/19/30946726
```

If you don't have a JSON viewer, you can use Python...

```
[hadoop@ip-172-31-49-78 ~]$ python3.6
Python 3.6.2 (default, Nov  2 2017, 19:34:31)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-11)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import json
>>> data = json.loads(open("trails.json","r").read())
>>> data.keys()
dict_keys(['awsAccountId', 'digestStartTime', 'digestEndTime', 'digestS3Bucket',
'digestS3Object', 'digestPublicKeyFingerprint', 'digestSignatureAlgorithm',
'newestEventTime', 'oldestEventTime', 'previousDigestS3Bucket',
'previousDigestS3Object', 'previousDigestHashValue',
'previousDigestHashAlgorithm', 'previousDigestSignature', 'logFiles'])
>>> len(data["logFiles"])
46
>>> data["logFiles"][1]
{'s3Bucket': 'crfs780-trails', 's3Object': 'AWSLogs/309467262965/CloudTrail/us-
east-1/2018/02/19/309467262965_CloudTrail_us-
east-1_20180219T0135Z_8bl6vp5CfHoaB2Lh.json.gz', 'hashValue':
'f4151aa6d39999c32e6a0be3adc9807839861291f75ac30d4a376698caef5230',
'hashAlgorithm': 'SHA-256', 'newestEventTime': '2018-02-19T01:30:35Z',
'oldestEventTime': '2018-02-19T01:28:39Z'}
```

—

We can get that file...

```
$ aws s3 cp s3://crfs780-trails/AWSLogs/309467262965/CloudTrail/us-east-1/2018/02/19/309467262965_CloudTrail_us-east-1_20180219T0135Z_8bl6vp5CfHoaB2Lh.json.gz data2.json.gz
download: s3://crfs780-trails/AWSLogs/309467262965/CloudTrail/us-east-1/2018/02/19/309467262965_CloudTrail_us-east-1_20180219T0135Z_8bl6vp5CfHoaB2Lh.json.gz to ./records.json.gz
$ gunzip records.json.gz
```

```
hadoop@ip-172-31-49-78 ~]$ python3
Python 3.6.2 (default, Nov 2 2017, 19:34:31)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-11)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import json
>>> records = json.loads(open("records.json").read())
>>> len(records['Records'])
11
>>>
```

Here is one record (keys in bold)

```
>>> records['Records'][0]
{'eventVersion': '1.05', 'userIdentity': {'type': 'AWSService', 'invokedBy':
'elasticmapreduce.amazonaws.com'}, 'eventTime': '2018-02-19T01:29:36Z', 'eventSource':
'sts.amazonaws.com', 'eventName': 'AssumeRole', 'awsRegion': 'us-east-1',
'sourceIPAddress': 'elasticmapreduce.amazonaws.com', 'userAgent':
'elasticmapreduce.amazonaws.com', 'requestParameters': {'roleArn':
'arn:aws:iam::309467262965:role/EMR_DefaultRole', 'roleSessionName': 'CCSSession',
'durationSeconds': 1500}, 'responseElements': {'credentials': {'accessKeyId':
'ASIAJPGNR7QOAKKOM5HQ', 'expiration': 'Feb 19, 2018 1:54:36 AM', 'sessionToken':
'AgoGb3JpZ2luEA0aCXVzLWVhc3QtMSKAAjI5KkvkNsX7M9SG1nHC/
yT6DFQDLZZ8Ek9dty2WpgQbpgPmn6TE+VLEUz1ftZPJxrHL0OH5RLm27G9xIxQXpRrrwInleeZ50yw62cp9BO8ex
TwNViYCiSVWN3AnqQqlmHwKU+Wdd/mVX+c4Uu7EKYY1ijLB6pFUEcW6rHxWCWqQLyJ/
OzIgbZqHpHd7nUyaCMOz1inte450FKew/eSfvI2LIyE0OM/OSh/X3p2/dwa5rIDwuSdxnN/
aTyTHiLTrMF54J+Z7R4zA+qV5KSeHHjS/So4MJoELkWIO1MjpxqMd4GADIkUV/
whzREuTjqvxsQ1l8tVvm90coRc1LL1cd7wqhgmI4//////////
ARAAGgwzMDk0NjcyNjI5NjUiDHfH08I1YiqjkzXwOyraAu/Xisb8QiIZg13/JG81uLwS1tEw/
QFuqtPNQOpJ7s0z6m3DMubseO0c6E6HEjThNf9VMRjrNsQEH2JuMPG9TzDgjFivoc9X2QEH3SWrK68E7waPTH+35
6nqIueXWiLiatJiWCEUS526URB78kQyvC2KdYaTDnmtqml4/3fM98a+eV3iRBw9Grz1+RGvZ2OZ9Sm+RgFOtMtLc
MYEIH05qn6odJpid4wIF+qdtkl/Qp/
Uny1qf1FJJ0E80oemhcwJFr23R0IBDNW9c71NK+MEYwZE5XaXSF3J7lg9IH74X1w05dScPEipiKUJKTTXoxbUk5v
USbPWM9FWeKYtbGHRovOkjaDRqgD5y5ehiF2qF8mDBXLZ2FlW0HaktfC1wnZ6GKznVGyCTzJsa2R3/EG/
dqx6YaZXRo0utlijlMPtAg+azHhhVJBjGiMtL9PKFXwCN3jh7ENwYlyJZvIwgNGolAU='}},
'assumedRoleUser': {'assumedRoleId': 'AROAILDJTGTBO4WJMRFDS:CCSSession', 'arn':
'arn:aws:sts::309467262965:assumed-role/EMR_DefaultRole/CCSSession'}}, 'requestID':
'58b0a175-1514-11e8-9166-8b01e23a16a9', 'eventID': '6bc67694-fef9-4f49-85e3-
efdbbcb0d855', 'resources': [{'ARN': 'arn:aws:iam::309467262965:role/EMR_DefaultRole',
'accountId': '309467262965', 'type': 'AWS::IAM::Role'}], 'eventType': 'AwsApiCall',
'recipientAccountId': '309467262965', 'sharedEventID': '6921e72e-aadc-4336-
a822-928ea82f2c3a'}
```


Working with CloudTrails

—If you can avoid it, don't write your own analysis tools

Amazon partners have tools for processing CloudTrails:

- <https://aws.amazon.com/cloudtrail/partners/>

Amazon Athena is an analysis platform provided by Amazon:

- <https://aws.amazon.com/blogs/big-data/aws-cloudtrail-and-amazon-athena-dive-deep-to-analyze-security-compliance-and-operational-activity/>


Amazon Glue and QuickSight also provide tools:

- <https://aws.amazon.com/blogs/big-data/streamline-aws-cloudtrail-log-visualization-using-aws-glue-and-amazon-quicksight/>

Multiple Amazon Blog entries:

- <https://aws.amazon.com/blogs/mt/category/management-tools/aws-cloudtrail/>

aws.amazon.com/athena/

Menu  [Contact Sales](#) [Products](#) [Solutions](#) [More](#) [English](#) [My Account](#) [Sign In to the Console](#)

Amazon Athena

Start querying data instantly. Get results in seconds. Pay only for the queries you run.

[Get Started with Amazon Athena](#)

[Product Details](#) [Pricing](#) [Getting Started](#) [FAQs](#) [Documentation](#) [Big Data Blog](#)

Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.

Athena is easy to use. Simply point to your data in Amazon S3, define the schema, and start querying using standard SQL. Most results are delivered within seconds. With Athena, there's no need for complex ETL jobs to prepare your data for analysis. This makes it easy for anyone with SQL skills to quickly analyze large-scale datasets.

[Display a menu](#)

Amazon Athena looks like SQL...

The screenshot shows the Amazon Athena User Guide page in a web browser. The browser's address bar displays the URL `docs.aws.amazon.com/athena/latest/ug/cloudtrail-logs.html`. The page has a sidebar on the left with the 'Amazon Athena' logo and a search bar. Below the search bar, there is a list of navigation links: 'What is Amazon Athena?', 'Release Notes', 'Setting Up', 'Getting Started', 'Integration with AWS Glue', 'Connecting to Amazon Athena with ODBC and JDBC Drivers', 'Security', 'Working with Source Data', 'Querying Data in Amazon Athena Tables', 'Querying Geospatial Data', and 'Querying AWS Service Logs'. The main content area on the right contains a paragraph stating: 'of these fields, use `JSON_EXTRACT` functions. For more information, see [Extracting Data From JSON](#).' Below this paragraph is a code block containing SQL syntax for creating an external table named `cloudtrail_logs`. The code defines fields such as `eventversion` (STRING), `userIdentity` (STRUCT with fields like `type`, `principalid`, `arn`, `accountid`, `invokedby`, `accesskeyid`, and `username`), `sessioncontext` (STRUCT with `attributes`), and `sessionIssuer` (STRUCT). It also lists other fields like `eventTime`, `eventSource`, `eventName`, `awsRegion`, `sourceIpAddress`, `userAgent`, `errorCode`, `errorMessage`, `requestParameters`, and `responseElements`, all of which are of type `STRING`. At the bottom of the page, there is a footer with the text 'Use | © 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.' and a feedback section asking 'Did this page help you?' with 'Yes' and 'No' buttons, and a 'Feedback' button.

docs.aws.amazon.com/athena/latest/ug/cloudtrail-logs.html

Menu

aws

English

Sign In to the Console

Amazon Athena

User Guide

Documentation - This Guide

Search

- What is Amazon Athena?
- Release Notes
- Setting Up
- Getting Started
- Integration with AWS Glue
- Connecting to Amazon Athena with ODBC and JDBC Drivers
- Security
- Working with Source Data
- Querying Data in Amazon Athena Tables
- Querying Geospatial Data
- Querying AWS Service Logs
- Querying AWS CloudTrail

of these fields, use `JSON_EXTRACT` functions. For more information, see [Extracting Data From JSON](#).

```
CREATE EXTERNAL TABLE cloudtrail_logs (  
  eventversion STRING,  
  userIdentity STRUCT<  
    type:STRING,  
    principalid:STRING,  
    arn:STRING,  
    accountid:STRING,  
    invokedby:STRING,  
    accesskeyid:STRING,  
    username:STRING,  
  sessioncontext:STRUCT<  
    attributes:STRUCT<  
      mfaauthenticated:STRING,  
      creationdate:STRING>,  
    sessionIssuer:STRUCT<  
      type:STRING,  
      principalId:STRING,  
      arn:STRING,  
      accountId:STRING,  
      userName:STRING>>>,  
  eventTime STRING,  
  eventSource STRING,  
  eventName STRING,  
  awsRegion STRING,  
  sourceIpAddress STRING,  
  userAgent STRING,  
  errorCode STRING,  
  errorMessage STRING,  
  requestParameters STRING,  
  responseElements STRING,
```

Use | © 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Did this page help you? Yes No Feedback

Once you have a the tables created, you can SQL away!

```
SELECT
  useridentity.arn,
  eventname,
  sourceipaddress,
  eventtime
FROM cloudtrail_logs
LIMIT 100;
```

Pricing:

- \$5 per TB of data scanned
- Can scan CloudTrails, text files, Parquet files. and more.

Other options: AWS Glue & Quicksight (and splunk, too!)

The screenshot shows a web browser window displaying an AWS Big Data Blog article. The browser's address bar shows the URL: <https://aws.amazon.com/blogs/big-data/streamline-aws-cloudtrail-log-visi>. The page header includes the AWS logo, navigation links (Menu, Products, Solutions, Pricing, Getting Started, More), and a 'Sign In to the Console' button. Below the header, the article title 'Visualize AWS Cloudtrail Logs using AWS Glue and Amazon Quicksight' is prominently displayed, followed by the author 'Luis Caro Perez' and the date '10 NOV 2017'. The article text explains the benefits of visualizing AWS CloudTrail logs and outlines a solution using AWS Glue and Amazon Quicksight. A 'Solution overview' section includes a diagram of the architecture. The diagram shows a flow from 'S3 log files Several Folders json.gz format' to 'Lambda Function', then to 'S3 files log files Single Folder json.gz Format', then to 'AWS GLUE DATA CATALOG Crawler', then to 'GLUE ETL', then to 'S3 files Parquet format compressed', then to 'AWS GLUE DATA CATALOG Crawler', then to 'Athena Query', and finally to 'Quick Sight Visualization'. The article text below the diagram explains that CloudTrail delivers log files in an Amazon S3 bucket folder, and that a Lambda function is used to transform the files into a single folder. It also mentions that AWS Glue scans the data, converts it into Apache Parquet format, and catalogs it for querying and visualization using Amazon Athena and Amazon Quicksight.

AWS Big Data Blog

Visualize AWS Cloudtrail Logs using AWS Glue and Amazon Quicksight

by Luis Caro Perez | on 10 NOV 2017 | in [Amazon Athena*](#), [Amazon QuickSight*](#), [AWS CloudTrail*](#), [AWS Glue*](#) | [Permalink](#) | [Comments](#) | [Share](#)

Being able to easily visualize [AWS CloudTrail](#) logs gives you a better understanding of how your AWS infrastructure is being used. It can also help you audit and review AWS API calls and detect security anomalies inside your AWS account. To do this, you must be able to perform analytics based on your CloudTrail logs.

In this post, I walk through using [AWS Glue](#) and [AWS Lambda](#) to convert AWS CloudTrail logs from JSON to a query-optimized format dataset in [Amazon S3](#). I then use [Amazon Athena](#) and [Amazon QuickSight](#) to query and visualize the data.

Solution overview

To process CloudTrail logs, you must implement the following architecture:

```
graph LR; A[S3 log files  
Several Folders  
json.gz format] --> B[Lambda Function]; B --> C[S3 files  
log files  
Single Folder  
json.gz Format]; C --> D[AWS GLUE  
DATA CATALOG  
Crawler]; D --> E[GLUE  
ETL]; E --> F[S3 files  
Parquet format  
compressed]; F --> G[AWS GLUE  
DATA CATALOG  
Crawler]; G --> H[Athena  
Query]; H --> I[Quick Sight  
Visualization]
```

CloudTrail delivers log files in an Amazon S3 bucket folder. To correctly crawl these logs, you modify the file contents and folder structure using an Amazon S3-triggered Lambda function that stores the transformed files in an S3 bucket single folder. When the files are in a single folder, AWS Glue scans the data, converts it into Apache Parquet format, and catalogs it to allow for querying and visualization using Amazon Athena and Amazon Quicksight.

QuickSight can log to CloudTrail...

The screenshot shows a web browser window displaying the AWS Big Data Blog. The URL in the address bar is aws.amazon.com/blogs/big-data/amazon-quicksight-now-supports-audit-logging-with-aws-cloudtrail/. The browser's tab bar shows several open tabs, including 'Amazon QuickSight Now Supports Audit Loggi...', 'Querying AWS CloudTrail Logs - Amazon Athena', 'AWS CloudTrail - Splunk', and 'New AWS Big Data Blog Post: Analyze Secur...'. The AWS navigation bar at the top includes links for Menu, Products, Solutions, Pricing, Getting Started, More, My Account, and a Sign In to the Console button. Below the navigation bar, there's a search bar and a 'Search Blogs' button. The main content area features a grid of related posts, including 'Build a social media dashboard using machine learning and BI services', 'Amazon QuickSight Update – Geospatial Visualization, Private VPC Access, and More', 'Audit Amazon Aurora Database Logs for Connections, Query Patterns, and More, using Amazon Athena and Amazon QuickSight', 'Visualize AWS Cloudtrail Logs using AWS Glue and Amazon QuickSight', 'Amazon QuickSight Adds Support for Combo Charts and Row-Level Security', and 'Use Amazon QuickSight Federated Single Sign-On with Amazon Cognito User Pools'. The main article is titled 'Amazon QuickSight Now Supports Audit Logging with AWS CloudTrail' by Jose Kunnackal, dated 28 APR 2017. The article text states: 'We launched Amazon QuickSight to democratize BI. Our goal is to make it easier and cheaper to roll out advanced business analytics capabilities to everyone in an organization. Overall, this enables better understanding of business, and allows faster data-driven decisions in an organization. In the past, the ability to share data presented an administrative challenge – that of knowing who has access to what data. Solving this problem ensures compliance with policies, and also provides an opportunity for businesses to see how employees use data to drive crucial decisions. Today, we are happy to announce support for AWS CloudTrail in Amazon QuickSight, which allows logging of QuickSight events across an AWS account. Whether you have an enterprise setting or a small team scenario, this integration will allow QuickSight administrators to accurately answer questions such as who last changed an analysis, or who has connected to sensitive data. With CloudTrail, administrators have better governance, auditing and risk management of their QuickSight usage. You can get started with CloudTrail with just a few clicks. Any AWS account that is enabled for CloudTrail will automatically see QuickSight activity included in the CloudTrail logs. When enabled, CloudTrail starts logging events including:' followed by a list of events: Account subscribe/unsubscribe, Data source create/update/delete, and Data source delete.

Display a menu

Amazon Simple Notification Service (SNS)

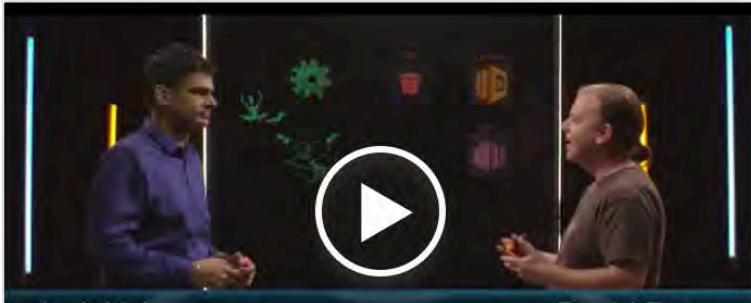
Pub/Sub messaging and mobile notifications for microservices, distributed systems, and serverless applications

Start today for free

Pub/Sub Messaging Mobile Notifications Product Details Getting Started FAQs Pricing

Amazon Simple Notification Service (SNS) is a flexible, fully managed [pub/sub messaging](#) and mobile notifications service for coordinating the delivery of messages to subscribing endpoints and clients. With SNS you can fan-out messages to a large number of subscribers, including distributed systems and services, and mobile devices. It is easy to set up, operate, and reliably send notifications to all your endpoints – at any scale. You can get

<https://aws.amazon.com/sns/getting-started/>



AWS SNS

Secure | https://us-west-2.console.aws.amazon.com/sns/v2/home?refid=ha_AWSSM-411®ion=us-west-2#/home

Apps Bb CFRS780 Bb Discussion Board -... Schedule AWS EC2 Management C... EC2 Instance Pricing... \$\$ GMU






aws Services Resource Groups EC2 EMR sgarfin2@gmu.edu Oregon Support

SNS dashboard

SNS dashboard

- Topics
- Applications
- Subscriptions
- Text messaging (SMS)

Common actions

-  **Create topic**
Create a communication channel to send messages and subscribe to notifications
-  **Create platform application**
Create a platform application for mobile devices
-  **Create subscription**
Subscribe an endpoint to a topic to receive messages published to that topic
-  **Publish message**
Publish a message to a topic or as a direct publish to a platform endpoint
-  **Publish text message (SMS)**
Publish a text message to a phone number

Resources

You are using the following Amazon SNS resources in the us-west-2 region:

Topic	0
Subscriptions	0
Applications	0
Endpoints	0

More info

- [Getting started](#)
- [Documentation](#)
- [API reference](#)
- [Forums](#)
- [Service health](#)

Feedback English (US) © 2008 - 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use

AWS SNS

Secure | https://us-west-2.console.aws.amazon.com/sns/v2/home?refid=ha_AWSSM-411®ion=us-west-2#/home

Apps Bb CFRS780 Discussion Board -... Schedule AWS EC2 Management C... EC2 Instance Pricing... \$\$ GMU

aws Services Resource Groups EC2 EMR sgarfin2@gmu.edu Oregon Support

SNS dashboard

Create new topic


A topic name will be used to create a permanent unique identifier called an Amazon Resource Name (ARN).

Topic name ⓘ

Display name ⓘ

[Cancel](#) [Create topic](#)

platform endpoint

 **Publish text message (SMS)**
Publish a text message to a phone number

[API reference](#)
[Forums](#)
[Service health](#)

[Feedback](#) [English \(US\)](#) © 2008 - 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved. [Privacy Policy](#) [Terms of Use](#)

AWS EFS

Options for storing forensic data results:

EBS — Elastic Block Store

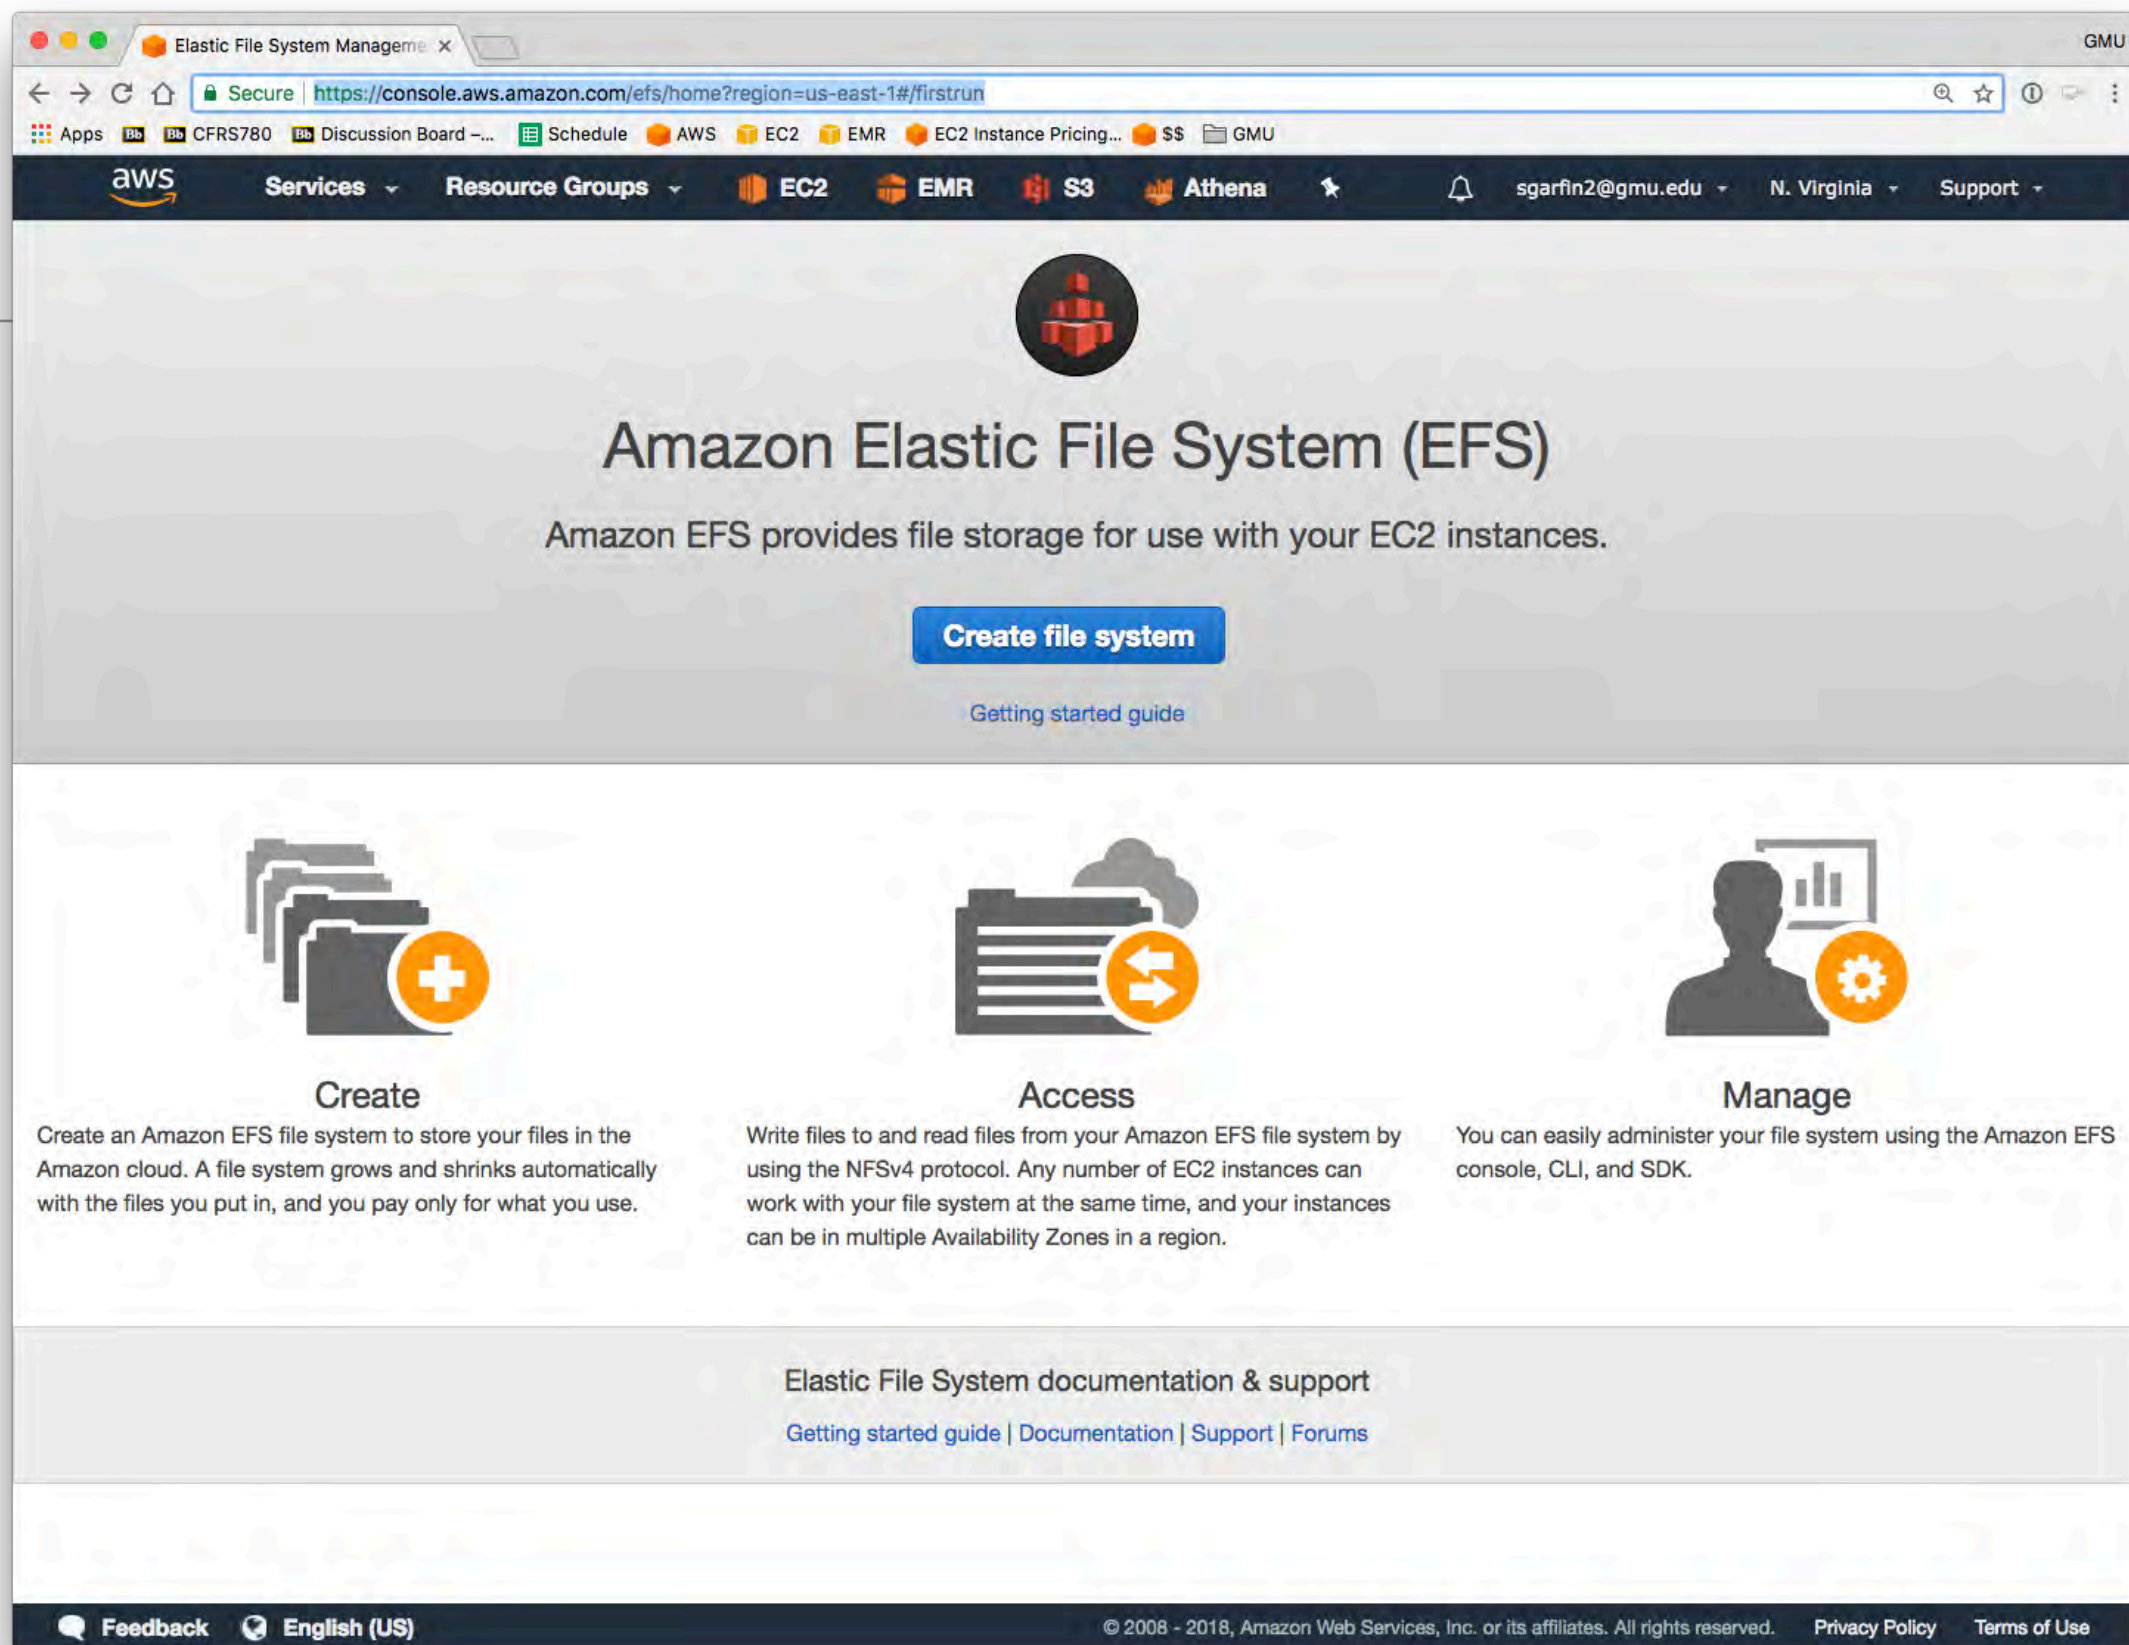
- Very fast
- Read/write on a single VM at a time
- Snapshot capability
- Restricted to an availability zone within a region

S3 — Simple Storage Service

- Stores objects, logfiles, etc.
- Integrates with Amazon Athena
- Doesn't have file permissions

EFS — Elastic File System

- A consistent read-write file system
- Can be mounted on many servers at once
- Works across availability zones
- NFS4 (Network File System) interface



Read the documentation:

- <https://docs.aws.amazon.com/efs/latest/ug/whatisefs.html>

Elastic File System Manage

Secure | https://console.aws.amazon.com/efs/home?region=us-east-1#/wizard/1

Apps | CFRS780 | Discussion Board | Schedule | AWS | EC2 | EMR | EC2 Instance Pricing | GMU

aws | Services | Resource Groups | EC2 | EMR | S3 | Athena | sgarfin2@gmu.edu | N. Virginia | Support

Create file system

Step 1: Configure file system access

Step 2: Configure optional settings

Step 3: Review and create

Configure file system access

An Amazon EFS file system is accessed by EC2 instances running inside one of your VPCs. Instances connect to a file system by using a network interface called a mount target. Each mount target has an IP address, which we assign automatically or you can specify.

VPC

vpc-1aaba062 (default)

Create mount targets

Instances connect to a file system by using mount targets you create. We recommend creating a mount target in each of your VPC's Availability Zones so that EC2 instances across your VPC can access the file system.

	Availability Zone	Subnet	IP address	Security groups
<input checked="" type="checkbox"/>	us-east-1a	subnet-e22613a9 (default)	Automatic	sg-29f4f65d - default
<input checked="" type="checkbox"/>	us-east-1b	subnet-e51676b8 (default)	Automatic	sg-29f4f65d - default
<input checked="" type="checkbox"/>	us-east-1c	subnet-92c892f6 (default)	Automatic	sg-29f4f65d - default
<input checked="" type="checkbox"/>	us-east-1d	subnet-90523ebf (default)	Automatic	sg-29f4f65d - default
<input checked="" type="checkbox"/>	us-east-1e	subnet-66324259 (default)	Automatic	sg-29f4f65d - default
<input checked="" type="checkbox"/>	us-east-1f	subnet-ec31cce3 (default)	Automatic	sg-29f4f65d - default

Cancel

Next Step

Feedback

English (US)

© 2008 - 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved. | Privacy Policy | Terms of Use

You must use security groups to restrict access to your access points.

Create mount targets

Instances connect to a file system by using mount targets you create. We recommend creating a mount target in each of your VPC's Availability Zones so that EC2 instances across your VPC can access the file system.

	Availability Zone	Subnet i	IP address i	Security groups i
<input checked="" type="checkbox"/>	us-east-1a	subnet-e22613a9 (default) ▾	Automatic	sg-29f4f65d - default ✕
<input checked="" type="checkbox"/>	us-east-1b	subnet-e51676b8 (default) ▾	Automatic	sg-29f4f65d - default ✕
<input checked="" type="checkbox"/>	us-east-1c	subnet-92c892f6 (default) ▾	Automatic	sg-29f4f65d - default ✕
<input checked="" type="checkbox"/>	us-east-1d	subnet-90523ebf (default) ▾	Automatic	sg-29f4f65d - default ✕
<input checked="" type="checkbox"/>	us-east-1e	subnet-66324259 (default) ▾	Automatic	sg-29f4f65d - default ✕
<input checked="" type="checkbox"/>	us-east-1f	subnet-ec31cce3 (default) ▾	Automatic	sg-29f4f65d - default ✕

Cancel

Next Step

I created a SLG Analysis security group

I manually added to it the IP address of each of my VMs.

The screenshot shows the AWS Management Console interface. At the top, there's a navigation bar with various services like EC2, EMR, S3, and Athena. Below this, the 'Create Security Group' page is visible. A table lists several security groups, with 'SLG Analysis' (ID: sg-74498502) highlighted. Below the table, the details for the selected security group are shown, including its inbound rules. The 'Inbound' tab is active, displaying a single rule: a Custom TCP Rule allowing traffic from 172.31.26.133/32 on port 0.

Name	Group ID	Group Name	VPC ID	Description
	sg-19cdb26e	ElasticMapReduce-master	vpc-1aaba062	Master group for Elastic MapReduce created on 2018-02-13T00:57:...
	sg-29f4f65d	default	vpc-1aaba062	default VPC security group
	sg-2a94595c	launch-wizard-13	vpc-1aaba062	launch-wizard-13 created 2018-03-03T09:49:33.380-05:00
Demo01	sg-2d373959	Demo01-WebServerSecurity...	vpc-1aaba062	Enable connection from your IP
	sg-74498502	SLG Analysis	vpc-1aaba062	For bulk_extractor analysis
	sg-91bec9e6	open	vpc-1aaba062	open
	sg-c59459b3	launch-wizard-11	vpc-1aaba062	launch-wizard-11 created 2018-03-03T09:45:17.235-05:00
	sg-f0c5ba87	ElasticMapReduce-slave	vpc-1aaba062	Slave group for Elastic MapReduce created on 2018-02-13T00:57:5...
	sg-f68a4780	launch-wizard-12	vpc-1aaba062	launch-wizard-12 created 2018-03-03T09:47:39.379-05:00

Security Group: sg-74498502

Description Inbound Outbound Tags

Edit

Type	Protocol	Port Range	Source	Description
Custom TCP Rule	TCP	0	172.31.26.133/32	

Now, create the EFS with my added security group

Elastic File System Manage...

What Is Amazon Elastic File S...

Mounting Automatically - Ama...

EC2 Management Console

Secure | https://console.aws.amazon.com/efs/home?region=us-east-1#/wizard/1

AppsBbBbCFRS780BbDiscussion Board -...ScheduleAWS EC2 EMR EC2 Instance Pricing...\$\$\$GMU

awsServicesResource GroupsEC2EMRS3Athena

sgarfin2@gmu.eduN. VirginiaSupport

Create file system

Step 1: Configure file system access

Step 2: Configure optional settings

Step 3: Review and create

Configure file system access

An Amazon EFS file system is accessed by EC2 instances running inside one of your VPCs. Instances connect to a file system by using a network interface called a mount target. Each mount target has an IP address, which we assign automatically or you can specify.

VPCvpc-1aaba062 (default)

Create mount targets

Instances connect to a file system by using mount targets you create. We recommend creating a mount target in each of your VPC's Availability Zones so that EC2 instances across your VPC can access the file system.

	Availability Zone	Subnet	IP address	Security groups
<input checked="" type="checkbox"/>	us-east-1a	subnet-e22613a9 (default)	Automatic	sg-29f4f65d - default sg-74498502 - SLG Analysis
<input checked="" type="checkbox"/>	us-east-1b	subnet-e51676b8 (default)	Automatic	sg-29f4f65d - default sg-74498502 - SLG Analysis
<input checked="" type="checkbox"/>	us-east-1c	subnet-92c892f6 (default)	Automatic	sg-29f4f65d - default sg-74498502 - SLG Analysis
<input checked="" type="checkbox"/>	us-east-1d	subnet-90523ebf (default)	Automatic	sg-29f4f65d - default sg-74498502 - SLG Analysis
<input checked="" type="checkbox"/>	us-east-1e	subnet-66324259 (default)	Automatic	sg-29f4f65d - default sg-74498502 - SLG Analysis
<input checked="" type="checkbox"/>	us-east-1f	subnet-ec31cce3 (default)	Automatic	sg-29f4f65d - default sg-74498502 - SLG Analysis

CancelNext Step

My file system info:

Name	File system ID	Metered size	Number of mount targets	Creation date
	fs-1a3ac552	6.0 KiB	6	2018-03-03T21:27:12Z

Other details

Owner ID 309467262965

Life cycle state Available

Performance mode General Purpose

Encrypted No

Tags

No tags added

[Manage tags](#)

File system access

DNS name fs-1a3ac552.efs.us-east-1.amazonaws.com ?

[Amazon EC2 mount instructions](#)

[AWS Direct Connect mount instructions](#)

[Manage file system access](#)

Mount targets

VPC	Availability Zone	Subnet	IP address	Mount target ID	Network interface ID	Security groups	Life cycle state
vpc-1aaba062 (default)	us-east-1c	subnet-92c892f6 (default)	172.31.10.234	fsmt-31589779	eni-3a715cba		Creating
	us-east-1e	subnet-66324259 (default)	172.31.51.122	fsmt-3358977b	eni-29a76493		Creating
	us-east-1a	subnet-e22613a9 (default)	172.31.23.151	fsmt-3658977e	eni-3b82d622		Creating
	us-east-1d	subnet-90523ebf (default)	172.31.82.201	fsmt-3f589777	eni-49a05186		Creating
	us-east-1b	subnet-e51676b8 (default)	172.31.44.94	fsmt-c9589781	eni-cae9091b		Creating
	us-east-1f	subnet-ec31cce3 (default)	172.31.69.36	fsmt-cb589783	eni-1811f2bd		Creating

Mount the file system

```
$ sudo yum install -y nfs-utils
```

```
$ df
```

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
devtmpfs	72071020	84	72070936	1%	/dev
tmpfs	72080004	0	72080004	0%	/dev/shm
/dev/nvme0n1p1	8123812	1075712	6947852	14%	/
/dev/nvme1n1	82438832	3356568	74871576	5%	/mnt

```
$ sudo mkdir /efs
```

```
$ sudo mount -t nfs -o
```

```
nfsvers=4.1,rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2
```

```
fs-1a3ac552.efs.us-east-1.amazonaws.com:/ /efs
```

```
$ df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
devtmpfs	69G	88K	69G	1%	/dev
tmpfs	69G	0	69G	0%	/dev/shm
/dev/nvme0n1p1	7.8G	1.1G	6.7G	14%	/
/dev/nvme1n1	79G	3.3G	72G	5%	/mnt
fs-1a3ac552.efs.us-east-1.amazonaws.com:/	8.0E	0	8.0E	0%	/efs

```
$
```

You can set the file system to mount on reboot.

Be careful — if you damage /etc/fstab, the system won't boot!

Add the boot instructions to the /etc/fstab

```
$ sudo bash
# cat >> /etc/fstab
fs-1a3ac552.efs.us-east-1.amazonaws.com:/ /efs      nfs4
nfsvers=4.1,rsiz=1048576,wsiz=1048576,hard,timeo=600,retrans=2,_netdev 0 0
^d
```

and test:

```
# umount /efs
# mount /efs
# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
devtmpfs	69G	88K	69G	1%	/dev
tmpfs	69G	0	69G	0%	/dev/shm
/dev/nvme0n1p1	7.8G	1.1G	6.7G	14%	/
/dev/nvme1n1	79G	7.3G	68G	10%	/mnt
fs-1a3ac552.efs.us-east-1.amazonaws.com:/	8.0E	0	8.0E	0%	/efs

```
#
```

Running bulk_extractor in AWS

There is no reason to run bulk_extractor on a small machine.

The screenshot shows the AWS EC2 Management Console in the 'us-east-1' region. The 'Launch Instance Wizard' is at 'Step 2: Choose an Instance Type'. The table below lists available instance types, with 'm5.12xlarge' selected.

	General purpose	Instance type	vCPUs	Memory (GiB)	EBS only	Yes	Low to moderate	Yes
<input type="checkbox"/>	General purpose	t2.xlarge	4	16	EBS only	-	Moderate	Yes
<input type="checkbox"/>	General purpose	t2.2xlarge	8	32	EBS only	-	Moderate	Yes
<input type="checkbox"/>	General purpose	m5.large	2	8	EBS only	Yes	Up to 10 Gigabit	Yes
<input type="checkbox"/>	General purpose	m5.xlarge	4	16	EBS only	Yes	Up to 10 Gigabit	Yes
<input type="checkbox"/>	General purpose	m5.2xlarge	8	32	EBS only	Yes	Up to 10 Gigabit	Yes
<input type="checkbox"/>	General purpose	m5.4xlarge	16	64	EBS only	Yes	Up to 10 Gigabit	Yes
<input checked="" type="checkbox"/>	General purpose	m5.12xlarge	48	192	EBS only	Yes	10 Gigabit	Yes
<input type="checkbox"/>	General purpose	m5.24xlarge	96	384	EBS only	Yes	25 Gigabit	Yes
<input type="checkbox"/>	General purpose	m4.large	2	8	EBS only	Yes	Moderate	Yes
<input type="checkbox"/>	General purpose	m4.xlarge	4	16	EBS only	Yes	High	Yes

Navigation buttons: Cancel, Previous, Review and Launch, Next: Configure Instance Details

The spot advisor is a useful tool for finding spot instances.

The screenshot shows the AWS Spot Advisor interface in the EC2 Management Console. The browser address bar shows the URL `https://console.aws.amazon.com/ec2sp/v1/spot/advisor?region=us-east-1`. The page title is "Spot Advisor" with a "Cancel" button in the top right.

What kind of application or task will these instances support?

Three options are available:

- Default** (selected): Obtain instances of a consistent CPU/memory ratio at the best price.
- Web Service**: Secure and maintain capacity using instances of a similar size to run your Web service.
- MapReduce job**: Acquire low-cost, single AZ instances for your MapReduce job. *e.g. Hadoop, Spark, etc.*

Enter your compute specifications (minimal requirements) and we'll recommend a fleet

Fields for configuration:

- vCPU: 64
- Memory GiB: 120
- Platform: Linux
- Availability Zone: Any
- Amount required: 1

Or inherit values from [an instance type](#)

Your recommended fleet

The selected instance pools will be used interchangeably to fulfill and maintain your specified compute requirements. The actual instance pools and quantities used from this fleet are dynamic to ensure that your capacity is maintained and that your specified fulfillment priority is honored.

Instance type	Average Spot price	Interruption likelihood
<input checked="" type="checkbox"/> m4.16xlarge 64 vCPU, 256GiB U x1	\$0.9827/hr	Low
<input checked="" type="checkbox"/> c5.18xlarge 72 vCPU, 144GiB U x1	\$1.0953/hr	Low
<input checked="" type="checkbox"/> r4.16xlarge 64 vCPU, 488GiB U x1	\$1.1111/hr	Low
Total (18 instance pools) ✓ Strong breadth of instance types to fulfill/maintain your request		Low
Availability Zones us-east-1a, us-east-1b, us-east-1c, us-east-1d, us-east-1e, us-east-1f		
Estimated fleet price \$1.063 /hr 71% savings		

Configure your fleet

Be careful! If you have a standing spot request, it will relaunch instances when you terminate them.

The screenshot shows the AWS EC2 Management Console interface. The left sidebar contains navigation links for various AWS services, with 'Spot Requests' highlighted under the 'INSTANCES' category. The main content area displays the 'Request Spot Instances' page. At the top, there are tabs for 'Request Spot Instances', 'Spot Advisor', 'Actions', and 'Pricing History'. Below these tabs, there are filters for 'Request type: all' and 'State: all', along with a search bar. A table lists the spot requests, showing one request with the ID 'sir-zdzi74bg'. The table columns are: Request Id, Request type, Instance type, State, Capacity, Status, Persistence, and Created. The request is in an 'active' state and has a 'fulfilled' status. Below the table, there is a message: 'Select a Spot request above to see more details'.

Request Id	Request type	Instance type	State	Capacity	Status	Persistence	Created
sir-zdzi74bg	instance	t2.micro	active	i-0607a3045ec9...	fulfilled	one-time	14 hours ago

Hibernate interrupt behavior

Instance details

Monitoring ⓘ

☐ Enable CloudWatch detailed monitoring

Health check ⓘ

☐ Replace unhealthy instances

Interruption behavior ⓘ

✓ Terminate

Stop

Hibernate

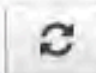


User data ⓘ

Instance tags ⓘ

Key	Value
-----	-------

Give the EC2 instance a sufficient IAM role.

Set keypair and role

Key pair name ⓘ	<input type="text" value="mucha"/>		Create new key
IAM instance profile ⓘ	<div>(optional) ✓ Bulk_Extractor_Experiment EMR_EC2_DefaultRole</div>		Create new IAM
IAM fleet role ⓘ	aws-ec2-spot-fleet-tagging-role 		

Manage firewall rules

Here is standing request creates the analysis VM.

The screenshot displays the AWS Management Console interface. The left sidebar contains navigation links for EC2 Dashboard, INSTANCES, IMAGES, ELASTIC BLOCK STORE, NETWORK & SECURITY, LOAD BALANCING, and AUTO SCALING. The main content area shows the 'Request Spot Instances' page for a specific request ID: sfr-8807fbf6-fdff-48a0-9b69-c0456192d82f. A table lists two requests: a 'fleet' request (submitted) and an 'instance' request (active). Below the table, the 'Description' tab is selected, showing detailed configuration for the selected request.

Request Id	Request type	Instance type	State	Capacity	Status	Persistence	Created	Max price
sfr-8807fbf6-fdff-48...	fleet	m4.16xlarge,c5....	submitted	0 of 1		maintain	a few seconds a...	\$4.256
sir-zdzi74bg	instance	t2.micro	active	i-0607a3045ec9...	fulfilled	one-time	14 hours ago	\$0.02

Request Id: sfr-8807fbf6-fdff-48a0-9b69-c0456192d82f	
Request Id	sfr-8807fbf6-fdff-48a0-9b69-c0456192d82f
Request type	fleet
Created	3/3/2018, 10:06:30 AM
State	submitted
Status	
Target capacity	1
Allocation strategy	diversified
Instance type(s)	m4.16xlarge \$3.83, c5.18xlarge \$3.06, r4.16xlarge \$4.256
AMI ID	ami-55ef662f
Subnet	subnet-66324259,subnet-90523ebf,subnet-ec31cce3,subnet-92c892f6,subnet-e22613a9,subnet-e51676b8
IAM fleet role	aws-ec2-spot-fleet-tagging-role
Max price	\$4.256
Persistence	maintain
Key pair name	mucha
IAM role	Bulk_Extractor_Experiment
EBS-optimized	no
Monitoring	no
Health check	no
Tenancy	default
Interruption behavior	terminate
Classic load balancers	-
Target groups	-

The request is submitted, and a VM has been created.

The screenshot shows the AWS EC2 Management Console. The left sidebar contains navigation links for EC2 Dashboard, Events, Tags, Reports, Limits, INSTANCES, IMAGES, ELASTIC BLOCK STORE, NETWORK & SECURITY, LOAD BALANCING, and AUTO SCALING. The main content area displays the 'Request Spot Instances' dashboard. A table lists the spot requests, with the first request (sfr-8807fbf6-fdff-48a0-9b69-c0456192d82f) highlighted. Below the table, the details for this request are shown, including its state (submitted), capacity (0 of 1), and various configuration options.

Request Id	Request type	Instance type	State	Capacity	Status	Persistence	Created	Max price
sfr-8807fbf6-fdff-48a0-9b69-c0456192d82f	fleet	m4.16xlarge,c5...	submitted	0 of 1		maintain	a few seconds a...	\$4.256
sir-wssr6a6h	instance	c5.18xlarge	active	i-0d1ab43350cd...	fulfilled	persistent	a minute ago	\$3.06
sir-zdzi74bg	instance	t2.micro	active	i-0607a3045ec9...	fulfilled	one-time	14 hours ago	\$0.02

Request Id: sfr-8807fbf6-fdff-48a0-9b69-c0456192d82f

Description | **Instances** | **History** | **Auto Scaling**

Request Id	sfr-8807fbf6-fdff-48a0-9b69-c0456192d82f	Max price	\$4.256
Request type	fleet	Persistence	maintain
Created	3/3/2018, 10:06:30 AM	Key pair name	muchu
State	submitted	IAM role	Bulk_Extractor_Experiment
Status		EBS-optimized	no
Target capacity	1	Monitoring	no
Allocation strategy	diversified	Health check	no
Instance type(s)	m4.16xlarge \$3.83, c5.18xlarge \$3.06, r4.16xlarge \$4.256	Tenancy	default
AMI ID	ami-55ef662f	Interruption behavior	terminate
Subnet	subnet-66324259,subnet-90523ebf,subnet-ec31cce3,subnet-92c892f6,subnet-e22613a9,subnet-e51676b8	Classic load balancers	-
IAM fleet role	aws-ec2-spot-fleet-tagging-role	Target groups	-

lscpu shows that this instance has 72 CPUs!

```
[ec2-user@ip-172-31-26-133 ~]$ lscpu
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):               72
On-line CPU(s) list:   0-71
Thread(s) per core:    2
Core(s) per socket:    18
Socket(s):             2
NUMA node(s):          2
Vendor ID:             GenuineIntel
CPU family:            6
Model:                85
Model name:            Intel(R) Xeon(R) Platinum 8124M CPU @ 3.00GHz
Stepping:              3
CPU MHz:               3000.000
BogoMIPS:              6000.00
Hypervisor vendor:    KVM
Virtualization type:   full
L1d cache:            32K
L1i cache:            32K
L2 cache:             1024K
L3 cache:             25344K
NUMA node0 CPU(s):    0-17,36-53
NUMA node1 CPU(s):    18-35,54-71
[ec2-user@ip-172-31-26-133 ~]$
```


The IAM Role works:

No configuration required for 'aws' command!

```
[ec2-user@ip-172-31-26-133 ~]$ aws s3 ls
2018-02-24 18:28:03 aws-athena-query-results-309467262965-us-east-1
2018-02-13 02:31:54 aws-logs-309467262965-us-east-1
2018-02-19 00:39:21 cfrs780
2018-02-19 00:55:04 crfs780-trails
[ec2-user@ip-172-31-26-133 ~]$ ls -al .aws
ls: cannot access .aws: No such file or directory
[ec2-user@ip-172-31-26-133 ~]$ aws s3 cp s3://cfrs780/bin/bulk_extractor
bulk_extractor
download: s3://cfrs780/bin/bulk_extractor to ./bulk_extractor
```

Change the hostname to make identifying the hosts easier

```
[ec2-user@ip-172-31-26-133 ~]$ sudo hostname beefy
[ec2-user@ip-172-31-26-133 ~]$ logout
Connection to 107.23.131.99 closed.
[Dance ~ 10:19:54]$ ssh -A ec2-user@107.23.131.99
Last login: Sat Mar  3 15:16:22 2018 from 68.33.78.230
```

```
  _|  _|_ )
 _|  (    /  Amazon Linux AMI
 _|\_||_|
```

```
https://aws.amazon.com/amazon-linux-ami/2017.09-release-notes/
6 package(s) needed for security, out of 12 available
Run "sudo yum update" to apply all updates.
[ec2-user@beefy ~]$
```

Performing an experiment with bulk_extractor and EC2

Experimental design: are we getting clean EBS boot volumes?

Experimental outline:

- Create VMs with 40GB, 100GB and 1000GB boot volumes
- Specify that volumes are to be kept after machine termination.
- Terminate each machine after boot
- Attach the volumes to a “beefy” VM
- Run bulk_extractor on each volume, storing the results in a single file system.
- Compare the results.

Wrong availability zone...

```
$ zone=`curl -s http://169.254.169.254/latest/meta-data/placement/availability-zone`  
$ region=`echo $zone | sed s/.$//`  
$ id=`curl -s http://169.254.169.254/latest/meta-data/instance-id`  
$ aws ec2 attach-volume --device /dev/sdf --instance-id $id --region $region --  
volume-id vol-08baaf673f8051ab6
```

An error occurred (InvalidVolume.ZoneMismatch) when calling the AttachVolume operation: The volume 'vol-08baaf673f8051ab6' is not in the same availability zone as instance 'i-0d1ab43350cd3e1ee'
[ec2-user@beefy ~]\$

To move a volume to a different availability zone: Create a snapshot, and create a volume from the snapshot.

```
[ec2-user@beefy ~]$ aws ec2 create-snapshot --description stats_snap --volume-id
vol-08baaf673f8051ab6 --region $region
{
  "Description": "stats_snap",
  "Tags": [],
  "Encrypted": false,
  "VolumeId": "vol-08baaf673f8051ab6",
  "State": "pending",
  "VolumeSize": 80,
  "StartTime": "2018-03-03T15:24:36.000Z",
  "Progress": "",
  "OwnerId": "309467262965",
  "SnapshotId": "snap-0fa3c10efa5cda4ef"
}
[ec2-user@beefy ~]$ aws ec2 create-volume --availability-zone $zone --region
$region --snapshot-id snap-0fa3c10efa5cda4ef
{
  "AvailabilityZone": "us-east-1a",
  "Encrypted": false,
  "VolumeType": "standard",
  "VolumeId": "vol-07ac3927385281d53",
  "State": "creating",
  "SnapshotId": "snap-0fa3c10efa5cda4ef",
  "CreateTime": "2018-03-03T15:27:30.962Z",
  "Size": 80
}
[ec2-user@beefy ~]$
```

Now we can attach the volume

```
[ec2-user@beefy ~]$ aws ec2 attach-volume --device /dev/sdf --instance-id $id --
region $region --volume-id vol-07ac3927385281d53
{
  "AttachTime": "2018-03-03T15:28:08.830Z",
  "InstanceId": "i-0d1ab43350cd3e1ee",
  "VolumeId": "vol-07ac3927385281d53",
  "State": "attaching",
  "Device": "/dev/sdf"
}
[ec2-user@beefy ~]$ ls -l /mnt/
total 28
-rw-r--r-- 1 root root      0 Mar  3 14:57 0_VOLUME_FOR_STATS
drwx----- 2 root root 16384 Mar  3 02:22 lost+found
drwxr-xr-x 5 root root  4096 Mar  3 02:35 vm_exp1
drwxr-xr-x 5 root root  4096 Mar  3 03:08 vm_exp2
drwxr-xr-x 5 root root  4096 Mar  3 03:21 vm_exp3
[ec2-user@beefy ~]$
```

Now attach the volumes for the 40GB, 100GB and 1000GB

All needed to be moved to the availability zone of the beefy:

```
$ aws ec2 create-volume --availability-zone $zone --region $region --snapshot-id  
snap-0e0bfd05ae61cc002  
$ aws ec2 create-volume --availability-zone $zone --region $region --snapshot-id  
snap-07f5081edb3e0f6c5  
$ aws ec2 create-volume --availability-zone $zone --region $region --snapshot-id  
snap-0e70b2b58d662e10b
```

```
$ aws ec2 attach-volume --device /dev/sdi --instance-id $id --region $region --volume-id  
vol-0a9d462b1f4229dcd  
$ aws ec2 attach-volume --device /dev/sdj --instance-id $id --region $region --volume-id  
vol-0c7e4831ea1e04585  
$ aws ec2 attach-volume --device /dev/sdk --instance-id $id --region $region --volume-id  
vol-04297e43473766836
```

Don't mount them!

```
[ec2-user@beefy ~]$ df -h  
Filesystem      Size  Used Avail Use% Mounted on  
devtmpfs        69G   84K   69G   1% /dev  
tmpfs           69G    0    69G   0% /dev/shm  
/dev/nvme0n1p1  7.8G  1.1G   6.7G  14% /  
/dev/nvme1n1    79G  1.7G   73G   3% /mnt  
[ec2-user@beefy ~]$ ls -l /mnt/  
total 28  
-rw-r--r-- 1 root root      0 Mar  3 14:57 0_VOLUME_FOR_STATS  
drwx----- 2 root root 16384 Mar  3 02:22 lost+found  
[ec2-user@beefy ~]$
```


Get the docs for create-volume

<https://docs.aws.amazon.com/cli/latest/search.html>

The screenshot shows a web browser window with the URL `https://docs.aws.amazon.com/cli/latest/search.html?q=create+volume&check_keywords=yes&area=default`. The page is titled "AWS CLI Command Reference" and features a search bar with the text "create volume" and a "search" button. Below the search bar, the results are displayed under the heading "Searching..". The results list several commands, including `add-instance-groups`, `add-tags-to-resource`, `bundle-instance`, and `copy-image`. The `add-instance-groups` result is highlighted, showing its description: "...nd establishes the bid price for the instances. * ``[EbsConfiguration]`` - Specifies additional Amazon EBS storage **volumes** attached to EC2 instances using an inline JSON structure. * ``[AutoScalingPolicy]`` - Specifies an automati...". The `add-tags-to-resource` result is also visible, showing its description: "...u can add tags to the following AWS Storage Gateway resources: * Storage gateways of all types * Storage **Volumes** * Virtual Tapes You can **create** a maximum of 10 tags for each resource. Virtual tapes and storage...". The `bundle-instance` result shows its description: "...tion ===== Bundles an Amazon instance store-backed Windows instance. During bundling, only the root device **volume** (C:\) is bundled. Data on other instance store volumes is not preserved. ... note:: This action is not...". The `copy-image` result shows its description: "...An identifier for the AWS Key Management Service (AWS KMS) customer master key (CMK) to use when creating the encrypted **volume**. This parameter is only required if you want to use a ... default CMK; if this parameter is not specified, the de...".

Search — AWS CLI 1.14.50 Col x

Secure | https://docs.aws.amazon.com/cli/latest/search.html?q=create+volume&check_keywords=yes&area=default

Apps Bb CFRS780 Bb Discussion Board —... Schedule AWS EC2 EMR EC2 Instance Pricing... \$\$ GMU

AWS CLI Command Reference Home User Guide Forum GitHub Star 6,090

amazon web services

Feedback

Did you find this page useful?
Do you have a suggestion?
[Give us feedback](#) or send us a [pull request](#) on GitHub.

User Guide

First time using the AWS CLI?
See the [User Guide](#) for help getting started.

Search

From here you can search these documents. Enter your search words into the box below and click "search". Note that the search function will automatically search for all of the words. Pages containing fewer words won't appear in the result list.

create volume search

Searching..

- add-instance-groups**
...nd establishes the bid price for the instances. * ``[EbsConfiguration]`` - Specifies additional Amazon EBS storage **volumes** attached to EC2 instances using an inline JSON structure. * ``[AutoScalingPolicy]`` - Specifies an automati...
- add-tags-to-resource**
...u can add tags to the following AWS Storage Gateway resources: * Storage gateways of all types * Storage **Volumes** * Virtual Tapes You can **create** a maximum of 10 tags for each resource. Virtual tapes and storage...
- bundle-instance**
...tion ===== Bundles an Amazon instance store-backed Windows instance. During bundling, only the root device **volume** (C:\) is bundled. Data on other instance store volumes is not preserved. ... note:: This action is not...
- copy-image**
...An identifier for the AWS Key Management Service (AWS KMS) customer master key (CMK) to use when creating the encrypted **volume**. This parameter is only required if you want to use a ... default CMK; if this parameter is not specified, the de...

Now we can run bulk_extractor!

```
$ aws s3 cp s3://cfrs780/bin/bulk_extractor bulk_extractor
[root@beefy ec2-user]# ls -l
total 21924
-rw-rw-r-- 1 ec2-user ec2-user 22450072 Mar  3 02:04 bulk_extractor
[root@beefy ec2-user]# chmod +x bulk_extractor
[root@beefy ec2-user]#

[root@beefy ec2-user]# ./bulk_extractor -o /mnt/vm_40G -e wordlist /dev/sdi
bulk_extractor version: 1.6.0-dev
Hostname: beefy
Input file: /dev/sdi
Output directory: /mnt/vm_40G
Disk Size: 42949672960
Threads: 72
Attempt to open /dev/sdi
15:44:35 Offset 67MB (0.16%) Done in 0:47:28 at 16:32:03
15:44:39 Offset 150MB (0.35%) Done in 0:39:35 at 16:24:14
```

Once we have results in EBS, we could copy them to S3 for safe keeping...

```
[ec2-user@beefy mnt]$ du -sh .
du: cannot read directory './lost+found': Permission denied
3.2G .
[ec2-user@beefy mnt]$ time aws s3 cp --recursive . s3://cfrs780/
bulk_extractor_results/
warning: Skipping file /mnt/lost+found. File/Directory is not readable.
upload: ./0_VOLUME_FOR_STATS to s3://cfrs780/bulk_extractor_results/
0_VOLUME_FOR_STATS
upload: vm_1000G/ether.txt to s3://cfrs780/bulk_extractor_results/vm_1000G/
ether.txt
upload: vm_1000G/exif.txt to s3://cfrs780/bulk_extractor_results/vm_1000G/
exif.txt
upload: vm_1000G/find.txt to s3://cfrs780/bulk_extractor_results/vm_1000G/
find.txt
...
upload: vm_exp3/wordlist_split_000.txt to s3://cfrs780/bulk_extractor_results/
vm_exp3/wordlist_split_000.txt
upload: vm_exp3/wordlist.txt to s3://cfrs780/bulk_extractor_results/vm_exp3/
wordlist.txt

real 0m31.065s
user 0m13.972s
sys 0m7.476s
[ec2-user@beefy mnt]$
```

We could also save the result in EFS

```
[ec2-user@beefy mnt]$ sudo mkdir /efs/results
[ec2-user@beefy mnt]$ sudo chown $USER /efs/results
[ec2-user@beefy mnt]$ time cp -r * /efs/results/
cp: cannot access 'lost+found': Permission denied
```

```
real 0m50.788s
```

```
user 0m0.016s
```

```
sys 0m3.176s
```

```
[ec2-user@beefy mnt]$ ls -l /efs/results/
```

```
total 32
```

```
-rw-r--r-- 1 ec2-user ec2-user 0 Mar 3 21:33 0_VOLUME_FOR_STATS
```

```
drwx----- 2 ec2-user ec2-user 6144 Mar 3 21:33 lost+found
```

```
drwxr-xr-x 5 ec2-user ec2-user 6144 Mar 3 21:33 vm_1000G
```

```
drwxr-xr-x 5 ec2-user ec2-user 6144 Mar 3 21:33 vm_100G
```

```
drwxr-xr-x 5 ec2-user ec2-user 6144 Mar 3 21:34 vm_40G
```

```
drwxr-xr-x 5 ec2-user ec2-user 6144 Mar 3 21:34 vm_exp1
```

```
drwxr-xr-x 5 ec2-user ec2-user 6144 Mar 3 21:34 vm_exp2
```

```
drwxr-xr-x 5 ec2-user ec2-user 6144 Mar 3 21:34 vm_exp3
```

```
[ec2-user@beefy mnt]$
```

S3 copy times:

```
real 0m31.065s
```

```
user 0m13.972s
```

```
sys 0m7.476s
```

Compared to copying to S3

- Slower to send data
- Easier to use data once sent

Speed:

Let's compare running bulk_extractor out of EBS and EFS.

```
$ sudo mkdir /mnt/work /efs/work
$ sudo chown $USER /mnt/work /efs/work
$ time wget -O /mnt/work/nps-2009-domexusers.E01 \
  http://downloads.digitalcorpora.org/corpora/drives/nps-2009-domexusers/
nps-2009-domexusers.E01
```

I started the download...

- To /mnt/work and got 27MB/sec
- After 1GB, I started a simultaneous download to /efs/work in another window
- Download rate remained 27MB *in both windows...*

I ran bulk_extractor sequentially, same conditions:

- reboot
- mount file system
- run bulk_extractor from /efs/work or /mnt/work

from /mnt (EBS)

```
$ cd /mnt/work
$ ls -al
total 4268484
drwxr-xr-x  2 ec2-user root          4096 Mar  3 21:42 .
drwxr-xr-x 10 root      root          4096 Mar  3 21:40 ..
-rw-rw-r--  1 ec2-user ec2-user 4370913825 May 14 2012 nps-2009-domexusers.E01
$ ~/bulk_extractor -o out -e wordlist nps-2009-domexusers.E01
bulk_extractor version: 1.6.0-dev
Hostname: ip-172-31-18-68
Input file: nps-2009-domexusers.E01
Output directory: out
Disk Size: 42949672960
Threads: 72
 2:16:06 Offset 67MB (0.16%) Done in 0:38:19 at 02:54:25
 2:16:09 Offset 150MB (0.35%) Done in 0:30:14 at 02:46:23
 2:16:11 Offset 234MB (0.55%) Done in 0:26:31 at 02:42:42
 2:16:13 Offset 318MB (0.74%) Done in 0:23:54 at 02:40:07
 2:16:15 Offset 402MB (0.94%) Done in 0:22:47 at 02:39:02
 2:16:17 Offset 486MB (1.13%) Done in 0:21:50 at 02:38:07
 2:16:20 Offset 570MB (1.33%) Done in 0:21:22 at 02:37:42
 2:16:22 Offset 654MB (1.52%) Done in 0:20:44 at 02:37:06
...
```

```
3. ec2-user@ip-172-31-18-68:~ (ssh)
top - 02:17:54 up 4:07, 2 users, load average: 7.07, 2.76, 1.13
Tasks: 586 total, 1 running, 584 sleeping, 1 stopped, 0 zombie
Cpu(s): 15.3%us, 0.0%sy, 0.0%ni, 83.9%id, 0.8%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 144160008k total, 6518480k used, 137641528k free, 35388k buffers
Swap: 0k total, 0k used, 0k free, 3574404k cached
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
2696	ec2-user	20	0	5826m	2.0g	9m	D	1103.7	1.5	14:53.52	bulk_extractor
26	root	20	0	0	0	0	S	0.3	0.0	0:00.05	kworker/3:0
2794	ec2-user	20	0	15716	2648	1928	R	0.3	0.0	0:00.05	top
3687	root	20	0	6476	100	0	S	0.3	0.0	0:00.27	rngd
1	root	20	0	19648	2652	2320	S	0.0	0.0	0:02.06	init
2	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kthreadd
3	root	20	0	0	0	0	S	0.0	0.0	0:00.00	ksoftirqd/0
5	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kworker/0:0H
6	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kworker/u144:0
8	root	20	0	0	0	0	S	0.0	0.0	0:00.99	rcu_sched
9	root	20	0	0	0	0	S	0.0	0.0	0:00.00	rcu_bh
10	root	RT	0	0	0	0	S	0.0	0.0	0:00.00	migration/0
11	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	lru-add-drain
12	root	20	0	0	0	0	S	0.0	0.0	0:00.00	cpuhp/0
13	root	20	0	0	0	0	S	0.0	0.0	0:00.00	cpuhp/1
14	root	RT	0	0	0	0	S	0.0	0.0	0:00.00	migration/1
15	root	20	0	0	0	0	S	0.0	0.0	0:00.00	ksoftirqd/1
16	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kworker/1:0

```
2:21:36 Offset 42849MB (99.77%) Done in 0:00:00 at 02:21:36
2:21:36 Offset 42932MB (99.96%) Done in 0:00:00 at 02:21:36
All data are read; waiting for threads to finish...
Time elapsed waiting for 3 threads to finish:
    (timeout in 60 min.)
All Threads Finished!
Producer time spent waiting: 0 sec.
Average consumer time spent waiting: 302.728 sec.
*****
** bulk_extractor is probably I/O bound. **
**          Run with a faster drive          **
**          to get better performance.        **
*****
MD5 of Disk Image: 8e7176524a64376631cd7dc9d90339f1
Phase 2. Shutting down scanners
Phase 3. Uniquifying and recombining wordlist
Phase 3. Creating Histograms
Elapsed time: 391.168 sec.
Total MB processed: 42949
Overall performance: 109.799 MBytes/sec (1.52498 MBytes/sec/thread)
Total email features found: 8757
$
```


Run bulk_extractor from/to EFS

Reboot instance, then:

```
$ sudo mkdir /efs
$ sudo mount -t nfs -o
nfsvers=4.1,rsz=1048576,wsz=1048576,hard,timeo=600,retrans=2
fs-1a3ac552.efs.us-east-1.amazonaws.com:/ /efs

[ec2-user@ip-172-31-18-68 work]$ ~/bulk_extractor/src/bulk_extractor -o out -e
wordlist nps-2009-domexusers.E01
bulk_extractor version: 1.6.0-dev
Hostname: ip-172-31-18-68
Input file: nps-2009-domexusers.E01
Output directory: out
Disk Size: 42949672960
Threads: 72
2:30:42 Offset 67MB (0.16%) Done in 0:17:58 at 02:48:40
2:30:42 Offset 150MB (0.35%) Done in 0:10:05 at 02:40:47
2:30:43 Offset 234MB (0.55%) Done in 0:07:59 at 02:38:42
```

```
6. ec2-user@ip-172-31-18-68:~ (ssh)
top
Task top - 02:32:29 up 4 min, 2 users, load average: 17.31, 7.67, 2.87
Cpu Tasks: 565 total, 2 running, 561 sleeping, 2 stopped, 0 zombie
Mem: Cpu(s): 16.1%us, 0.2%sy, 0.0%ni, 83.4%id, 0.2%wa, 0.0%hi, 0.0%si, 0.0%st
Swap Mem: 144159548k total, 8094272k used, 136065276k free, 9172k buffers
Swap: 0k total, 0k used, 0k free, 4752948k cached

PID
398
410
3986 ec2-user 20 0 6334m 2.6g 9924 R 1171.7 1.9 47:47.95 bulk_extractor
3961 root 0 -20 0 0 0 S 5.0 0.0 0:02.64 kworker/7:1H
2504 root 20 0 0 0 0 S 1.0 0.0 0:00.44 kworker/7:1
466 root 20 0 0 0 0 S 0.3 0.0 0:00.04 kworker/u145:3
3592 root 20 0 6476 100 0 S 0.3 0.0 0:00.11 rngd
4059 root 20 0 0 0 0 S 0.3 0.0 0:00.04 kworker/19:2
4104 ec2-user 20 0 15708 2468 1768 R 0.3 0.0 0:00.08 top
1 root 20 0 19644 2500 2176 S 0.0 0.0 0:02.08 init
2 root 20 0 0 0 0 S 0.0 0.0 0:00.00 kthreadd
3 root 20 0 0 0 0 S 0.0 0.0 0:00.00 ksoftirqd/0
4 root 20 0 0 0 0 S 0.0 0.0 0:00.00 kworker/0:0
5 root 0 -20 0 0 0 S 0.0 0.0 0:00.00 kworker/0:0H
6 root 20 0 0 0 0 S 0.0 0.0 0:00.00 kworker/u144:0
7 root 20 0 0 0 0 S 0.0 0.0 0:00.00 kworker/u145:0
8 root 20 0 0 0 0 S 0.0 0.0 0:00.03 rcu_sched
9 root 20 0 0 0 0 S 0.0 0.0 0:00.00 rcu_bh
10 root RT 0 0 0 0 S 0.0 0.0 0:00.00 migration/0
11 root 0 -20 0 0 0 S 0.0 0.0 0:00.00 lru-add-drain
```

```
2:33:52 Offset 42765MB (99.57%) Done in 0:00:00 at 02:33:52
2:33:52 Offset 42849MB (99.77%) Done in 0:00:00 at 02:33:52
2:33:52 Offset 42932MB (99.96%) Done in 0:00:00 at 02:33:52
All data are read; waiting for threads to finish...
Time elapsed waiting for 5 threads to finish:
    (timeout in 60 min.)
All Threads Finished!
Producer time spent waiting: 0 sec.
Average consumer time spent waiting: 142.452 sec.
*****
** bulk_extractor is probably I/O bound. **
**      Run with a faster drive      **
**      to get better performance.    **
*****
MD5 of Disk Image: 8e7176524a64376631cd7dc9d90339f1
Phase 2. Shutting down scanners
Phase 3. Uniquifying and recombining wordlist
Phase 3. Creating Histograms
Elapsed time: 263.151 sec.
Total MB processed: 42949
Overall performance: 163.213 MBytes/sec (2.26685 MBytes/sec/thread)
Total email features found: 8757
[ec2-user@ip-172-31-18-68 work]$
```

Another example:

m5.4xlarge

```
[ec2-user@beefy2 work]$ ./bulk_extractor -e wordlist -S write_feature_sqlite3=NO
-o jo_nosql jo-2009-12-11-001.E01
bulk_extractor version: 1.6.0-dev
Hostname: beefy2
Input file: jo-2009-12-11-001.E01
Output directory: jo_nosql
Disk Size: 15382241280
Threads: 16
0:35:39 Offset 67MB (0.44%) Done in 0:05:42 at 00:41:21
0:35:40 Offset 150MB (0.98%) Done in 0:03:45 at 00:39:25
0:35:41 Offset 234MB (1.53%) Done in 0:03:16 at 00:38:57
0:35:44 Offset 318MB (2.07%) Done in 0:04:49 at 00:40:33
0:35:44 Offset 402MB (2.62%) Done in 0:04:09 at 00:39:53
...
```



```
6. ec2-user@beefy2:~ (ssh)
top - 00:39:57 up 3:44, 3 users, load average: 12.32, 6.13, 2.54
Tasks: 195 total, 1 running, 194 sleeping, 0 stopped, 0 zombie
Cpu(s): 99.0%us, 0.3%sy, 0.0%ni, 0.5%id, 0.1%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 64471744k total, 14419864k used, 50051880k free, 33020k buffers
Swap: 0k total, 0k used, 0k free, 13089796k cached

  PID USER      PR  NI  VIRT  RES  SHR  S  %CPU  %MEM    TIME+  COMMAND
 19892 ec2-user   20   0 1716m 939m 10m  S 1587.8  1.5   42:29.42 bulk_extractor
 2686 root        0 -20    0    0    0  S   1.3   0.0    0:03.51 kworker/5:1H
   680 root       20   0    0    0    0  S   0.3   0.0    0:02.86 kworker/5:1
     1 root       20   0 19644 2448 2116  S   0.0   0.0    0:01.82 init
     2 root       20   0    0    0    0  S   0.0   0.0    0:00.00 kthreadd
     3 root       20   0    0    0    0  S   0.0   0.0    0:00.00 ksoftirqd/0
     4 root       20   0    0    0    0  S   0.0   0.0    0:00.00 kworker/0:0
     5 root        0 -20    0    0    0  S   0.0   0.0    0:00.00 kworker/0:0H
     7 root       20   0    0    0    0  S   0.0   0.0    0:00.28 rcu_sched
     8 root       20   0    0    0    0  S   0.0   0.0    0:00.00 rcu_bh
     9 root       RT    0    0    0    0  S   0.0   0.0    0:00.00 migration/0
    10 root        0 -20    0    0    0  S   0.0   0.0    0:00.00 lru-add-drain
    11 root       20   0    0    0    0  S   0.0   0.0    0:00.00 cpuhp/0
    12 root       20   0    0    0    0  S   0.0   0.0    0:00.00 cpuhp/1
    13 root       RT    0    0    0    0  S   0.0   0.0    0:00.00 migration/1
    14 root       20   0    0    0    0  S   0.0   0.0    0:00.00 ksoftirqd/1
    16 root        0 -20    0    0    0  S   0.0   0.0    0:00.00 kworker/1:0H
    17 root       20   0    0    0    0  S   0.0   0.0    0:00.00 cpuhp/2
```

Time elapsed waiting for 1 thread to finish:

18 sec (timeout in 59 min42 sec.)

Thread 14: Processing 11475615744

All Threads Finished!

Producer time spent waiting: 198.711 sec.

Average consumer time spent waiting: 31.1094 sec.

** bulk_extractor is probably CPU bound. **

** Run on a computer with more cores **

** to get better performance. **

MD5 of Disk Image: a6c44b7387a67333b8566955dcad6f50

Phase 2. Shutting down scanners

Phase 3. Uniquifying and recombining wordlist

Phase 3. Creating Histograms

Elapsed time: 416.436 sec.

Total MB processed: 15382

Overall performance: 36.9378 MBytes/sec (2.30861 MBytes/sec/thread)

Total email features found: 8480

Try same run with SQL output as well

```
[ec2-user@beefy2 work]$ bulk_extractor -e wordlist -S write_feature_sqlite3=YES
-o jo_sql jo-2009-12-11-001.E01
bulk_extractor version: 1.6.0-dev
Hostname: beefy2
Input file: jo-2009-12-11-001.E01
Output directory: jo_sql
Disk Size: 15382241280
Threads: 16
0:50:32 Offset 67MB (0.44%) Done in 0:52:05 at 01:42:37
0:50:33 Offset 150MB (0.98%) Done in 0:24:17 at 01:14:50
0:50:33 Offset 234MB (1.53%) Done in 0:16:23 at 01:06:56
0:50:36 Offset 318MB (2.07%) Done in 0:14:20 at 01:04:56
0:50:37 Offset 402MB (2.62%) Done in 0:11:37 at 01:02:14
```



```
6. ec2-user@beefy2:~ (ssh)
top - 00:50:55 up 3:55, 3 users, load average: 3.29, 1.99, 2.01
Tasks: 189 total, 1 running, 188 sleeping, 0 stopped, 0 zombie
Cpu(s): 89.6%us, 0.7%sy, 0.0%ni, 9.7%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 64471744k total, 15685696k used, 48786048k free, 34792k buffers
Swap: 0k total, 0k used, 0k free, 14142948k cached

  PID USER      PR  NI  VIRT  RES  SHR  S  %CPU  %MEM    TIME+  COMMAND
 25912 ec2-user  20   0 2013m 1.1g 10m  S   1443.6   1.8   3:22.97 bulk_extractor
     1 root      20   0 19644 2448 2116  S    0.0   0.0    0:01.82 init
     2 root      20   0     0     0     0  S    0.0   0.0    0:00.00 kthreadd
     3 root      20   0     0     0     0  S    0.0   0.0    0:00.00 ksoftirqd/0
     4 root      20   0     0     0     0  S    0.0   0.0    0:00.00 kworker/0:0
     5 root       0 -20     0     0     0  S    0.0   0.0    0:00.00 kworker/0:0H
     7 root      20   0     0     0     0  S    0.0   0.0    0:00.32 rcu_sched
     8 root      20   0     0     0     0  S    0.0   0.0    0:00.00 rcu_bh
     9 root      RT    0     0     0     0  S    0.0   0.0    0:00.00 migration/0
    10 root       0 -20     0     0     0  S    0.0   0.0    0:00.00 lru-add-drain
    11 root      20   0     0     0     0  S    0.0   0.0    0:00.00 cpuhp/0
    12 root      20   0     0     0     0  S    0.0   0.0    0:00.00 cpuhp/1
    13 root      RT    0     0     0     0  S    0.0   0.0    0:00.00 migration/1
    14 root      20   0     0     0     0  S    0.0   0.0    0:00.00 ksoftirqd/1
    16 root       0 -20     0     0     0  S    0.0   0.0    0:00.00 kworker/1:0H
    17 root      20   0     0     0     0  S    0.0   0.0    0:00.00 cpuhp/2
    18 root      RT    0     0     0     0  S    0.0   0.0    0:00.00 migration/2
    19 root      20   0     0     0     0  S    0.0   0.0    0:00.00 ksoftirqd/2
```


The SQLite3 output has 25% performance penalty when combined with with feature file output .

All Threads Finished!

Producer time spent waiting: 260.394 sec.

Average consumer time spent waiting: 28.6439 sec.

** bulk_extractor is probably CPU bound. **

** Run on a computer with more cores **

** to get better performance. **

MD5 of Disk Image: a6c44b7387a67333b8566955dcad6f50

Phase 2. Shutting down scanners

Phase 3. Uniquifying and recombining wordlist

Phase 3. Creating Histograms

Elapsed time: 567.171 sec.

Total MB processed: 15382

Overall performance: 27.121 MBytes/sec (1.69506 MBytes/sec/thread)

Total email features found: 8480

MD5 of Disk Image: a6c44b7387a67333b8566955dcad6f50

Phase 2. Shutting down scanners

Phase 3. Uniquifying and recombining wordlist

Phase 3. Creating Histograms

Elapsed time: 416.436 sec.

Total MB processed: 15382

Overall performance: 36.9378 MBytes/sec (2.30861 MBytes/sec/thread)

Total email features found: 8480

Class demonstration: show the results of the runs
